

ÉCOLE POLYTECHNIQUE – ÉCOLES NORMALES SUPÉRIEURES
ÉCOLE SUPÉRIEURE DE PHYSIQUE ET DE CHIMIE INDUSTRIELLES

CONCOURS D'ADMISSION 2016

FILIÈRES MP, PC ET PSI

ÉPREUVE ÉCRITE DE LANGUE VIVANTE – (XEULCR)

ANGLAIS

Durée totale de l'épreuve écrite de langue vivante (A+B) : 4 heures

Documents autorisés : aucun

PREMIÈRE PARTIE (A)
SYNTHÈSE DE DOCUMENTS

Contenu du dossier : trois articles et un document iconographique pour chaque langue. Les documents sont numérotés 1, 2, 3 et 4.

Sans paraphraser les documents proposés dans le dossier, le candidat réalisera une synthèse de celui-ci, en mettant clairement en valeur ses principaux enseignements et enjeux dans le contexte de l'aire géographique de la langue choisie, et en prenant soin de n'ajouter aucun commentaire personnel à sa composition.

La synthèse proposée devra comprendre entre 600 et 675 mots et sera rédigée intégralement dans la langue choisie. Elle sera en outre obligatoirement précédée d'un titre proposé par le candidat.

SECONDE PARTIE (B)
TEXTE D'OPINION

En réagissant aux arguments exprimés dans cet éditorial (document numéroté 5), le candidat rédigera lui-même dans la langue choisie un texte d'opinion d'une longueur de 500 à 600 mots.

A - Document 1

Artificial intelligence might be a threat to humans but not for the reasons you think

By Nigel Shadbolt¹; *The Guardian*,
Thursday 22 January 2015

[...] Our computers are getting better thanks to the exponential developments that drive this area of science and engineering. The computer you buy today is obsolete in R&D terms and yet is roughly twice as powerful as the one the same money could buy 18 months earlier. This has been happening for decades.

My students have access to computers that are 1 million times more powerful than the ones I began my AI research on back in the late 1970s. If we had improved air travel as fast I would fly from London to Sydney in less than a tenth of a second.

As well as more powerful computers, we have learned how to write software that “learns” to get better, “understands” human speech, and “navigates” from one place to another. I put the verbs in quotes because for the most part in AI we are not claiming that the algorithms operate in the way that we do when we solve similar tasks.

A founding father of AI once said “there are lots of ways being smart that aren’t smart like us”. What we have built in AI are numerous slivers of smart behaviour, a digital ecosystem populated with adaptive systems narrowly crafted to a particular niche.

When a high-end computer beat Garry Kasparov, the world chess champion, in the 90s it didn’t usher in a new age of intelligent machines. It did demonstrate what you could do with large amounts of computer power, large databases full of moves and good heuristics to look ahead and search possible moves. The overall effect on the world chess champion was unnerving. Kasparov felt as if Deep Blue was reading his mind. Deep Blue had no concept there was another mind involved.

But it is easy to endow our AI systems with general intelligence. If you watch the performance of IBM’s Watson as it beats reigning human champions in the popular US TV quiz show you feel you are in the presence of a sharp intelligence. Watson displays superb general knowledge – but it has been exquisitely trained to the rules and tactics of that game and loaded with comprehensive data sources from Shakespeare to the Battle of Medway. But Watson couldn’t play Monopoly. Doubtless it could be trained – but it would be just another specialised skill.

We have no clue how to endow these systems with overarching general intelligence. DeepMind, a British company acquired by Google, has programs that learn to play old arcade games to superhuman levels. All of this shows what can be achieved with massive computer power, torrents

1. About the author : Sir Nigel Richard Shadbolt is Principal of Jesus College, Oxford, and Professorial Research Fellow in the Department of Computer Science, University of Oxford. He is Chairman of the Open Data Institute which he co-founded with Sir Tim Berners-Lee. He is also a Visiting Professor in the School of Electronics and Computer Science at the University of Southampton.

of data and AI learning algorithms. But our programs are not about to become self-aware. They are not about to apply a cold calculus to determine that they and the planet would be better off without us.

What of “emergence” – the idea that at a certain point many AI components together display a collective intelligence – or the concept of “hard take off” a point at which programs become themselves self-improving and ultimately self-aware? I don’t believe we have anything like a comprehensive idea of how to build general intelligence – let alone self-aware reflective machines.

