AI: Background, History and Future Opportunities

Raj Reddy
Carnegie Mellon University
Pittsburgh, PA 15213
Jul 20, 2018
Speech at SSN ACM Conference, Jul 20, 2018

Hype and Misinformation About AI

- Lot of Hype and Misinformation
- Incorrect Mistaken Assumptions
  - AI will NOT Replace Humans
  - AI will NOT Kill Us
  - AI will NOT Enslave Us
  - AI will Lead to Loss of Some Jobs
    - But More New Jobs will Ultimately be Created

Background

What is AI?

- AI is an attempt to automate tasks that are usually thought to be uniquely Human
  - Requiring Intelligence, Intuition, Creativity, Innovation, Emotion, Empathy
  - Usually Human coding using Heuristics, Rules, and
  - Statistical Models - HMMs
  - Non Sequential Algorithms
- Early Attempts
  - Attempts to Solve Problem not Expressible as Classical Algorithms
  - Proving Theorems, Playing Chess
  - Understand Language, Speech, Images
  - Create Robots that Sense, Think and Act
- Later Systems attempted
  - Compose music, Painting
  - Automate trading in Stock Market
- Usually leads to Imperfect Solutions
  - OCR error rates 99%
  - 2010 Flash Crash: Dow Lost 1000 points in 36 minutes

AI Principles: Use of Knowledge to Solve Problems

- An Intelligent System must
  - Learn from Experience
  - Use Vast Amount of Knowledge
  - Tolerate Error and Ambiguity
  - Respond in Real Time
  - Communicate with Humans using Natural Language
- Search Compensates for Lack of Knowledge
  - Puzzles
- Knowledge Compensates for Lack of Search
  - $F=ma$; $E = mc^2$
- Traditional Sources of Knowledge
  - Formal Knowledge: Learnt in Schools and Universities
    - Books, Manuals,
  - Informal Knowledge: Heuristics provided by People
- Initially Everything is Informal and Approximate Knowledge
  - As We Formulate Mathematical Formulas and Theories, Informal Knowledge Becomes Formal, Precise, Predictable and Repeatable
  - Invention of Zero and Algorithms for Addition and Multiplication took Five Hundred Years

History: 20th Century AI

Techniques and Systems of AI in 20th Century

- Human Encoding of Knowledge
  - Expert Systems
  - Knowledge Based Systems
  - Rule Based Systems
- Systems That Do Tasks that Require Intelligence
  - Play Chess
  - Prove Theorems
  - Discover Molecular Structure
- Systems That Do Tasks That Humans Do Effortlessly
  - Speak and Hear: Speech Understanding and Dialog
  - See: Computer Vision and Image Understanding
  - Use Language: Ambiguous and Non-grammatical Language
  - Drive a Car

Major Breakthroughs in AI of the 20th Century

- Enabled by Brute-force, Heuristics, Human Coding of Rules and Knowledge, and Simple Machine Learning (Pattern Recognition)
  - World Champion Chess Machine
    - IBM Deep Blue
  - Mathematical Discovery
    - Proof Checkers
  - Accident Avoiding Car
    - CMU: No Hands Across America
  - Robotics
    - Manufacturing Automation
    - Disaster Rescue Robots
  - Speech Recognition Systems
    - Dictation Machine
  - Computer Vision and Image Processing
    - Medical Image Processing
  - Expert Systems
    - Rule Based Systems
    - Knowledge Based Systems

Recent History: Advances in 21st Century

Recent Advances in AI:
Discover and Use Data Driven Knowledge Sources

- Paradigm Shift in Science
  - First 3 Paradigms: Experiment, Theory, Simulation
    - Rutherford, Bohr, Oppenheimer
  - 4th Paradigm: Data Driven Science
- Create Next Generation AI systems
  - Data Driven AI systems
  - To Solve Previously Unsolved Problems
- AI 2.0: Using Previously Unavailable Sources of Data
  - Knowledge from Big Data:
    - Data Driven Learning of Models and Algorithms
  - Knowledge from Multiple (Cross) Media:
    - Social Media Intelligence Gathering
  - Knowledge from Crowd Intelligence:
    - Global Brain: from Individual Intelligence to Collective Intelligence
  - Knowledge from Augmented Intelligence:
    - Human-Machine Hybrid Intelligence for Collaborative Problem Solving
  - Knowledge from Unmanned Autonomous Vehicles:
    - Intelligence from Collaborating Teams of Robots
- Automatic Discovery of New Knowledge
  - Machine Learning using Big Data
    - Deep Learning

Recent Breakthroughs in AI

- Enabled by Big Data and Machine Learning
  - Language Translation
    - Google Translate: Any Language to Any Language
  - Speech to Speech Dialog
    - Siri, Cortana, Alexa
  - Autonomous Vehicles
    - CMU, Stanford, Google, Tesla
  - Deep Question Answering
    - IBM’s Watson
  - Robo Soccer
  - World Champion Poker
    - CMU Libratus
      - No Limit Texas Hold’em Poker

