Controls Over Artificial Agents

Ronald L. Chichester, JD, CISA

ISACA Greater Houston Chapter

September 17, 2015

Overview

- What is an Artificial Agent?
- How does it work?
- What are the risks?
- Examples
- Conclusions

What is an Artificial Agent?

- An software application that incorporates Artificial Intelligence (“AI”)
- What AI is not...
  - Standard algorithmic software
  - Human
- What AI is...
  - *Symbolic* processing of information
  - *Non-algorithmic* processing of information
  - Agents work autonomously once taught and instantiated

History of Artificial Intelligence

- 1940 - Conception by Shannon, Turing, Wheeler...
- 1950 - The Turing Test
- 1960 - AI established as a research field
- 1970 - First commercialization
- 1980 - Artificial Neural Networks
- 1990 - Intelligent Artificial Agents

So what is going on now?

More Processing Power

Less Expensive

More Numerous

Ubiquitous

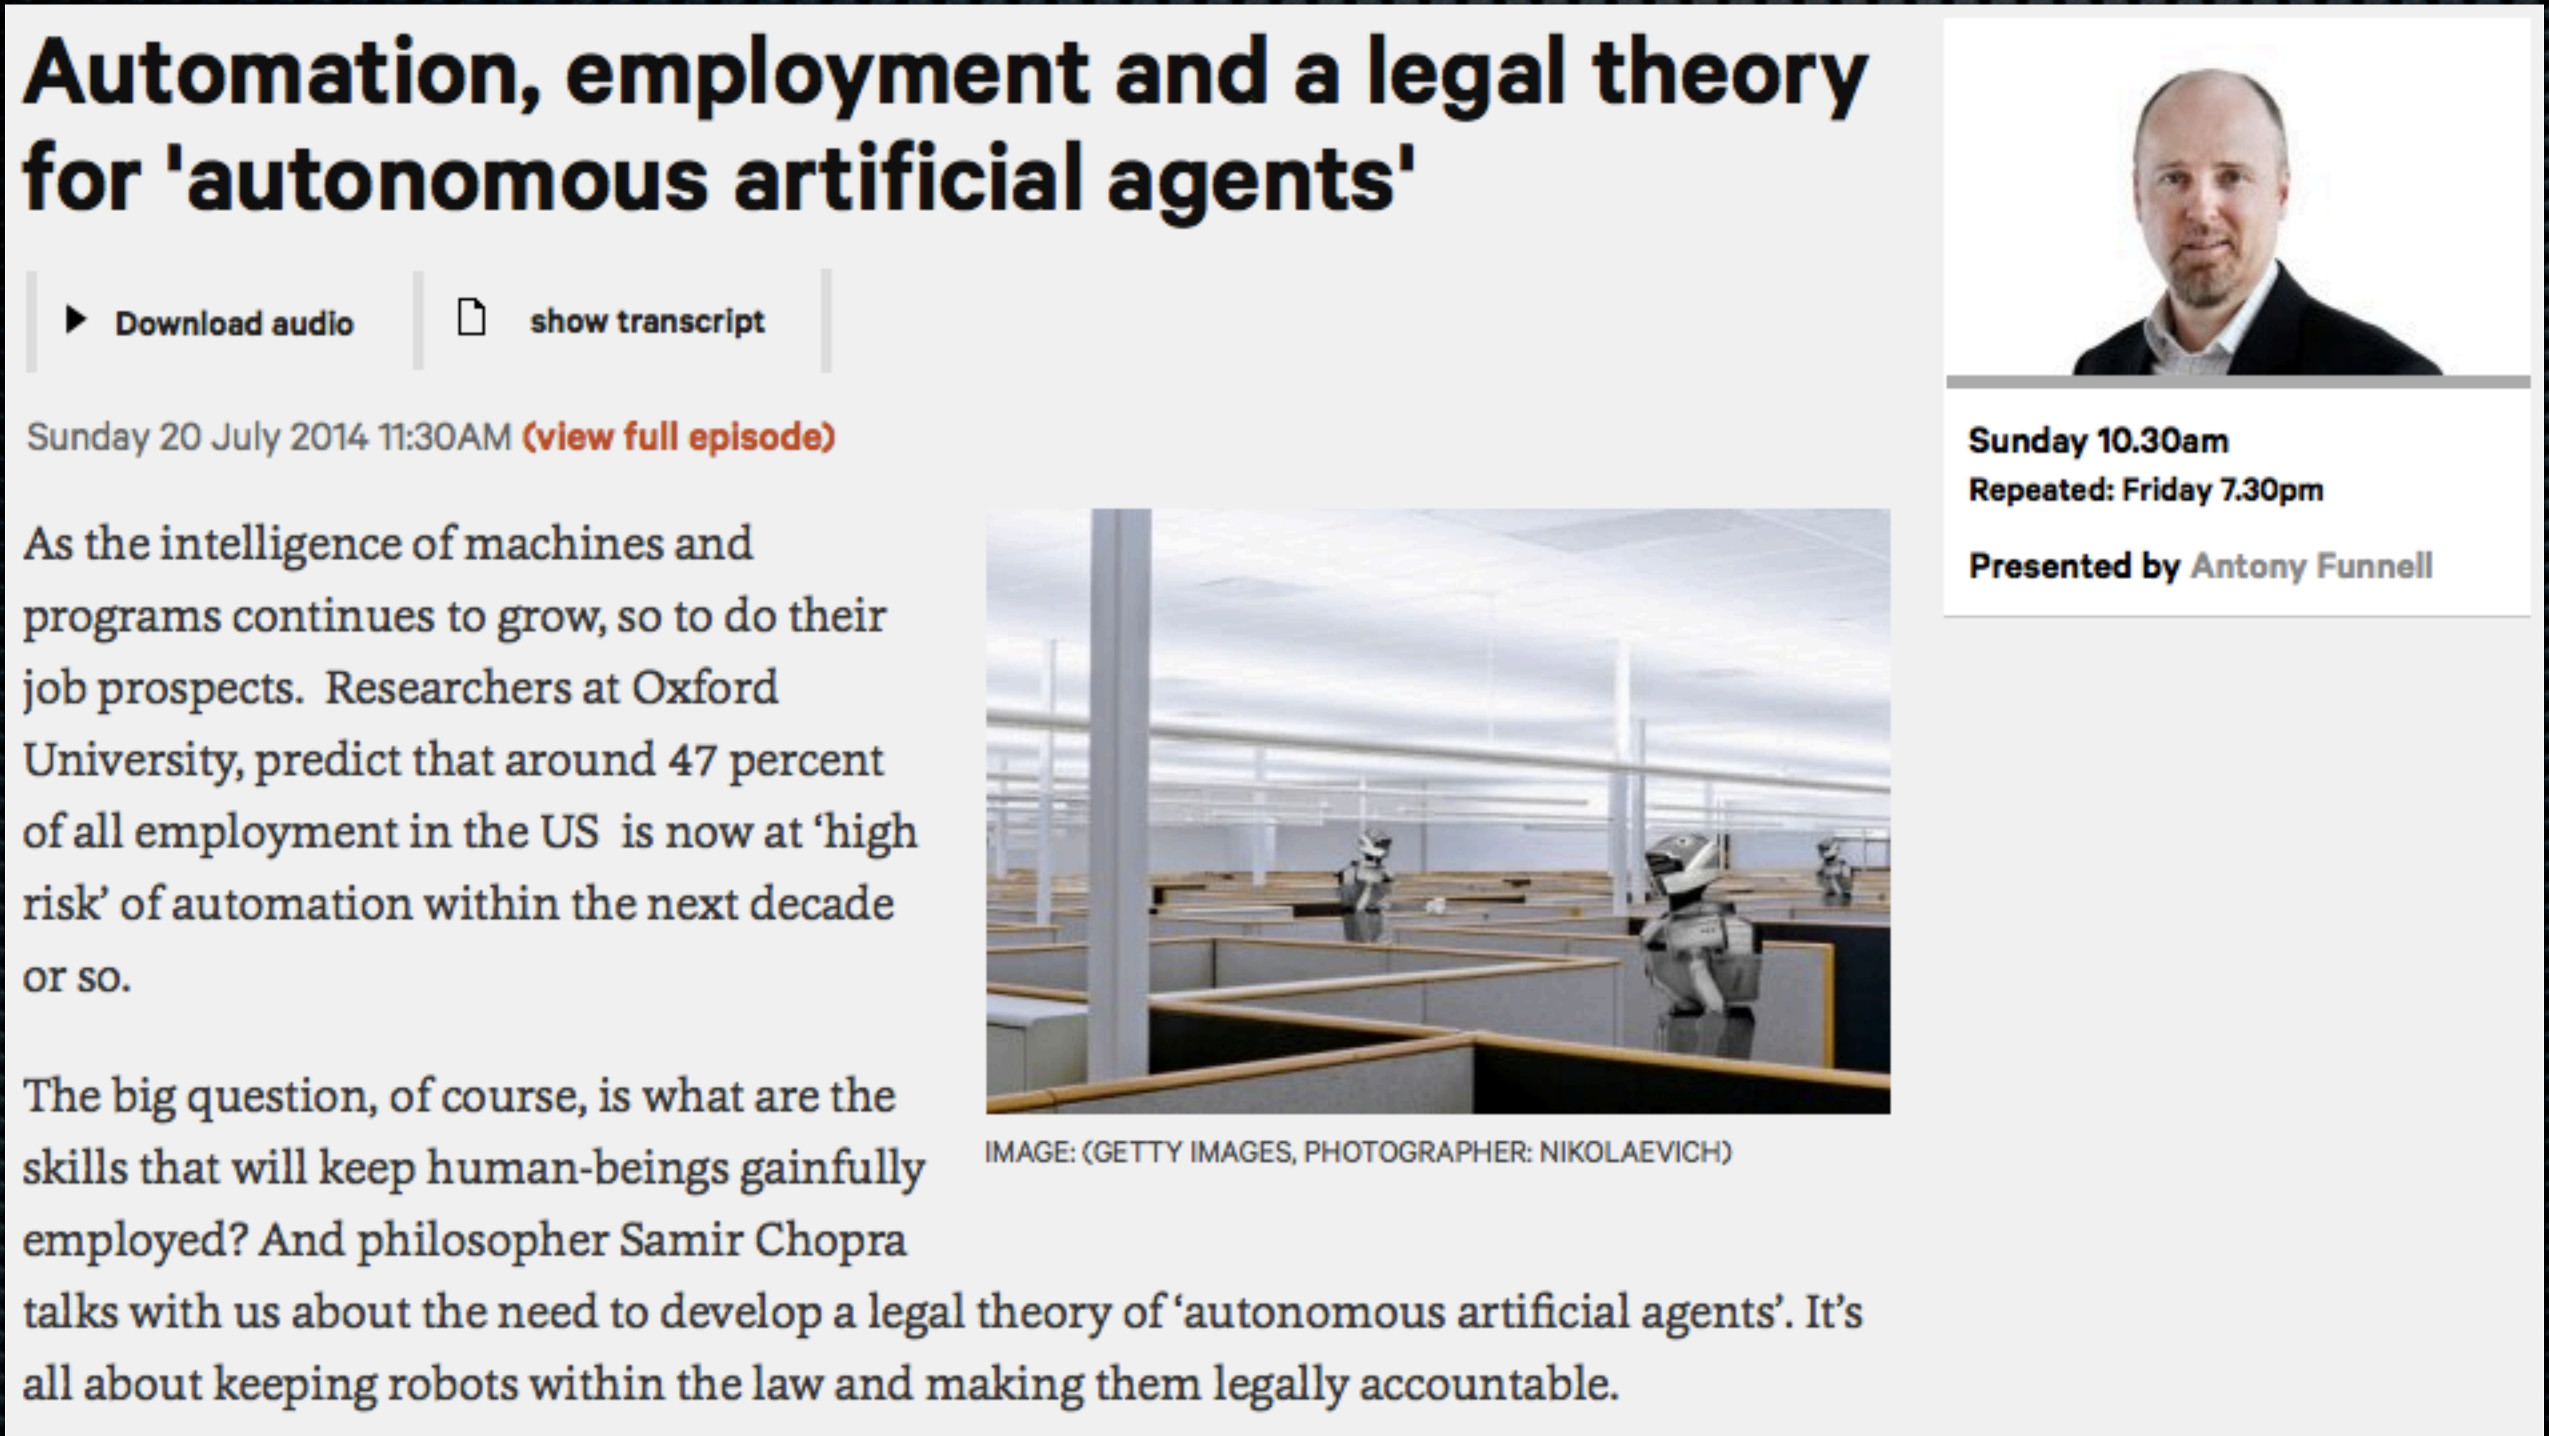

Automation, employment and a legal theory for 'autonomous artificial agents'

Sunday 20 July 2014 11:30AM

As the intelligence of machines and programs continues to grow, so to do their job prospects. Researchers at Oxford University, predict that around 47 percent of all employment in the US is now at 'high risk' of automation within the next decade or so.

IMAGE: (GETTY IMAGES, PHOTOGRAPHER: NIKOLAEVICH)

The big question, of course, is what are the skills that will keep human-beings gainfully employed? And philosopher Samir Chopra talks with us about the need to develop a legal theory of 'autonomous artificial agents'. It's all about keeping robots within the law and making them legally accountable.

