Artificial Intelligence

Augustin Cosse.



Fall 2020

September 2, 2020

◆□▶ ◆□▶ ◆臣▶ ◆臣▶ 臣 のへぐ

Schedule

- Class and labs: Monday/Wednesday, 12:30pm-1:45pm.
- Recitation (Mandatory): Wednesday 2.15pm 3.45pm.
- Office hour : Monday 2.15pm 3.45pm.
- Location: NYU Paris, 57 Boulevard Saint-Germain, Room 401

- Combination between programming sessions (python) and lectures
- Final Exam: Midterm: 30%, Final : 30%
- Assignements throughout the semester: 30%
- Independent project: 10%

Course organization

- Notes + Sample exam questions will be posted soon on the course webpage
- Sample exam questions = help you with the study but not comprehensive
- If a section of the notes is not covered in class, you don't have to study it for the exam

< □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > <

 See http://www.augustincosse.com/teaching for details (select "AI > Fall 2020" semester)

Bootcamp (to be confirmed)

- General revision on Math/Programming
- Book the dates !
 - Friday Sept. 4 all day
 - Friday Sept. 11 morning
 - Saturday Sept. 12 Morning
- Detailed schedule
 - Friday Sept. 4, 9am-12pm: vector and complex numbers
 - Friday Sept. 4, 1pm-4pm: univariate calculus
 - Friday, Sept 11, 9am-11am: linear algebra
 - Friday, Sept 11, 11am-1pm: multivariate calculus
 - Friday, Sept 11, 2pm-4pm: probability and statistics
 - Saturday Sept 12, 9am-12pm: basics of Python / Numpy

• Bootcamp webpage to be opened soon..

Course website (I)

Artificial Intelligence (NYU Paris, Fall 2020)





Tentative schedule:

Legend: Lab sessions are in green, Homeworks are in red (right side of the table), dates related to the project are in orange.

Week #	date	Торіс	Assignements
Week 1	08/31, 09/02	General Intro. Critical and Philosophical perspectives, History,	Readings
Week 2	09/07, 09/09	Intelligent Agent, Search Methods (including informed/uninformed, A*,) Part I Lab 2	Assign. 1, Readings Solutions
Week 3	09/14,09/16	Intelligent Agents, Search Methods (including informed/uninformed, A*,) Part II Lab 3	Readings

General Organization (I): Website + Project

- All the material will be posted on the website (including notebooks, slides and coursenotes)
- Current Password : singularity2045
- One end of semester project: Implementation and/or reading of pioneering paper (see course website for some ideas)

General Organization (II): Reference Books

 Main reference for the course: Artificial Intelligence: A Modern Approach Livre de Peter Norvig et Stuart Russell (available online + physical copy at NYUP library)



 If you want to go deeper/complement : E.g. Artificial Intelligence, P. H. Winston, Introduction to Artificial Intelligence, Mariusz Flasinski, ... (see course website)

Tentative Program

Week 1	General Introduction, Intelligent Agents (RN 1-2)			
Week 2	Solving problems by searching, Breadth/Depth-first search,			
	A* search, (Part I) (W 4-6) (RN 3-4)			
Week 3	Solving problems by searching, Breadth/Depth-first search,			
	A* search, (Part II) (W 4-6) (RN 3-4)			
Week 4	Logic, propositional logic, First order logic			
	Acting logically (Part I) (W 13) (NJ 2-4) (RN 6,7,9,11)			
Week 5	Logic, propositional logic, First order logic			
	Acting logically (Part II) (W 13) (NJ 2-4) (RN 6,7,9,11)			
Week 6	Learning through Neural Networks, (W 22, 23, 24) (RN 19)(NJ 15)			
Week 7	Deep Learning, Alpha Go and Alpha Go Zero, link 1, link 2 (RN 20)			
Week 8	Unsupervised Learning and Reinforcement Learning, (RN 17, 20)(NJ 12)			
Week 9	Multiagent Framework + Guest lecture (e.g. Minds and Machines)			
Week 10	Learning under uncertainty, Probabilistic Reasoning Systems, Hidden Markov models,			
	Bayesian and Belief networks (RN 14, 15, 19) (NJ 6-11)			
Week 11	Evolutionary Computation, Genetic algorithms,			
	Swarm computation, Ant colonies (Part I) (W 25)(NJ 13-14) (RN 20.8)			
Week 12	Evolutionary Computation, Genetic algorithms,			
	Swarm computation, Ant colonies (Part II) (W 25) (NJ 13-14) (RN 20.8)			
Week 13	Vision and Language I, (W 26-29)(RN 22-24)(NJ 16)			
Week 14	Vision and Language II, (W 26-29) (RN 22-24) (NJ 16)			
Week 15	Conclusions, philosophical foundations, Present and Future of AI (RN 26-27)			
Week 16	Final Exams, Projects presentations			
Bonus (If Time)	Additional Applications (e.g. [1])			

