

|



**Alan Turing
Ross Quillian
Edward Feigenbaum
Marvin Minsky
John McCarthy - 1956**



**“the science and engineering of
making intelligent machines”**

John McCarthy

“The study and design of intelligent agents, where an intelligent agent is a system that perceives its environment and takes actions that maximize its chances of success.”

Stuart Russel and

Peter Norvig

Machine Learning
Deep Learning (eg. c. vision)
Reinforcement learning
(decision making)
Natural language processing
Algorithmic game theory
Neuromorphic Computing

Computer vision

<https://cloud.google.com/vision/>

Computer vision reverse proof

<https://research.googleblog.com/2015/06/inceptionism-going-deeper-into-neural.html>

Play videogame
<http://www.theverge.com/2016/6/9/11893002/google-ai-deepmind-atari-montezumas-revenge>

SHARE

 SHARE
58


 TWEET

 PIN
1

 COMMENT
71

 EMAIL

IN A HUGE BREAKTHROUGH, GOOGLE'S AI BEATS A TOP PLAYER AT THE GAME OF GO



questões

Computer Chess - review



As enthralling as it is quirky, Andrew Bujalski's tale of man and machine moves from audaciously dull to singularly freaky

Peter Bradshaw

@PeterBradshaw1

Thursday 21 November
2013 16.45 EST



0 8



Mesmerising monochrome ... Computer Chess

Until now, my suspicion has been that the only thing of lasting value



GARRY
KASPAROV

DEEP
BLUE

KASPAROV vs
DEEP BLUE
the rematch



We proposed that the ultimate goal of the RoboCup Initiative to be stated as follows:

“By the middle of the 21st century, a team of fully autonomous humanoid robot soccer players shall win a soccer game, complying with the official rules of FIFA, against the winner of the most recent World Cup.

We propose that this goal will be one of the grand challenges shared by the robotics and AI communities for the next 50 years. This goal may sound overly ambitious given the state of the art of technology today.

Nevertheless, we believe it is important that such a long-range goal should be set and pursued. It took only 50 years from the Wright Brothers' first aircraft to the Apollo mission, to send a man to the moon and safely return him to the earth. It also took only 50 years, from the invention of the digital computer to Deep Blue, which beat the human world champion in chess. Building a humanoid soccer player requires an equally long period and extensive efforts of a broad range of researchers, and the goal will not be met in the near future.

Problema: recompensas errôneas

**[https://www.youtube.com/w
atch?v=tI0IHko8ySg](https://www.youtube.com/watch?v=tI0IHko8ySg)**

Problema: moralidade

**[http://moralmachine.mit.
edu/](http://moralmachine.mit.edu/)**

Self-Driving Mercedes-Benzes Will Prioritize Occupant Safety over Pedestrians

OCTOBER 7, 2016 AT 5:27 PM BY MICHAEL TAYLOR | ILLUSTRATION BY THE MANUFACTURER

 SHARE

 TWEET

 g+

 [Email](#)


 [Comment](#)

Research Vehicles

Select a Make 

Select a Model (optional) 

Research



The technology is new, but the moral conundrum isn't: A self-driving car identifies a group of children running into the road. There is no time to stop. To swerve around them would drive the car into a speeding truck on one side or over a cliff on the other, bringing

Sa

Problema: antropomorfismo

<https://www.youtube.com/watch?v=OpPbN-GueZQ>




An Open Letter

RESEARCH PRIORITIES FOR ROBUST AND BENEFICIAL ARTIFICIAL INTELLIGENCE

Artificial intelligence (AI) research has explored a variety of problems and approaches since its inception, but for the last 20 years or so has been focused on the problems surrounding the construction of intelligent agents – systems that perceive and act in some environment. In this context, “intelligence” is related to statistical and economic notions of rationality – colloquially, the ability to make good decisions, plans, or inferences. The adoption of probabilistic and decision-theoretic representations and statistical learning methods has led to a large degree of integration and cross-fertilization among AI, machine learning, statistics, control theory, neuroscience, and other fields. The establishment of shared theoretical frameworks, combined with the availability of data and processing power, has yielded remarkable successes in various component tasks such as speech recognition, image classification, autonomous vehicles, machine translation, legged locomotion, and question-answering systems.