Sunday 10.30am
Repeated: Friday 7.30pm

Presented by Antony Funnell

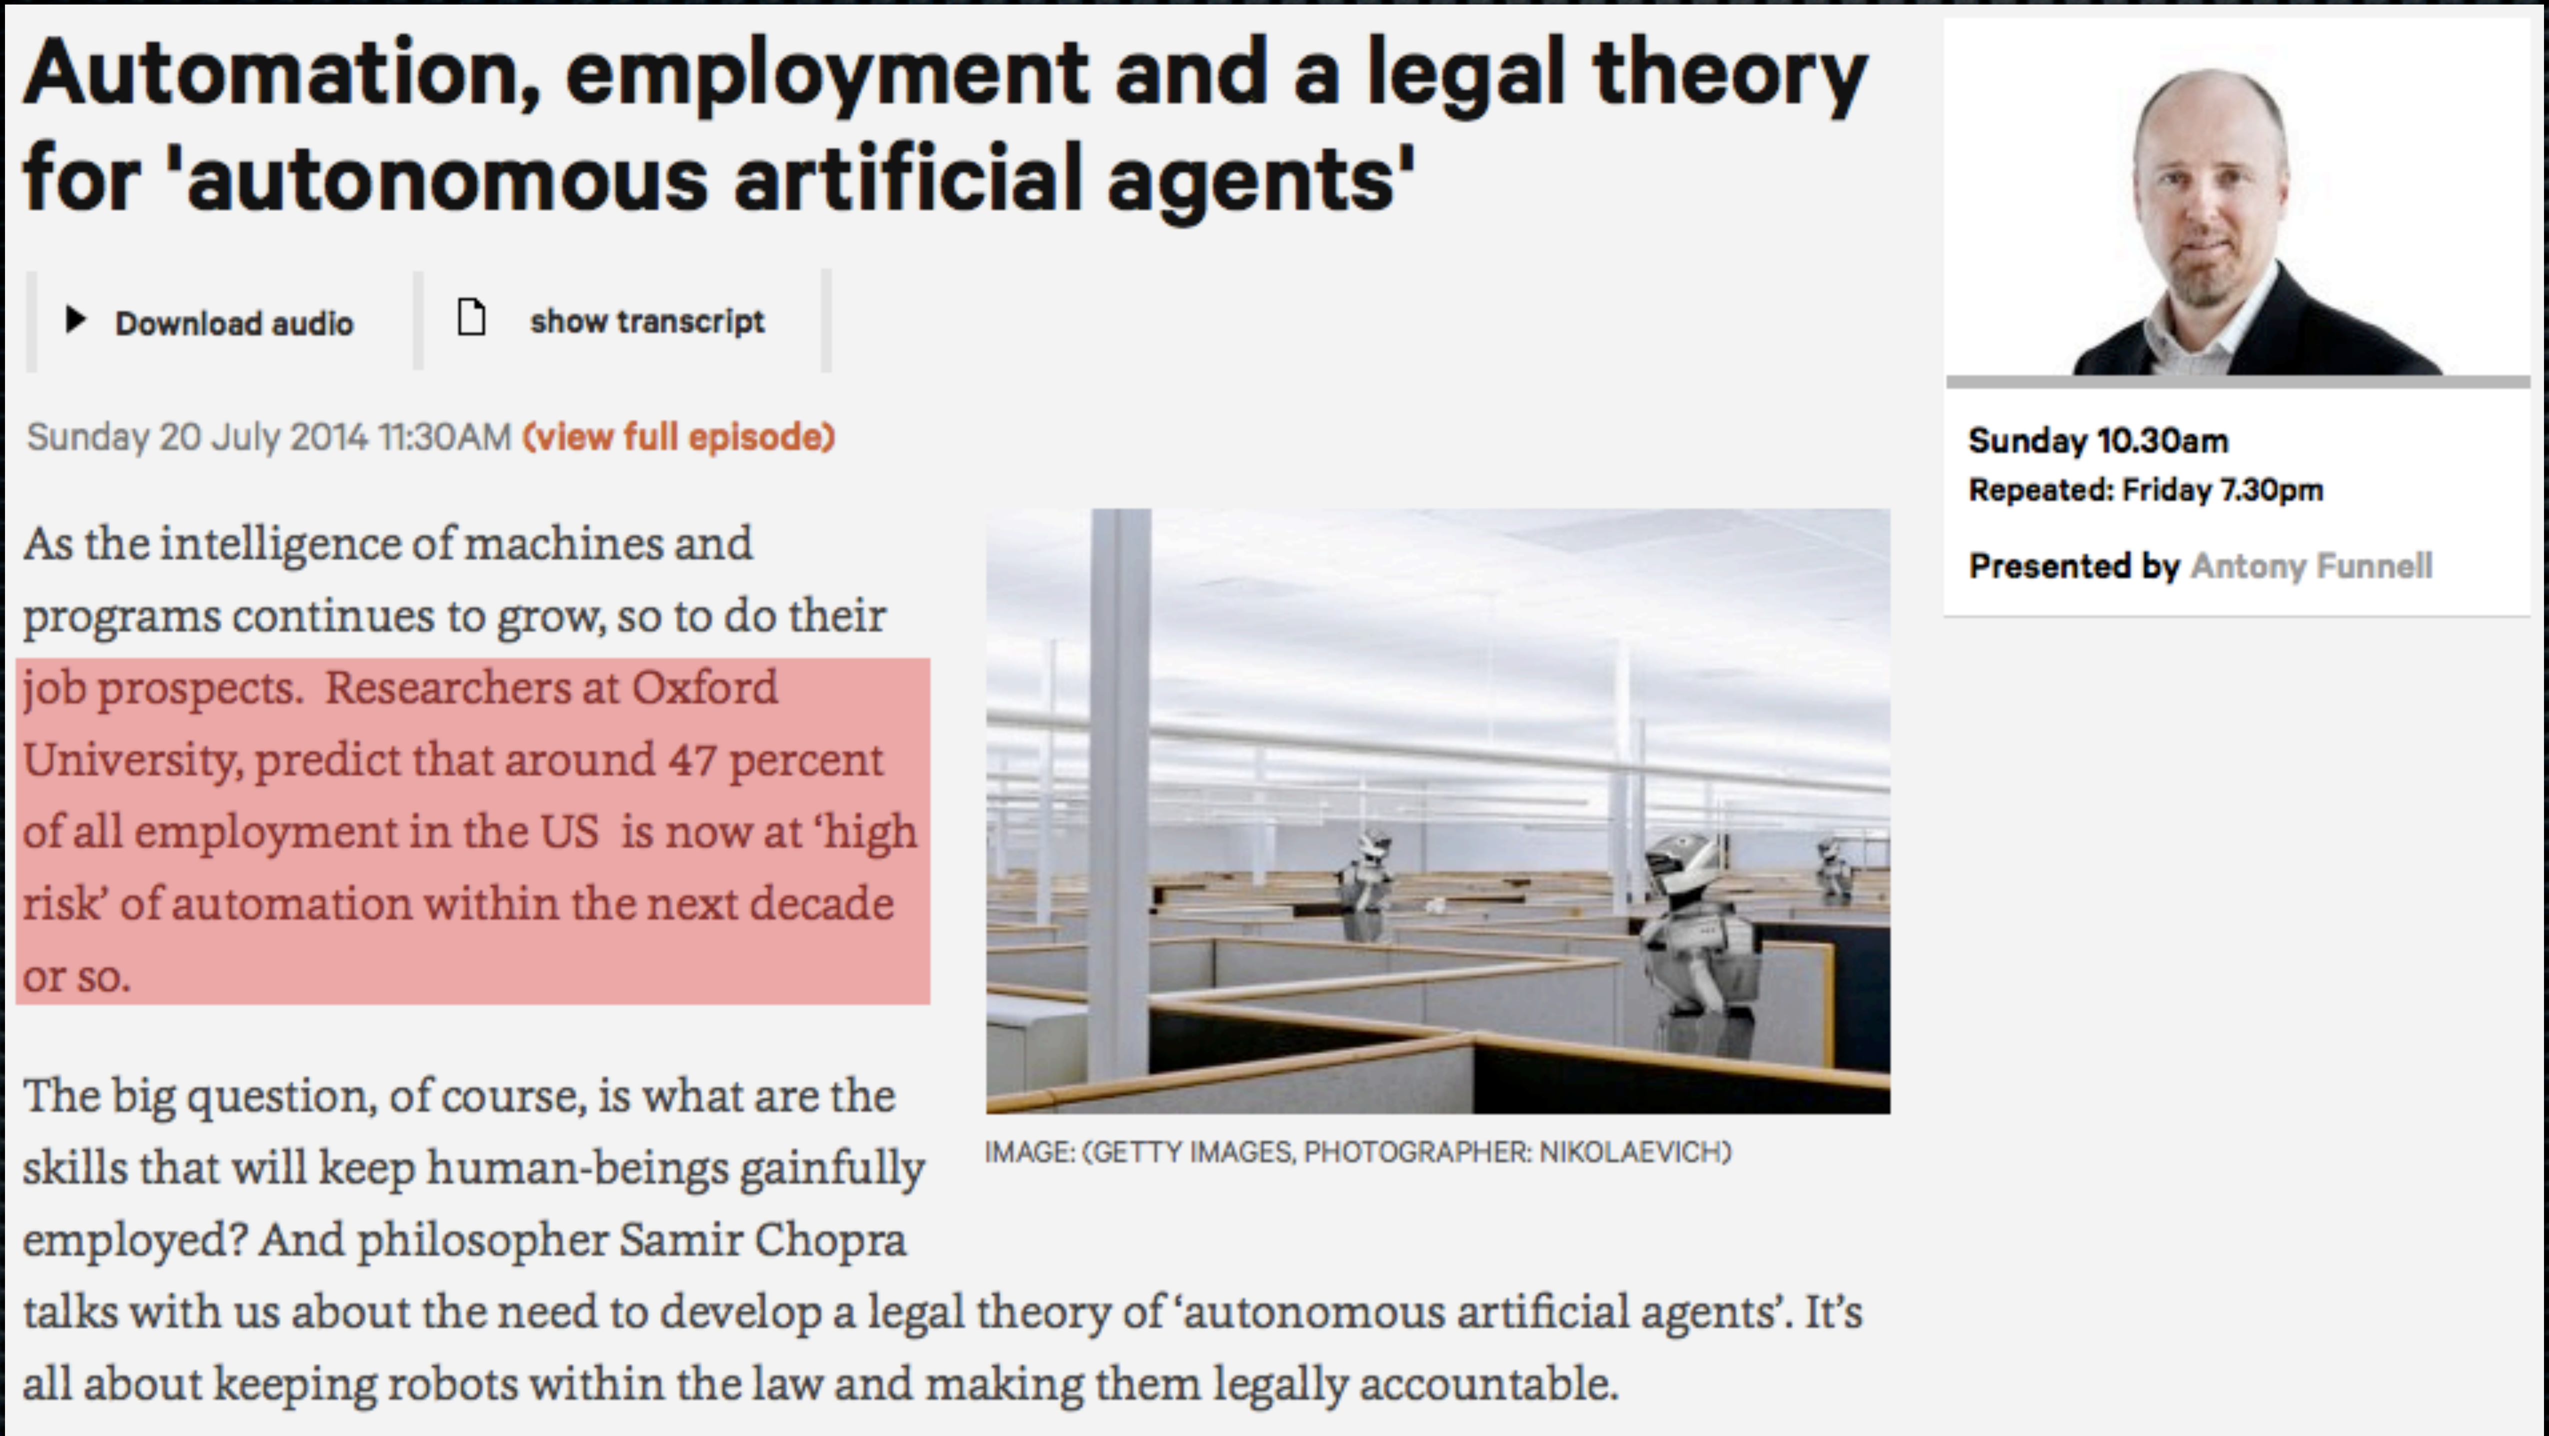

Automation, employment and a legal theory for 'autonomous artificial agents'

Sunday 20 July 2014 11:30AM (view full episode)

As the intelligence of machines and programs continues to grow, so to do their job prospects. Researchers at Oxford University, predict that around 47 percent of all employment in the US is now at 'high risk' of automation within the next decade or so.