Today

- General Introduction. Critical and Philosophical perspectives, History.
- No recitation. Take the time to install python anaconda
- Open a github account and make sure you can open and work with jupyter notebooks

・ロト・日本・モート モー うへぐ



WIRED

▲□▶ ▲圖▶ ★ 国▶ ★ 国▶ - 国 - のへで

artificial intelligence index 2019 annual report



Stanford

Human-Centered Artificial Intelligence



Fig. 1.1.

Annual Growth in Al papers on Scopus Source: Elsevier, 2019.



World Map of Academic-Corporate Collaboration: Total Number of AI papers



・ロト ・回ト ・ヨト ・ヨ



Total Volume and average annual per capita AI Journal Publications, 2015-2018

Source: MAG, 2019.



Al Patents



Fig. 1.13. Note: Percentage of patent citations to unknown country is 37.2%.

(日)、

500

ж

Women in AI -



Many Western European countries as well as Argentina, Canada, and Iran show relatively high presence of women in AI research.





Attendance at large conferences (1984-2019)

Source: Conference provided data.



Note: IJCAI occurred every other year till 2014. The missing year between 1984 and 2014 are interpolated as the mean between the two known conference attendance dates to provide a comparative view across conferences.



Number of paper titles at AI conferences mentioning Ethics keywords, 1969-2018 Source: Prates et al., 2019.





æ

Visual-Question Answering (VQA) -

The VQA challenge incorporates both computer vision and natural language understanding. The VQA challenge tests how well computers can jointly reason over these two distinct data distributions.



GLUE

Being able to analyze text is a crucial, multipurpose AI capability. In recent years, progress in natural language processing and natural language understanding has caused the AI community to develop new, harder tests for AI capabilities. In the language domain, a good example is GLUE, the General Language Understanding Evaluation benchmark.



Fig. 3.7.

▲ロト ▲圖 ト ▲ 画 ト ▲ 画 ト の Q @

SQuAD

One way to highlight recent progress in natural language processing is to examine performance on the Stanford Question Answering Dataset (SQuAD) challenge. SQuAD is a reading comprehension dataset, consisting of questions posed by crowdworkers on a set of Wikipedia articles. The answer to every question is a segment of text, or span, from the corresponding reading passage, or the question might be unanswerable. SQuAD1.1 is the SQuAD dataset and contains 100,000+ questionanswer pairs on 500+ articles. SQuAD2.0 combines the 100,000 questions in SQuAD1.1 with over 50,000 unanswerable questions written adversarially by crowdworkers to look similar to answerable ones.



Allen Institute for Artificial Intelligence: ARC Reasoning Challenge

Source: AI2 Leaderboard.



A new dataset of 7,787 genuine grade-school level, multiple-choice science questions, assembled to encourage research in advanced question-answering. The dataset is partitioned into a Challenge Set and an Easy Set, where the former contains only questions answered incorrectly by both a retrieval-based algorithm and a word co-occurrence algorithm. We are also including a corpus of over 14 million science sentences relevant to the task, and an implementation of three neural baseline models for this dataset. We pose ARC as a challenge to the community.

(日)、

Allen Institute for Al

Human-Level Performance Milestones

The inaugural 2017 AI Index report included a timeline of circumstances where AI reached or beat human-level performance. The list outlined game playing achievements, accurate medical diagnoses, and other general, but sophisticated, human tasks that AI performed at a human or superhuman level. This year, two new achievements are added to that list. It is important not to over-interpret these results. The tasks below are highly specific, and the achievements, while impressive, say nothing about the ability of the systems to generalize to other tasks.



Othello

In the 1980s Kai-Fu Lee and Sanjoy Mahajan developed BLL, a Bayesian learningbased system for playing the board game Othello. In 1989, the program won the US national tournament of computer players, and beat the highest ranked US player, Brian Rose, 56—8. In 1997, a program named Logistello won every game in a six game match against the reigning Othello world champion.