But there are lots of ways of being smart that aren’t smart like us, and there is the danger that arises from a world full of dull, pedestrian dumb-smart programs. Of hunter kill drones that just do one thing very well – take out human targets. Done at scale this becomes an existential risk. How reflective does a system have to be to wreak havoc. Not at all if we look to nature and the self-replicating machines of biology such as Ebola and HIV.

AI researchers are becoming aware of the perils as well as the benefits of their work. Drones full of AI recognition and target acquisition software alarm many. We need restraints and safeguards built into the heart of these devices. In some cases we might seek to ban their development altogether.

We might also want to question the extent and nature of the great processing and algorithmic power that can be applied to human affairs, from financial trading to surveillance, to managing our critical infrastructure. What are those tasks that we should give over entirely to our machines? These are ethical questions we need to attend to. [...]

A - Document 2

The Master Algorithm : A world remade by machines that learn

By Anil Ananthaswamy, *New Scientist*,
October 28, 2015

WHEN machine learning algorithms that replace newspaper reporters became fodder for a recent episode of Comedy Central's *The Daily Show*, it was clear that the technology had gone mainstream.

But as Pedro Domingos points out in *The Master Algorithm*, machines that learn have been deeply involved with our lives for a while. If you use Google, Netflix, Amazon, Pandora, Yelp, Xbox or just about any online dating service, your life is being run by algorithms that are learning more and more about you by chomping on the data you, sometimes unwittingly, provide.

"Society is changing, one learning algorithm at a time. Machine learning is remaking science, technology, business, politics and war," writes Domingos, a computer scientist at the University of Washington, Seattle. For people in his field, the problem is that there are myriad such algorithms, each trying to discern patterns in the masses of data we produce. "Machine learning is about prediction," he writes, "predicting what we want, the results of our actions, how to achieve our goals, how the world will change."

The book is about the quest for that one master algorithm which would change machine learning, and hence our lives, irrevocably. If it exists, says Domingos, the master algorithm can derive all knowledge in the world "past, present, and future – from data". In theory, such an algorithm could derive Newton's laws from the astronomical observations of Tycho Brahe, with no a priori knowledge of such laws.

But why should such an algorithm even exist? Domingos provides compelling arguments from neuroscience, evolution, physics, statistics and computer science. For instance, the cerebral cortex might be an instance of such an algorithm : some neuroscientists think that it implements the same algorithm all over, just tweaked to learn to see or hear, or to make sense of touch.

Depending on your world view, the development of a master algorithm is either really thrilling or downright scary. It's not surprising that Domingos, an expert in machine learning, has a very optimistic view. He clearly sees the master algorithm as desirable and maybe even inevitable. This cheery outlook shines through large parts of the book, when he writes that such an algorithm will "speed poverty's decline", that routine jobs "will be automated and replaced by more interesting ones", that the health of our planet will "take a turn for the better", and that our own lives will be "longer, happier and more productive".

Domingos has few doubts, and those he has mainly concern whether the technology will really happen as promised. "Maybe," he muses, "the master algorithm will take its place among the great chimeras, alongside the philosopher's stone and the perpetual motion machine."

But what about the future that lies in store for us, should machine learning take over our lives (if it hasn't already)? Again, Domingos sees it all as a positive. "Someday there'll be a robot in

every house, doing the dishes, making the beds, even looking after the children while the parents work. How soon depends on how hard finding the Master Algorithm turns out to be.” [...]

It’s hard to avoid the feeling that machine learning is only going to increase the rift between the haves and the have-nots, as we enter a new phase of survival of the fittest. As Domingos writes, “He who learns fastest wins”, and machine learning “is the latest chapter in the arms race of life on Earth”.

But he’s still not worried. As machine learning does away with most jobs, the world Domingos envisions consists of a large class of unemployed people living on a permanent basic income doled out by the government, while those in the few remaining human occupations will be stupendously wealthy. “For those of us not working, life will not be meaningless, any more than life on a tropical island where nature’s bounty meets all needs is meaningless.” [...]

