Looking to the future on... Logic and Learning

Marta Kwiatkowska

Department of Computer Science, University of Oxford

Logic and Learning, ATI,12th Jan 2018

- <u>Logic</u> and/for/…

- Precision

- Hard/soft guarantees

- Soundness

- <u>(Deep) Learning</u>

- Performance

- No guarantees!

- Hidden assumptions…

- Logic   and/for/...

- Precision

- Hard/soft guarantees

- Soundness

- (Deep) Learning
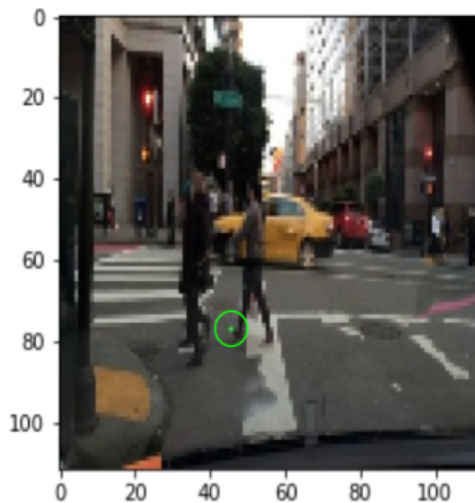
- Performance

- No guarantees!

- Hidden assumptions...

Different cultures!

- Complex scenarios
  - goals
  - perception
  - situation awareness
  - context (social, regulatory)

- Safety–critical, so guarantees needed

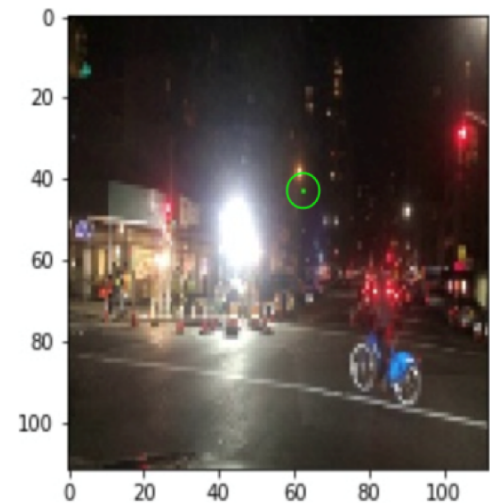- Should failure occur, accountability needs to be established

(a)   (b)   (c)

- Neural networks are unstable wrt adversarial perturbations
  - Nexar Traffic Light Challenge: Red light classified as green with 68%/95%/78% confidence after one pixel change.
  - CAV 2017, https://arxiv.org/abs/1610.06940
  - TACAS 2018, https://arxiv.org/abs/1710.07859

5

Summit on…

- ## Machine Learning Meets Formal Methods!

- Date: 13 July 2018
- Venue: University of Oxford
- Talks and panel discussion by academics and industrialists

- https://www.turing.ac.uk/events/summit-machine-learning-meet-formal-methods/

- http://www.floc2018.org/

Summit on ML Meets FM

Machine learning has revolutionised computer science and AI: deep neural networks have been shown to match human ability in various tasks and solutions based on machine learning are being deployed in real-world systems, from automation and self-driving cars to security and banking. Undoubtedly, the potential benefits of AI systems are immense and wide ranging. At the same time, recent accidents of machine learning systems have caught the public's attention, and as a result several researchers are beginning to question their safety. Traditionally, safety assurance methodologies are the realm of formal methods, understood broadly as the rigorous, mathematical underpinning of software and hardware systems. They are rooted in logic and reasoning, and aim to provide guarantees that the system is behaving correctly, which is necessary in safety-critical contexts. Such guarantees can be provided automatically for conventional software/hardware systems using verification technologies such as model checking or theorem proving. However, machine learning does not offer guarantees, and reasoning techniques necessary to justify safety of its autonomous decisions are in their infancy.

The summit on machine learning meets formal methods will bring together academic and industrial leaders who will discuss the benefits and risks of machine learning solutions. The overall aim is to identify promising future directions for research and innovation of interest to The Alan Turing Institute and UK research councils and government agencies, which will be summarised in a written report that will be made public.

8

Summit speakers

- **Machine learning and logic: Fast and slow thinking**
  Moshe Vardi (Rice University, USA)

- **Transparency and accountability for machine learning**
  Anupam Datta (Carnegie Mellon University, USA)

- **Programming from examples: PL meets ML**
  Sumit Gulwani (Microsoft, USA)

- **Towards neural program synthesis and code repair**
  Pushmeet Kohli (DeepMind, UK)

- **Provably beneficial Artificial Intelligence**
  Stuart Russell (University of California at Berkeley, USA)

- **Data centric engineering: A new concept?**
  Mark Girolami (Imperial College London and The Alan Turing Institute, UK)

- **The realities of applied AI**
  Alison Lowndes (NVIDIA, USA)

Summit speakers

- Rupak Majumdar (Max Planck Institute for Software Systems, Germany)
- Francesca Rossi (University of Padova, Italy)
- Christian Szegedy (Google, USA)
- Martin Vechev (ETH Zurich, Switzerland)

FLoC – Inspiring lecturers

Keynote
Shafi Goldwasser (MIT and Weizmann)
Georges Gonthier (INRIA Saclay)

Plenary
Byron Cook (Amazon and UCL)
Peter O'Hearn (Facebook and UCL)

Public lecture 10 July

Stuart Russell (UC Berkeley)
Logic and Probability

FLoC –Debate

Oxford Union-style debate on Ethics for Robotics

Luciano Floridi (Oxford/ATI)
Francesca Rossi (Padova)
Ben Kuipers (Michigan)
Jeannette Wing (Columbia)
Matthias Scheutz (Tufts)
Sandra Wachter (Oxford/ATI)

Moderated by Judy Wajcman (LSE)