Tencent 腾讯 | 腾讯云

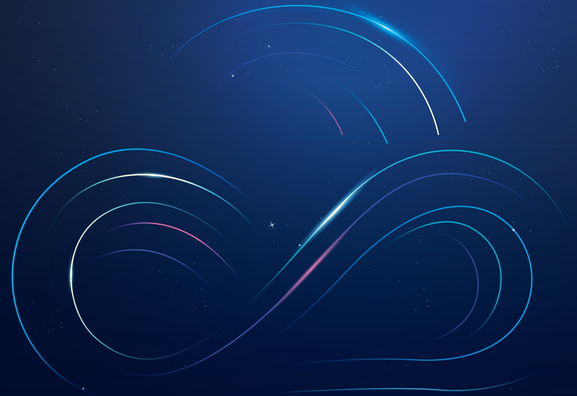

Artificial Intelligence: Perspectives and Challenges

Michael I. Jordan
University of California, Berkeley

June 21, 2017

Perspectives on AI

- The classical“I, robot”perspective
  cf. AI in the movies, interactive home robotics
- The“intelligence augmentation”(IA) perspective
  cf. search engines, recommendation systems, natural language translation
- The“intelligent infrastructure”(II) perspective
  cf. transportation, intelligent dwellings, urban planning

A Major Disconnect

- Problems studied from the“I, robot”perspective aren’ tnecessarily the same as those that arise in the IA or II perspectives
- But the“AI solutions”being deployed for the latter are often those developed in service of the former
- Some of the II problems that are far from solved:
  - error control for multiple, correlated decisions at massive scale
  - how to share data in competitive contexts?
  - cloud-edge interactions at massive scale
  - how to achieve fairness, quality and diversity?
  - robustness and security issues
- And there are“I, robot”problems that are far from solved

What Is Possible and Not Yet Possible?

- Computer vision
  Possible: labeling of objects in visual scenes
  Not Yet Possible: common-sense understanding of visual scenes

- Speech recognition
  Possible: speech-to-text and text-to-speech in a wide range of languages
  Not Yet Possible: common-sense understanding of auditory scenes

- Natural language processing
  Possible: minimally adequate translation and question-answering
  Not Yet Possible: semantic understanding, dialog

- Robotics
  Possible: industrial programmed robots
  Not Yet Possible: robots that interact with humans and can operate autonomously

What Will Be Possible in Ten Years?

- Everything that was listed as“Not Yet Possible”on the previous slide, at least in primitive form
  including“robots”such as self-driving cars and self-piloting air taxis
  including dialog systems that can stand in for humans in a range of service roles and secretarial roles
- But AI systems will still be very limited intelligences…

What Is Unlikely in our Lifetimes?

- It is unlikely that we will see AI systems that have the intellectual flexibility and creativity of humans
  - AI systems will have limited semantic understanding, and limited ability to cope with complex language (metaphor, irony, etc)
  - AI systems will have limited ability to reason abstractly, finding new abstractions on the fly
  - AI systems will have limited ability to plan in complex environments and adapt their plans on the fly
- They will seem like children who seem to know an amazing number of facts, and can have unusual insights, but nonetheless don't "understand"very much
- Certainly, we will not see"super-human AI"in the sense of AI that is clearly more intelligent than humans
  - I.e., I don't believe in the"singularity"in our lifetimes…

But Don’ t Computers Have Far More Processing Power than Humans?

- True, and indeed AI systems will have far more“patience”than humans to trawl through massive amounts of data and massive hypothesis spaces
  - so they will be able to have come up with unusual insights
  - but they won’t necessarily understand why an insight is an insight, and vet whether an insight is“real”or“useful” they will have limited abilities to form long chains of reasoning
  - they will have limited ability to explain themselves

- The AI systems that win at chess and Go are limited intelligences
  - note in particular that chess and Go are“fully observable”; there’s nothing hidden that the computer has to reason about
  - that’s not like real life

What Is Worth Worrying About?

- AI systems that seem smart but actually aren't
  and make decisions that create massive headaches and disasters
- The loss of jobs and incomes by large numbers of people
  especially in the service sectors, which have been the main hope for people displaced from classical industrial jobs
- Further disparities in wealth and lifestyle if the knowledge of how to build AI systems is not diffused
- The misuse of AI by people with bad intentions

Near-Term Challenges for ML

- Designing systems that can provide meaningful, calibrated notions of their uncertainty
- Designing systems that can explain their decisions
- Finding causes and performing causal reasoning
- Systems that pursue long-term goals, and actively collect data in service of those goals
- Achieving real-time performance goals

- Robustness in the face of unexpected situations
- Robustness in the face of adversaries
- Sharing data among individuals and organizations
- Protecting privacy and data ownership

Broader Challenges

- Semantics
  - real-world grounding
  - context
  - natural language taken seriously
- Cloud-edge interaction
  - cloud can aggregate (if privacy concerns are addressed)
  - and it can provide near-infinite processing
  - *but*, the cloud is further away from the real-world grounding and context that is precisely what's needed for real intelligence
- Uncertainty
  - representing uncertainty
  - communicating uncertainty
  - decision-making and planning under uncertainty
  - the social side of uncertainty

Conclusions

- ML (AI) has come of age
- But it is far from being a solid engineering discipline that can yield robust, scalable solutions to modern data-analytic problems
- There's no"magic"
- There are many hard problems involving uncertainty, inference, decision- making, robustness and scale that are far from being solved
- Not to mention economic, social and legal issues
- Great care will be needed to develop systems that work well and don't create headaches and disasters