The Master Algorithm : How the quest for the ultimate learning machine will remake our world
Pedro Domingos
Basic Books/Penguin

A - Document 3

How We Can Overcome the Risks of AI

Andrew Lohn, Andrew Parasiliti and William Welser IV, *Time Magazine*,
October 22, 2015

Apple's recent acquisition of Vocal IQ, an artificial intelligence company that specializes in voice programs, should not on its face lead to much fanfare: It appears to be a smart business move to enhance Siri's capabilities. But it is also another sign of the increased role of AI in our daily lives. While the warnings and promises of AI aren't new, advances in technology make them more pressing. Forbes reported this month: "The vision of talking to your computer like in *Star Trek* and it fully understanding and executing those commands are about to become reality in the next 5 years." Antoine Blondeau, CEO at Sentient Technologies Holdings, recently told *Wired* that in five years he expects "massive gains" for human efficiency as a result of artificial intelligence, especially in the fields of health care, finance, logistics and retail.

Blondeau further envisions the rise of "evolutionary intelligence agents," that is, computers which "evolve by themselves – trained to survive and thrive by writing their own code—spawning trillions of computer programs to solve incredibly complex problems."

While Silicon Valley enthusiasts hail the potential gains from artificial intelligence for human efficiency and the social good, Hollywood has hyped its threats. AI-based enemies have been box office draws at least since HAL cut Frank Poole's oxygen hose in *2001: A Space Odyssey*. And 2015 has truly been the year of fictional AI provocateurs and villains with blockbuster movies including *Terminator Genisys*, *Ex-Machina*, and *The Avengers: Age of Ultron*. But are the risks of AI the domain of libertarians and moviemakers, or are there red flags to be seen in the specter of "intelligence agents?" Silicon Valley cannot have "exponential" technological growth and expect only positive outcomes. Similarly, Luddites can't wish away the age of AI, even if it might not be the version we see in the movies.

The pace of AI's development requires an overdue conversation between technology and policy leaders about the ethics, legalities and real life disruptions of handing over our most routine tasks to what we used to just call "machines." But this conversation needs to focus increasingly on near-term risks, not just cinematic ones.

For example, even if a supercomputer's coding is flawless, and someday self-generated, and is protected from being infected by a warring nation-state, a hacktivist, or even an angry teenager, AI can still produce wrong answers. A *Wired* article from January 2015 showed just how wrong. When presented with an image of alternating yellow and black parallel, horizontal lines, state of the art AI saw a school bus and was 99% sure it was right.