Checkers

In 1952, Arthur Samuels built a series of programs that played the game of checkers and improved via self-play. However, it was not until 1995 that a checkers-playing program, <u>Chinook</u>, beat the world champion.

Chess

Some computer scientists in the 1950s predicted that a computer would defeat the human chess champion by 1967, but it was not until 1997 that <u>IBM's DeepBlue</u> <u>system</u> beat chess champion Gary Kasparov. Today, chess programs running on smartphones can play at the grandmaster level.

Jeopardy!

In 2011, the IBM Watson computer system competed on the popular quiz show Jeopardy! against former winners Brad Rutter and Ken Jennings. Watson won the first place prize of \$1 million.

Atari Games

In 2015, a team at Google DeepMind used a reinforcement learning system to learn how to play 49 Atari games. The system was able to achieve human-level performance in a majority of the games (e.g., Breakout), though some are still significantly out of reach (e.g., Montezuma's Revenge).



Object Classification in ImageNet

In 2016, the error rate of automatic labeling of <u>ImageNet</u> declined from 28% in 2010 to less than 3%. Human performance is about 5%.

Go

In March of 2016, the AlphaGo system developed by the Google DeepMind team <u>beat Lee Sedol</u>, one of the world's greatest Go players, 4—1. DeepMind then released <u>AlphaGo Master</u>, which defeated the top ranked player, Ke Jie, in March of 2017. In October 2017, a Nature paper detailed yet another new version, <u>AlphaGo Zero</u>, which beat the original AlphaGo system 100—0.

Skin Cancer Classification

In a 2017 <u>Nature article</u>, Esteva et al. describe an Al system trained on a data set of 129,450 clinical images of 2,032 different diseases and compare its diagnostic performance against 21 board-certified dermatologists. They find the Al system capable of classifying skin cancer at a level of competence comparable to the dermatologists.

Speech Recognition on Switchboard

In 2017, <u>Microsoft</u> and <u>IBM</u> both achieved performance within close range of "human-parity" speech recognition in the limited Switchboard domain

Poker

In January 2017, a program from CMU called <u>Libratus</u> defeated four to human players in a tournament of 120,000 games of two-player, heads up, no-limit Texas Hold'em. In February 2017, a program from the University of Alberta called DeepStack played a group of 11 professional players more than 3,000 games each. DeepStack won enough poker games to prove the statistical significance of its skill over the professionals.

・ロト ・ 理 ト ・ ヨ ト ・ ヨ ト



2017

2018









2018



Ms. Pac-Man

Maluuba, a deep learning team acquired by Microsoft, created an AI system that learned how to reach the game's maximum point value of 999,900 on Atari 2600.

Chinese - English Translation

A <u>Microsoft</u> machine translation system achieved human-level quality and accuracy when translating news stories from Chinese to English. The test was performed on newstest2017, a data set commonly used in machine translation competitions.

Capture the Flag

A DeepMind agent reached human-level performance in a modified version of Quake III Arena <u>Capture the Flag</u> (a popular 3D multiplayer first-person video game). The agents showed human-like behaviours such as navigating, following, and defending. The trained agents exceeded the win-rate of strong human players both as teammates and opponents, beating several existing state-of-the art systems.

DOTA 2

<u>OpenAl Five</u>, OpenAl's team of five neural networks, defeats amateur human teams at <u>Dota 2</u> (with <u>restrictions</u>). OpenAl Five was trained by playing 180 years worth of games against itself every day, learning via self-play. (OpenAl Five is not yet superhuman, as it failed to beat a professional human team)

Prostate Cancer Grading

Google developed a <u>deep learning system</u> that can achieve an overall accuracy of 70% when grading prostate cancer in prostatectomy specimens. The average accuracy of achieved by US board-certified general pathologists in study was 61%. Additionally, of 10 high-performing individual general pathologists who graded every sample in the validation set, the deep learning system was more accurate than 8.

イロト イポト イヨト イヨト

Al and Compute (log scale), 1959-2019

Source: Compiled by OpenAl, 2019.



Overall PetaFlop Days
 Overall PetaFlop Days
 2 year doubling time Moore's law, starting from Perceptron

Fig. 3.14a.

Web Search and World News

US search interest for "machine learning," "deep learning," "TensorFlow" and "artificial intelligence" via Google Trends Source: Google Trends, 2019.



・ロト ・ 日 ト ・ モ ト ・ モ ト

э.