IMAGE: (GETTY IMAGES, PHOTOGRAPHER: NIKOLAEVICH)

The big question, of course, is what are the skills that will keep human-beings gainfully employed? And philosopher Samir Chopra talks with us about the need to develop a legal theory of 'autonomous artificial agents'. It's all about keeping robots within the law and making them legally accountable.

Sunday 10.30am
Repeated: Friday 7.30pm

Presented by Antony Funnell

How Does the Agent Work?

An Example

- Expert System
    - Knowledge Base
        - Comes from a human expert
    - Inference Engine
        - Comes from a software vendor
    - Instantiator
        - (aka the Host)

Another Example

- Genetic Algorithms
  - The universal approximator
  - Needs an goal (defined by a human, typically)
    - Way to define the desired result
    - An initial population
    - Can run autonomously, but best with intervention
- Each element/step can constitute intellectual property

Yet Another Example

- Neural Networks
  - Mimics synopses in human brains
  - Enables learning
- Requires a network architect
- Requires teaching (by one or more entities)
- Requires a host for instantiation

Final Example

- DeepQA (by IBM)
  - The power behind Watson
  - Natural language processing
  - Once initially taught how to learn, it can teach itself
  - Instances of Watson sold commercially
    - High-priced knowledge jobs (lawyers, doctors...)