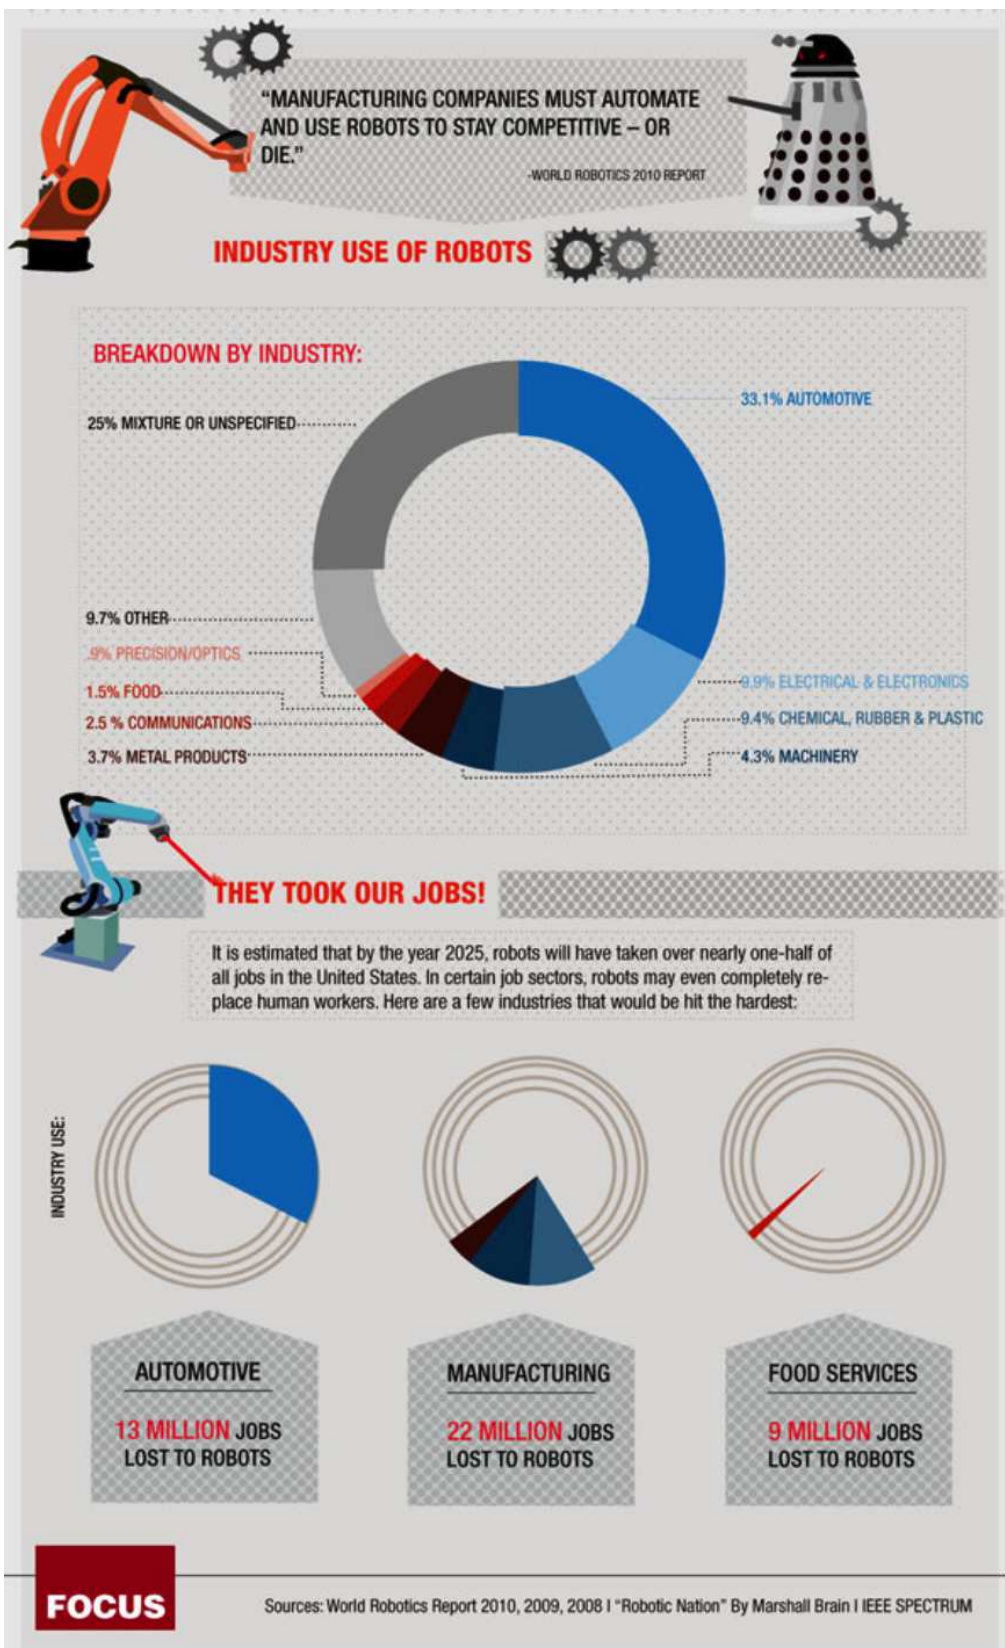
How far can we trust AI with such control over the Internet of Things, including our health, financial, and national defense decisions? There is a service to be done in developing a deeper understanding of the reasonable precautions needed to mitigate against coding flaws, attackers, infections and mistakes while enumerating the risks and their likelihoods.

Applied to military systems the risks are obvious, but commercial products designed by AI could produce a wide range of unexpected negative outcomes. One example might be designing fertilizers that help reduce atmospheric carbon. The Environmental Protection Agency tests such products before they are approved so dangerous ones can be discovered before they are released. But if AI only designs products that will pass the tests, is that AI designing inherently safe products or simply ones capable of bypassing the safeguards? [...]