Near Term Societal Impact of AI

Existing AI Technology Can be Used to Empower the 3 Billion People at The Bottom of the Pyramid

Bottom of the Pyramid: 3 Billion People

- 3 Billion People with Incomes of less than $3 a day
- The Bottom of The Pyramid is The Largest, But Poorest Socio-economic Group.
  - Globally that is the 3 billion people who live on less than say $2.50 per day.
- Most of Them Are Also Semi-literate, i.e., Cannot Read, Write and/or Understand Any Language
  - Cannot use Keyboard or Touch based Computing Apps
- For a Semi-literate Person, the Only Acceptable Mode of Communication is Speech
  - Voice Computing a la Amazon Echo and Enhancements is the Key
- Personal Assistants that Require Only Speech based Interaction is Essential for Such Populations

Applications for the Bottom of the Pyramid

- Voice Computing (No Keyboard or Touch) Can Help The Semi-literate to Read Newspapers, Watch Foreign Language Movies, Listen to Khan Academy Lectures, Vote Online and Order Groceries Online
  - A Mobile App for Entertainment and Education
    - Dynamic Real-time Translation of a Video Dialog from English to Telugu
    - Text to Speech App for Newspaper Reading Assistant
  - Enabling Digital Democracy
    - Vote Online (Authentication, Authorization and Audit)
  - Ecommerce and eBanking
    - Voice Authentication, Authorization and Audit
  - Learning Without a Teacher
    - Tutor for Listening and Speaking English
- Illiterate Populations Will be the Biggest Source of Customers for Speech Based Apps in the Future

Typical App for the Semi-literate: NextGen Alexa

- An Intelligent Agent That Anticipates What You Want To Do And Helps You To Do It Using Local Language and Clarification Dialog
  - Entertainment and Education: Streaming Video Translation
    - Alexa play Hamlet (BBC Shakespeare)
  - Reading Newspapers: Text to Speech Translation and/or Synthesis
    - Alexa read China Daily
  - Buying and Selling: Voice Dialog Management
    - Alexa order usual brand Rice, Meat and Vegetables
  - Communication: Voice and/or Video Email, Chat
    - Alexa call my Grandson in Shanghai
  - Banking: Monitor Bank account, Pay Bills
    - Alexa charge my mobile device with 1000 rupees
  - Online Voting
    - Voice Dialog to enable the Authorization, Authentication and Audit

Technology Exists to Create Voice-Only Apps

- Speech to Speech Exists (Microsoft, Facebook and Others)
  - BTW, Current Implementations are Based on Incorrect Business Assumptions
    - Available only for Commercial Languages
    - Will Never Result in Killer Apps
  - English to Chinese Speech to Speech Translation Demonstrated in 2012
    - Text Based “Translate” App of 2016 has to become Speech Based
  - Languages Supported Based on Commercial Viability
  - Not Need based
  - Unlikely to result in Killer App
- Apps Tailored to Semi-literate Populations Will Become Killer Apps
  - 1 Minute Learning Time; Two clicks; and Spoken Dialog
- All Such Apps Will Require Speech Recognition, Spoken Dialog, and Speech to Speech Translation (No Keyboard or Touch)
  - Speech to Speech Translation
    - Entertainment (Movies) and Education (Khan Academy)
    - Translate Live Dialog
  - QA Dialog (Siri and Cortana)
    - eCommerce and eBanking
    - English Language Learning - Detect Pronunciation Errors

AI: Future Opportunities

AI : Near term Future (2 to 3 Years)

Cognition Amplifiers That Enhance Human Capabilities

Do Tasks Faster and with Less Effort

Cognition Amplifiers in Service of Society

- A Cognition Amplifier (COG) is an Intelligent Agent that anticipates what you want to do and helps you to do it with less effort
- A Cognition Amplifier (COG) is a
  - Personal
  - Enduring
  - Autonomic Intelligent Agent that
  - Anticipates what you want to do and does it
- Cognition Amplifiers Enhance Human Capabilities
  - Do Tasks Faster and with Less Effort

Examples of COGs in Service of Society

- COGs personalized and mass customized agents as part of Knowledge as a Service (KaaS) may be Used by Everyone on the Planet for tasks such as
  - Buying and selling: Transact with multiple providers
  - Email: Filter spam, understand and respond to actionable email
  - News: Based on topic preferences, novelty, collaborative filtering
  - Banking: Monitor bank account, Credit Cards, Pay Bills
  - Travel: Flights, hotel, schedule disruptions, cancellations
- Each Person May Have Thousands of Cogs as Personal Assistants

Cloud Based Architecture of Cognition Amplifiers

- Always On
  - App Instantiation in the Cloud
  - Activity Monitoring: iPhone Sensor data
- Always Working
  - Diagnosis and Repair
- Always Learning
  - Publish/Subscribe

AI: Longer Term Future (5 to 10 Years)

Guardian Angels That Enable Humans to Do Tasks They Cannot Do Today.