Autonomous Military Systems Developed Worldwide, 1970-2016 Since 2000, development of systems for combat, targeting, and other purposes has sharply increased.



Number of Autonomous Miltary Systems Developed, 1950-2017





04

(日) ъ

World Map of Countries Testing AVs

Source: Online searches on nations testing AV's.





Annual Installations of Industrial Robots ('000 of units), 2012-2018 Source: International Federation of Robotics, 2019.



(日)



・ロト ・四ト ・ヨト ・ヨト 三日

Employment growth and decline by occupation, % change labor demand, midpoint automation



Technology professionals

China Germany India Japan Mexico United

Includes: Computer engineers, computer specialists

AI and automation will have a profound impact on work McKinsey





Employment growth and decline by occupation, % change labor demand, midpoint automation



-4% 21%

Predictable physical work

Includes: Fine-equipment installation and repair workers, protective services, garning- industry workers, dishwashers, cleaning- equipment operators, food-preparation workers, general mechanics

Jobs lost, jobs gained: What the future of work will mean for jobs, skills, and wages



-31%

China Germany India Japan Mexico United

-33%

Growth in Introduction to AI Enrollment (relative to 2010)

Source: University provided data, 2019.



◆□▶ ◆□▶ ◆三▶ ◆三▶ 三日 • ���



Fig. 5.3.

Percent of World AI Private Investment, Startup Cluster (2018-19) Source: CAPIQ, Crunchbase, Quid, 2019.



Fig. 4.2.4.

Note: The chart shows the sum of total private Al investments between January, 2018 - October, 2019.



Regional Dynamics of AI labor demand in the US

Source: Burning Glass, 2019



Fig. 4.1.10. Notes: Alaska and Hawaii have not been presented for presentational brevity.

・ロト ・回ト ・ヨト ・ヨト

æ

Breakdown of use cases by applicable techniques, %

Full value can be captured using non-Al techniques

Al necessary to capture value ("greenfield")

Al can improve performance over that provided by other analytics techniques

fii

Potential incremental value from AI over other analytics techniques, %



Global AI startups that have received funding within the last year (July 2018-July 2019)

Source: CAPIQ, Crunchbase, Quid, 2019.



Fig. 4.2.6a. Network showing 4,403 global AI startups that received investment between July 2018 and July 2019. Colored by sector with top five highlighted.

・ロト ・ 一下・ ・ ヨト・

Share of Total Online Job Postings, USA, 2010-2019 monthly

Source: BurningGlass, 2019.



Fig. 4.1.2.

(日) (四) (日) (日) (日)

æ

Al startups in the United States: Companies that have received any funding within the last year, by focus area (July 2018-July 2019)





Al startups in the European Union: Companies that have received any funding within the last year, by focus area (July 2018-July 2019)



	Fashion, Retail Tech	5.7%	
	MedTech	4.4%	
	Text Analytics	4.4%	
٠	Marketing, AdTech	4.3%	
	Autonomous Vehicles	4.0%	
•	Data Tools	4.0%	
	Energy Management	4.0%	
	Chatbots	3.7%	
٠	Robotic Process Automation	3.7%	
	Accounting, Payments	3.6%	
	FinTech	3.6%	
٠	Automation, Oil & Gas	3.3%	
	Real Estate	3.3%	
۰	Bioinformatics, Life Science	3.0%	
	InsureTech	2.9%	
٠	Sports, Games	2.8%	
	Speech Recognition	2.7%	
٠	AR/VR	2.6%	
	Fraud Detection	2.6%	
٠	Network Security	2.6%	
	Medical Imaging	2.5%	
•	Agriculture Tech	2.4%	
	Credit Cards, Lending	2.3%	
٠	Recruiting	2.3%	
٠	Supply Chain Management	2.2%	
۰	EdTech	2.1%	
٠	Marketing/Sales Automation	2.0%	
	Chips, Semiconductor	1.9%	
	Hospitality, Travel	1.9%	
	Music, Entertginment	1.8%	
	Food and Beverage	1.5%	
٠	Drones, UAV	1.3%	
	Facial Recognition	1.3%	
۰	Mental Health, Wellness	1.3%	
	Legal Tech	1.2%	
٠	Human Resources	1.0%	
 I 		< E .	►

3 •

Global AI Investment, Merger/Acquisition, Minority Stake, Private Investment and Public Offering Source: CAPIQ, Crunchbase, Quid, 2019.