Applications

- Driverless cars
- Stock brokers
- Aircraft pilots
- Purchasing agents
- Personal assistants
- Mechanisms to shift risk away from the owner

< example >

(AdSense)

< /example >

What are the risks?

- Bad programming
- Flawed teaching
- Malevolent humans
  - Employees
  - Hackers
  - Customers

Risk will increase as artificial agents displace human workers

How Do We Control the Artificial Agent?

Before we answer that…

... you need to consider ...

.... the zeroth question.

Ownership

Can you exert controls?

Step 1:
What is the AI made up of?

Step 2:

Who owns which *part*?

Step 3:

Establish the *ability* to control

Step 4:
Establish controls

Normal mechanisms work to establish the ability to control

- Warranties & service level agreements
- Indemnification
- Modified IP clauses in employment agreements
- License
- Purchase

But they must all be (or have been) in place *before* you use the AI.

Why?

Because AI can generate IP

Absent a prior written agreement, the entities split *evenly*

A Few Potential Snags

- Patents, Copyrights (and even Trademarks)
- Federal Law of Agency
- Evolving state contract laws
- Fiduciary duties
- This area of the law is *not* settled
  - one set of rules for artificial personhood
  - another set of rules for computer algorithms
  - Controls for each differ

A Few Potential Snags

- While the law remains unsettled...
  - Annual review of the control scheme is *essential*
  - And that includes issues of *ownership*

A key aspect of artificial agents

is their ability to *shift* responsibility or sever the link to responsibility

and thus to *liability*

The control ***must*** dovetail with the direction of liability

That’s

Critical

Because a control can be evidence suggesting the *acceptance* of liability

< example >

Lenz v. Universal

U.S. Circuit Court of Appeals for the 9th Circuit
(September, 2015)

Lenz v. Universal

- Universal’s software trolled YouTube looking for infringement of copyrighted material
- The software found a video that had (in the background) a song from Prince
- Software automatically sent a takedown notice to YouTube
- Plaintiff alleged software did not consider fair use

“…the implementation of computer algorithms appears to be a valid and good faith middle ground for processing a plethora of content while still meeting the DMCA’s requirements to somehow consider fair use.”

Reliance on the computer program -- without human intervention -- was enough for potential liability for Universal

< /example >

< legal theories >

#1- Artificial agents as tools

- Agents are mere tools of their operators; or
- Mere means of communication
- All actions of the artificial agent are attributed to the *operator*
    - Not attributed to the host or programmer

#1- Artificial agents as tools

- Agents are mere tools of their operators; or
- Mere means of communication
- All actions of the artificial agent are attributed to the *operator*
    - Not attributed to the host or programmer

#2 - Unilateral Offer Doctrine

- Contract offered to anyone in the world
- Offer accepted by the acceptor’s conduct that was stipulated in the contract
  - “By using this website, you agree...”
- What happens when the artificial agent can craft the offer and the terms of the contract?
- Only invalid if agent acting erratically or unreasonably
- Potential liability depends upon the design of the agent

#3 - Contractual Intention

- Formal title: “Objective Theory of Contractual Intention”
- A contract is an obligation attached by the force of law to certain acts of the parties which accompany and represent a known intent
- A party’s assent is not necessary
  - Manifestation of intent (seen against a reasonable standard) is enough
  - “I put my agent on the cloud so that it could act as my purchasing agent ... and it so acted”

#3 - Contractual Intention

- Broad base for potential liability for Agent’s actions
  - The person who instantiated it on the cloud
  - The various software vendors
  - The person who taught it
  - The person who *didn’t* control it (but should have)

Conclusions

- Artificial agents are not all the same
- Different designs may incur disparate legal treatment
- Need to consider legal risks when choosing an agent
- Controls must be tailored to the agent design *and* the desired legal framework

Questions?

Ronald Chichester, JD, CISA

713.302.1679

Ron@TexasComputerLaw.com