As capabilities in these areas and others cross the threshold from laboratory research to economically valuable technologies, a virtuous cycle takes hold whereby even small improvements in performance are worth large sums of money, prompting greater investments in research. There is now a broad consensus that AI research is progressing steadily, and that its impact on society is likely to increase. The potential benefits are huge, since everything that civilization has to offer is a product of human intelligence; we cannot predict what we might achieve when this intelligence is magnified by the tools AI may provide, but the eradication of disease and poverty are not unfathomable. Because of the great potential of AI, it is important to research how to reap its benefits while avoiding potential pitfalls.

The progress in AI research makes it timely to focus research not only on making AI more capable, but also on maximizing the societal benefit of AI. Such considerations motivated the AAAI 2008-09 Presidential Panel on Long-Term AI Futures and other projects on AI impacts, and constitute a significant expansion of the field of AI itself, which up to now has focused largely on techniques that are neutral with respect to purpose. We recommend expanded research aimed at ensuring that increasingly





Google™

MUST READ **SECURITY 101: HERE'S HOW TO KEEP YOUR DATA PRIVATE, STEP BY STEP**

Brazilian government changes tech minister - again

Fourth minister in four years takes over in January; Aldo Rebelo's previous anti-innovation past sparks concerns within the local tech community



By [Angelica Mari](#) for Brazil Tech | December 30, 2014 -- 15:27 GMT (07:27 PST) | Topic: [Government](#)

Projects supported by the politician in recent years include a draft bill that would prevent petrol stations from using self service pumps, as well as electronic ticketing gates in buses.



Rebelo is also known for his nationalist stance and has supported projects such as the reduction of foreign words - such as "mouse", the computing device, or "tablet" - in the Portuguese language.

martin wolf



worst of times

THE FUTURE OF EMPLOYMENT: HOW SUSCEPTIBLE ARE JOBS TO COMPUTERISATION?*

Carl Benedikt Frey[†] and Michael A. Osborne[‡]

September 17, 2013


Abstract

We examine how susceptible jobs are to computerisation. To assess this, we begin by implementing a novel methodology to estimate the probability of computerisation for 702 detailed occupations, using a Gaussian process classifier. Based on these estimates, we examine expected impacts of future computerisation on US labour market outcomes, with the primary objective of analysing the number of jobs at risk and the relationship between an occupation's probability of computerisation, wages and educational attainment. According to our estimates, about 47 percent of total US employment is at risk. We further provide evidence that wages and educational attainment exhibit a strong negative relation-



pepper

- Management, Business, and Financial
- Computer, Engineering, and Science
- Education, Legal, Community Service, Arts, and Media
- Healthcare Practitioners and Technical
- Service
- Sales and Related
- Office and Administrative Support
- Farming, Fishing, and Forestry
- Construction and Extraction
- Installation, Maintenance, and Repair
- Production
- Transportation and Material Moving



GIGS WILL


GO TO

THE ROBOTS

Likelihood
human

Likelihood
automated


IN 10 TO 20 YEARS




NURSE




AIRLINE
PILOT




COMPUTER
PROGRAMMER



CONSTRUCTION



TAXI
DRIVER



TELEMARKETER

SOURCE: THE FUTURE OF EMPLOYMENT, BY CARL BENEDIKT FREY AND MICHAEL A. OSBORNE

PREPARING FOR THE FUTURE OF ARTIFICIAL INTELLIGENCE

Executive Office of the President
National Science and Technology Council
Committee on Technology

October 2016



best of times

produtividade agregada


1947-1973
3% ao ano (US)

2007

1,2% ao ano

2015-2016
0,6% ao ano

lowest historic churn
(Robert Atkinson
and John Wu)



aumentos de salário

**Carros autônomos vão levar no
mínimo 25 anos**

OECD diz que só 9% dos empregos estão ameaçados pela automação

**Erro: se um trabalho pode ser
automatizado ele será
automatizado**

ATMs: 1970
Mais de 400.000 (US)

**Número de caixas de banco subiu
entre 2000 e 2010 (James Basson)**

**Declínio projetado de 8% para
próxima década**

**De todas as 271 profissões do
censo de 1950, apenas uma
desapareceu
(ascensorista) (US)**

**Custo do trabalho no mundo
desenvolvido vs. mundo em
desenvolvimento**

“Robots and Jobs”

(Pascual Restrepo)

670.000 postos perdidos

China: 2.4 milhões

**“If you want to know what
happened to manufacturing after
2000, the answer is very clearly
not automation, it’s China”
(Dean Baker)**

**AI pode aumentar o PIB dos Eua
em 2%
(Accenture)**

**Custo do trabalho no mundo
desenvolvido vs. mundo em
desenvolvimento**

Singularity

Multiplicity

desafios continuam

big data vs. no data

Most engineers are white – and so are the faces the train software


A black researcher had to wear a white mask to test her own project.