Public Offering = Private Investment = Minority Stake = Merger/Acquisition

▲□▶ ▲圖▶ ▲圖▶ ▲圖▶ 二直 - 釣��

Al and ML mentions in U.S. Congress (1995-2019)

Source: U.S. Congressional Record website, the McKinsey Global Institut



Al and ML mentions in Canadian Parliament (2002-2019)

Source: Parliament of Canada website, the McKinsey Global Institu



・ロト ・ 同ト ・ 日下 ・ 日

ъ

Al and ML mentions in U.K. Parliament (1980-2019)

Source: Parliament of UK website, the McKinsey Global Inst.



Quid network with 3,661 news articles on AI Ethics from August 12, 2018 to August 12, 2019. Colored by theme. Labeled by theme.



▲□▶ ▲□▶ ▲□▶ ▲□▶ = 三 のへぐ

Al use cases that support the UN Sustainable Development Goals

Source: 'Notes from the AI Frontier: Applying AI for social good', McKinsey Global Institute



Count of use cases

Al applicability to address UN SDG targets

Source: UN Global Indicator Framework, McKinsey Global Institute analysis

- 1. No poverty
- 2. Zero hunger
- 3. Good health and well-being
- 4. Quality education
- 5. Gender equality
- 6. Clean water and sanitation
- 7. Affordable and clean energy
- 8. Decent work and economic growth
- 9. Industry, innovation and infrastructure
- 10. Reduced inequalities
- 11. Sustainable cities and communities
- 12. Responsible consumption and production
- 13. Climate action
- 14. Life below water
- 15. Life on land
- 16. Peace, justice and strong institutions
- 17. Partnerships for the goals



Targets that could be supported by Al Total number of targets

19



◆□▶ ◆□▶ ◆臣▶ ◆臣▶ 三臣 - のへで

McKinsey Global Institute

AI in healthcare: Microsoft's Kevin Scott on how tech can treat a pandemic

July 28, 2020 | Podcast

McKinsey & Company

ET THE ECONOMIC TIMES

Automation, artificial intelligence to be central in the post-Covid world

▲□▶ ▲□▶ ▲□▶ ▲□▶ ▲□ ● ● ●



Base Pay (\$ Per Year)

Source: Glassdoor Economic Research. Active unique job listings on Glassdoor with "artificial intelligence" or "deep learning" job titles as of October 20, 2017.

HEATMAP: CONCENTRATION OF AI ACQUISITIONS BY CATEGORY (2011-2019 YTD)



CBINSIGHTS

▲ロト ▲冊 ▶ ▲ ヨ ▶ ▲ ヨ ▶ ● の Q @



AT THE CUSP OF THE 5TH INDUSTRIAL REVOLUTION

Recent rapid adoption and application of artificial intelligence algorithms — triggered by access to big data and better hardwareprocessing capabilities — are changing the face of blue and white collar jobs.

CBINSIGHTS

Ш MLTalks—Transhumanism: Searching for the Spirit in the **Aachine**

mit media lab

Artificial intelligence has become the catchphrase of our time, mediating great hopes and fears in anticipation of a new era in human history. Some expect that we're going to merge with machines and upgrade ourselves into God-like beings with divine abilities of creation and destruction. Others fear that machines will surpass their human creators and take control over oursworld.



The Ideology of Transhumanism:

The Ethics of Transcending Our Physical Bodies



How Humans Are Shaping Our Own Evolution

Like other species, we are the products of millions of years of adaptation. Now we're taking matters into our own hands.





Transhumanism: Meet the cyborgs and biohackers redefining beauty



Transhumanism And The Future Of Humanity: 7 Ways The World Will Change By 2030

Thought experiments

Forbes

The

How brains and machines can be made to work together

Brain-computer interfaces sound like the stuff of science fiction. Andrew Palmer sorts the reality from the hype



Neuralink Progress Update, Summer 2020

1890613 vues • Diffusé en direct le 29 août 2020

Elon Musk Announces Neuralink Advance Toward Syncing Our Brains With Al

Musk's company Neuralink demonstrates fully implantable chip in pigs

NEURALINK













REQUIREMENTS FOR WRITING TO THE BRAIN

Precise control of electric field in time and space

Wide range of current for different brain regions

No harm to brain over time

TWO-PHOTON MICROSCOPY



4,



TOWARDS HUMAN CLINICAL STUDIES

Received FDA Breakthrough Device designation in July

Preparing for first human implantation soon, pending required approvals and further safety testing

A,

Inside the Race to Hack the Human Brain

Bryan Johnson's insanely ambitious dream to create a neuroprosthesis would allow humans to "coevolve" with AI and even develop telepathy.