Super-Human AI?

Guardian Angels in Service of Society

- A Guardian Angel (GAT) is a
  - Personal
  - Enduring
  - Autonomic
  - Intelligent Agent
- That Discovers and Warns Humans About Unanticipated Events That Could Impact Safety, Security, and Happiness
- Guardian Angels Enable Humans to Do Tasks They Cannot Now Do.
  - Super-Human AI?

GATs in Service of Society

- GATs are personalized and mass customized agents as part of KaaS for
  - Just-in-Time Warnings: Hurricanes, Earthquakes, Extreme Weather
  - Act as Coach in Health and Education Matters
  - Accident Alerts and Rerouting; Transport Strikes
  - Scarcity of Essential Resources: Food, Energy, Water etc.
- Assume Everyone on the Planet has Personalized Guardian Angels (GATs)

Cloud Based Architecture of Guardian Angels

- Always On
  - App Instantiation in the Cloud
  - Activity Monitoring: iPhone Sensor data
- Always Working
  - Diagnosis and Repair
- Always Learning
  - Publish/Subscribe

Potential Economic Impact of COGs and GATs

Emergence of Knowledge as a Service (KaaS) Industry

- Every person on the planet will be able to perform many daily habits more effectively using Guardian Angels
  - Daily habits (routines, activities) include a wide spectrum from routine tasks (such as banking and travel planning) to tasks too difficult for the user
- Over 80% of all human activity will done by Guardian Angels by 2020
- 7 Billion People Market Vs 2 Billion Today
- Ultimately Humans Could be 10 Times More Efficient and Effective
- Global GDP is $100 Trillion
- Even 10% improvement will lead to $10T additional wealth creation

AI: Predicting the Future

“Those who have knowledge, don't predict. Those who predict, don't have knowledge”, Lao Tzu

“Any Sufficiently Advanced Technology is Indistinguishable from Magic”, Arthur Clarke

Back to the Future
Future will be Remarkably Same as the Past

- 50 Years Ago: List of Research Projects at Stanford AI Labs (SAIL), 1963-1969
  - Robotics: Led to Vision and Robotics Industry
  - Mobile Robotics: Mars Rover and Stanford Cart
  - Image Understanding: Led to Vision and Robotics at CMU and Penn
  - Knowledge Engineering: Led to Expert Systems, Knowledge Engineering, Knowledge Based Systems Industry, and Early Applications of AI
  - Speech
    - Led to the DARPA Speech Understanding Project during the years 1971-76
    - Most influential branch of Speech Recognition Industry: Dragon Systems, Apple, Microsoft. Indirectly IBM and Bell Labs
  - Language Understanding: Question Asking and Dialog Modeling
  - Computer Music: Led to Yamaha adopting digital synthesis for consumer products
  - Other AI Projects: Chess, Symbolic Mathematics, Correctness of Programs, Theorem Proving, Logical AI, Common Sense
  - Computer Science: Time Sharing, LISP, DEC Clones, Graphical Editors, Pieces of Glass, Theory of Computation

AI: Some Unsolved AI Problems

- Integrated Intelligent Systems That
  - Learn from Experience
  - Use Knowledge
  - Communicate using Speech and Language
  - Operate in real time
  - Tolerate Error and Ambiguity
- Imprecise AI Tasks
  - Unlike Chess and Go, Concepts such as Intuition, Creativity, Innovation, Emotion, Empathy and Consciousness are Imprecise
    - Can a Computer laugh at a Joke?
    - Can a Computer get Angry?
- Not All Learning is Deep Learning
  - Learning from Examples
  - Learning by Doing and from Experience
  - Learning from Teachers
  - Learning from Sparse Data
  - Learning by Analogy

In Conclusion…

- Many of the AI Tasks We are Working on were Started 6 Decades Ago.
  - Progress Has Been Slow and Problematic
  - Even Deep Learning Started 35 Years ago
    - Geoff Hinton Invented the Backprop Algorithm at CMU 1985
- AI Research will Continue to Make Incremental Progress on All the Tasks from the Past.
  - Industry on New Applications of Image Processing, Speech and Robotics
  - Research Labs on Previously Unexplored Tasks such as Anger, Pain and Humor
- The Main Difference is that, We Now Have A Million Times More Computing Power, Million Times Memory and Million Times Bandwidth!
- As McCarthy said in 1980s, We May Need 1.7 Einsteins, 3 Maxwells and Funding of the magnitude of the 0.7 Manhattan Project to get to Human Level AI
- There is No Danger of AI Taking Over the World ☺