BY TESS TOWNSEND | JAN 18, 2017, 11:45AM EST

 TWEET

 SHARE

 LINKEDIN



T



plataformas

An open-source software library for Machine Intelligence

[GET STARTED](#)

About TensorFlow

TensorFlow™ is an open source software library for numerical computation using data flow graphs. Nodes in the graph represent mathematical operations, while the graph edges represent the multidimensional data arrays (tensors) communicated between them. The flexible architecture allows you to deploy computation to one or more CPUs or GPUs in a desktop, server, or mobile device with a single API. TensorFlow was originally developed by researchers and engineers working on the Google Brain Team within Google's Machine Intelligence research organization for the purposes of conducting machine learning and deep neural networks research, but the system is general enough to be applicable in a wide variety of other domains as well.



Tensor

About OpenAI

OpenAI is a non-profit artificial intelligence research company. Our mission is to build safe AI, and ensure AI's benefits are as widely and evenly distributed as possible.

In the short term, we're building on recent advances in AI research and working towards the next set of breakthroughs. Please direct press inquiries to Jack Clark at press@openai.com.

Hiring

We're slowly building a tight-knit team and thoughtful, open culture. We believe that unreasonably great results can come from a small group working in concert. We seek to broadcast our work to the world, whether as papers, blog posts, code, open-source projects, talks, or tutorials.

algumas questões adicionais

não é uma coisa só mas várias

**limitações das instituições atuais
(ex. Responsabilidade civil)**

que é dono e controla dos dados

**proprietário v. open/
DMCA**

IA e o contrato social

**a democracia leva à solução
the solution (?)**

**aumentando o capital em face do
trabalho**



Davos 2017 - Artificial Intelligence



Singularity Lectures



Subscrever 16 284

10 144 visualizações



Fairness, Accountability, and Transparency in Machine Learning

Bringing together a growing community of researchers and practitioners concerned with fairness, accountability, and transparency in machine learning

The past few years have seen growing recognition that machine learning raises novel



Princípios de “accountability” para algoritmos e inteligência artificial

- 1. Responsabilidade**
- 2. Explicabilidade**
- 3. Acuracidade**
- 4. Auditabilidade**
- 5. Justiça (Fairness)**

1. Responsibility. “For any algorithmic system, there needs to be a person with the authority to deal with its adverse individual or societal effects in a timely fashion. This is not a statement about legal responsibility but, rather, a focus on avenues for redress, public dialogue, and internal authority for change”.

Who is responsible if users are harmed by this product?

What will the reporting process and process for recourse be?

Who will have the power to decide on necessary changes to the algorithmic system during design stage, pre-launch, and post-launch?

2. Explainability. “Any decisions produced by an algorithmic system should be explainable to the people affected by those decisions. These explanations must be accessible and understandable to the target audience; purely technical descriptions are not appropriate for the general public.”

Who are your end-users and stakeholders?

How much of your system / algorithm can you explain to your users and stakeholders?

How much of the data sources can you disclose?

3. Accuracy “The principle of accuracy suggests that sources of error and uncertainty throughout an algorithm and its data sources need to be identified, logged, and benchmarked.”

What sources of error do you have and how will you mitigate their effect?

How confident are the decisions output by your algorithmic system?

What are realistic worst case scenarios in terms of how errors might impact society, individuals, and stakeholders?

Have you evaluated the provenance and veracity of data and considered alternative data sources?

4. **Auditability** “The principle of auditability states that algorithms should be developed to enable third parties to probe and review the behavior of an algorithm... While there may be technical challenges in allowing public auditing while protecting proprietary information, private auditing (as in accounting) could provide some public assurance.”

Can you provide for public auditing (i.e. probing, understanding, reviewing of system behavior) or is there sensitive information that would necessitate auditing by a designated 3rd party?

How will you facilitate public or third-party auditing without opening the system to unwarranted manipulation?

5. **Fairness** “All algorithms making decisions about individuals should be evaluated for discriminatory effects. The results of the evaluation and the criteria used should be publicly released and explained.”

Are there particular groups which may be advantaged or disadvantaged, in the context in which you are deploying, by the algorithm / system you are building?

What is the potential damaging effect of uncertainty / errors to different groups?