MIT Technology Review

Humans and technology / Brain-computer interface

The Entrepreneur with the \$100 Million Plan to Link Brains to Computers



・ロト ・四ト ・ヨト ・ヨト ・ヨ

Putting a computer in your brain is no longer science fiction **The Washington Post**

<ロト 4 回 ト 4 回 ト 4 回 ト 回 の Q (O)</p>



A map of myelin content (red and yellow are high myelin; indigo and blue are low myelin) in the left hemisphere of cerebral cortex. (Matthew F. Glasser and David C. Van Essen/Nature via AFP/Getty Images)



By Elizabeth Dwoskin Silicon Valley Correspondent

Bryan Johnson invests \$100 million in Kernel to unlock the power of the human brain

TE TechCrun

First outside investment to accelerate on-demand access to world-leading, non-invasive brain recording technology



FUNDRAISING NEWS

Kernel Raises \$53M from General Catalyst, Khos Ventures, Eldridge, Manta Ray, Tiny Blue Dot



Read More

SPOT[®]

A nimble robot that climbs stairs and traverses rough terrain with unprecedented ease, yet is small enough to use indoors. Built to be a rugged and customizable platform, Spot has an industry track record in remote operation and autonomous sensing.

> SHOP SPOT

CONTACT SALES







EXPLORER \$74,500.00

The Spot Explorer kit puts the power of robotics into your hands and makes robotics easy, so you can focus on building your application.

▲ロト ▲帰 ト ▲ ヨ ト ▲ ヨ ト ・ ヨ ・ の Q ()

AI adoption advances, but foundational barriers remain

November 13, 2018 | Survey

McKinsey & Company

The analytics academy: Bridging the gap between human and artificial intelligence

ARTIFICIAL INTELLIGENCE KEY CHALLENGES & OPPORTUNITIES

McKinsey & Company

Human judgment is still needed to ensure AI supported decision making is fair

Minimizing bias will be critical if artificial intelligence is to reach its potential and increase people's trust in the systems.

Six potential ways forward for artificial-intelligence (AI) practitioners and business and policy leaders to consider

1	2	3	4	5	6
•					
Be aware of contexts in which AI can help correct for bias and those in which there is high risk for AI to exacerbate bias	Establish processes and practices to test for and mitigate bias in Al systems	Engage in fact-based conversations about potential biases in human decisions	Fully explore how humans and machines can best work together	Invest more in bias research, make more data available for research (while respecting privacy), and adopt a multidisciplinary approach	Invest more in diversifying the Al field itself

McKinsey & Company

・ロト ・ 理 ト ・ ヨ ト ・ ヨ ト ・ ヨ

Human judgment is still needed to ensure AI supported decision making is fair

Minimizing bias will be critical if artificial intelligence is to reach its potential and increase people's trust in the systems.

Six potential ways forward for artificial-intelligence (AI) practitioners and business and policy leaders to consider



McKinsey Global Institute

How to ensure artificial intelligence benefits society: A conversation with Stuart Russell and James Manyika

January 31, 2020 | Video



McKinsey & Company

・ロト・日本・モート モー うへぐ





SF MASTERWORKS

"PHILIP K. DICK IS AWE-INSPIRING." — THE WASHINGTON POST

INCLUDES THE STORES THAT INSPIRED TOTAL RECALL, SCREA MINORITY REPORT, PAY AND NEXT

PHILIP K. DICK READER #1 NEW YORK TIMES BESTSELLER AUTHOR OF SAPIENS

Yuval Noah Harari



21 Lessons for the 21st Century

"Offers a helpful framework for processing the news and thinking about the challenges we face." --BILL GATES, IN "5 BOOKS I LOVED IN 2018"

イロト 不得 トイヨト イヨト







э





NOTES ON STARTUPS, OR HOW TO BUILD THE FUTURE

MARYIN

with BLA





Uncovering the Spectrum of Consciousness

7

MINSKT

THE INTERNATIONAL BESTSELLER

Being human in the age of Artificial Intelligence

Life 3.0

THE SECOND MACHINE AGE

WORK, PROGRESS, AND PROSPERIT In a time of Brilliant technologies

ERIK BRYNJOLFSSON ANDREW MCAFEE