Can the risks posed by AI be completely eliminated? The short answer is no, but they are manageable, and need not be cause for alarm. The best shot at providing adequate safeguards would be regulating the AI itself: requiring the development of testing protocols for the design of AI algorithms, improved cybersecurity protections, and input validation standards—at the very least. Those protections would need to be specifically tailored to each industry or individual application, requiring countless AI experts who understand the technologies, the regulatory environment, and the specific industry or application. At the same time, regulatory proposals should be crafted to avoid stifling development and innovation.

AI needs to enter the public and political discourse with real-world discussion between tech gurus and policymakers about the applications, implications and ethics of artificial intelligence. Specialized AI for product design may be possible today, but answering broad questions such as, “Will this action be harmful?” is well outside the capabilities of AI systems, and probably their designers as well.

Answering such questions might seem like an impossible challenge, but there are signs of hope. First, the risks with AI, as with most technologies, can be managed. But the discussions have to start. And second, unlike in an AI-themed Hollywood thriller, these machines are built to work with humankind, not against it. It will take an army of human AI experts to keep it that way, but precautions can and should be sought now.



B - Document 5
Robot Rights Rule!

Artificial intelligence challenges the distinction between man and machine

By *The Washington Times*,
Sunday, July 26, 2015

The season of the Theater of the Absurd continues. After the Supreme Court twisted the clear meaning of plain English words to save Obamacare and bless same-sex marriage, after Iran hoodwinked Barack Obama into preserving and expanding its nuclear program, after Bruce Jenner remade himself (herself? itself?) into a buxom synthetic female, no one should be surprised when R2D2 wakes up to demand his civil rights, too. This might not be what Mr. Obama had in mind, but a conscientious radical accepts everything new, bad or not.

If self-awareness is the essence of what it means to be human, and humans merit rights, machines may soon be ready to claim their birthright (assembly-right?). Computer scientists at Rensselaer Polytechnic Institute in Troy, N.Y., have taught humanoid robots to recognize themselves as distinct from others. Taking a group of three robots, researchers administered a “dumbing pill” program to two of them, which told them they were unable to speak. When the group was asked which one could still speak, the third robot spoke up, recognized its own voice and announced that it was the one. It’s not exactly Descartes, “I think, therefore I am,” but for a robot, it’s not bad.

Robo-ethics, the morality of how robots are designed and tasked, is challenging scientists and engineers to ponder the possibility – some say the inevitability – of artificial intelligence advancing far beyond simple self-awareness, to outsmart the creators. Once they comprehend the concept of personhood, robots could grasp the idea that society is obligated to grant them rights, similar to the human rights described in the Declaration of Independence, such as “life, liberty and the pursuit of happiness.” But will a happy robot be a good robot?

An ethicist says that now is the time to ponder the enigmatic questions of cyber law: “Robotic systems accomplish tasks in ways that cannot be anticipated in advance; and robots increasingly blur the line between person and instrument,” says Ryan Calo, a professor at the University of Washington School of Law. If, in the future, a demonstrably sentient machine claims the right that humans have to procreate, or build copies of itself, who can say nay? When the multiplying machines petition for the right of representation in governance, men and women born of nature will face an ethical dilemma. “Which right do we take away from this sentient entity, then,” Professor Calo asks, “the fundamental right to copy, or the deep, democratic right to participate?”

Bringing down the curtain on this season of the Theater of the Absurd, by ordaining that rights are reserved for flesh-and-blood humans, may not be that simple. As replacing human hips and knees has become routine medicine in the 21st century, so might integration of bionic body parts to remedy the ravages of injury or disease in coming decades. Should society draw the line between man and machine when cyborgs – part living, part mechanical – show up at

the courthouse to register to vote? The befuddlement that accompanied the use of the “one-drop rule” in determining the race of Americans of mixed ancestry in years past, would be minor by comparison.

If robot rights seem a stretch, animal rights sound equally silly, but one nonhuman creature has won rudimentary human rights. In 2014, an orangutan in Argentina named Sandra was granted legal personhood through the imagination of the lawyers. A court ordered Sandra released from prison (a zoo, actually) on the grounds that as an intelligent, nonhuman primate, she is entitled to the freedom to live in a sanctuary rather than in a cage.

Believing that artificial intelligence will soon render robots to be humans of a different kind, one socially insensitive wag has taken up their cause with a slogan: “Robot lives matter.” Don’t laugh.