VOX

The Student Journal of
Politics, Economics & Philosophy

The Fourth Revolution

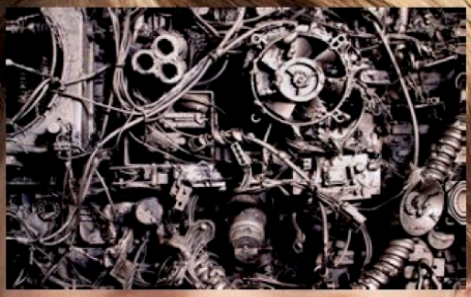

Autumn 2018

Editorial Note

The global tech and innovation race has proven to be one of the most critical and sensitive challenges for great powers, with no signs of slowing down anytime soon. Russian leader Putin once stated that 'the nation that leads in AI will be the ruler of the world' (September 2017). According to the 2018 Deloitte's yearly Technology, Media & Telecommunications report, the number of companies utilising AI will double by the end of the year. Although no nation has clearly won this competition yet, citizens surely are entitled to wonder what a daily life with the omnipresence of AI would be like.

On many levels, AI makes us look at future prospects with both enthusiasm and caution. The future of mankind is facing a robotic handful of opportunities, along with its ample risks and moral issues. By choosing the theme of AI for this issue, we encourage you to look at it from different angles to shape your own opinion through concise and reliable sources.
These include: a retrospective look at human rights applied to robots by Veer, a utopian vision of governance covered by Edward, as well as philosophical and ethics analyses led by both Sandrine and William. In addition to this, an insightful anthropological review will be provided by Elise, and last but not least, AI and its impact on climate change will be explored by Clara.

A special thanks to Werner Bonefeld without whom none of this academic journey would be possible, former Vox presidents Hannah Abban & Richard Wang for their support with the transition to the new committee and new committee members who show amazing dedication and excitement for this upcoming academic year!

As we launch this issue, we are in the process of expanding the exposure of Vox Journal, our young online blog named 'Voxensus'. For this new venture we strongly encourage student contribution, regardless of their academic backgrounds, as we are always seeking more diversity within our publications. All of this variety of support has been put in place in order to empower the University of York students' voices, enhance and promote the creation of ideas, debate and exchange within our academic community.

We like to think it is a right, an opportunity as well as a duty. We hope you enjoy our first issue of the year.

Idriss Bendrif

President of Vox Journal

CONTENTS

Source: https://www.shutterstock.com/search/robot+eyes

"The development of full artificial intelligence could spell the end of the human race... It would take off on its own, and re-design itself at an ever increasing rate. Humans, who are limited by slow biological evolution, couldn't compete, and would be superseded."

Stephen Hawking, Theoretical physicist and cosmologist

"A year spent in artificial intelligence is enough to make one believe in God"

Alan Perlis, American Computer Scientist and Professor at Yale University

"I'm most grounded on the role of technology. Ultimately to me it's about the human capital and the human potential and technology empowers humans to do great things. You have to be optimistic about what technology can do in the hands of humans."

Satya Nadella, CEO of Microsoft Corporation USA

The sad thing about artificial intelligence is that it lacks artifice and therefore intelligence."

Jean Baudrillard, French sociologist and philosopher

"Competition for AI superiority at national level is the most likely cause of WW3"

AI: The Good, the Bad and the Ugly

By Sandrine Chausson
2018 Graduate of the University of York

When people think of AI, they often do so in hypothetical terms. AI is perceived as a future possibility, not as a current reality. However, AI is very much present in our lives already, whether we realise it or not. You interact with an AI almost every time you play a video game (Lou, 2017). The personalised product, movie, and music recommendations you come across on a daily basis are most likely the product of an AI (Wei, 2017). Gmail uses an AI to detect and inform you when you forgot to attach a file to an email, and a semi-independent AI generates a number of Yahoo's news stories (Jenkin, 2016). While each of these AIs is an expert at a specific task, its "expertise" is not, in most cases, transferable to other applications. As such, they are called "narrow AIs". The real challenge, the one actively being taken on by AI researchers today, is to build a "general" AI, that is, an AI that is at least as competent as the average human at any given cognitive task (Future of Life Institute, n.d.). To understand some of the motivation behind const-ructing such a general AI, con-sider the following argument.

We have long been aware of the physical, emotional and intell-ectual limitations that cause us to make mistakes. Computer programs, however, do not suffer from these limitations. They do not lose track of their objectives or the motivation to pursue them and are able to keep playing for as long as necessary without the internal disruption a human might face, such as emotional distraction or fatigue. Building on these strengths, many narrow AIs perform overwhelmingly better than humans: for example, AIs have already largely surpassed humans at games such as Chess, Backgammon, or even *Jeopardy*!

This drive for performance might seem a bit futile if we look only at games. However, what make us underperform as chess players also make us under-perform in all spheres life, where the stakes are higher. Where these stakes amount to the wellbeing and safety of others, our mistakes can cause tragedies. Just think of road accidents, collateral casual-ties in military operations, or wrong-ful convictions. Surely, if we can build AIs that play games better than us, we could also build AIs that are better than us at driving,

planning and executing military operations), or making informed decisions during trials.

A complication arises here though. As long as the decisions made by AI where restricted to games, their outcomes were relatively inconsequential and unproblematic. However, in situations where the wellbeing or safety of others might be impacted by a decision, complex moral considerations arise that need to be taken into account. So how do we get an AI to factor in these moral considerations? At first sight, this might not seem more complicated than getting an AI to factor in any other variable. We could simply "upload" morality into our AI. That is, we could tell our AI exactly what variables to look for in morally sensitive scenarios, how to value them, and provide it with a method to come up with a decision.

Alternatively, we could train our AI to be moral through experience: we could present it with different morally sensitive decisions, tell it how it should choose, and let it infer moral rules through this process. Unfortunately, things are not this simple.

To begin with, there is a communication problem. Insinuations, figurative speech, references to experiential know-ledge and collectively shared ideas, or even reliance on the readers' ability to infer meaning from context are pervasive in the way we communicate. Moral discourse is not exempt from this. While these imprecisions do not hinder our own capacity to understand the content of our moral theories, they definitely hinder our capacity to communicate it with an AI. So we effectively need a new "formalised" moral language, which we do not have yet (Polonski, 2017; Bostrom, 2017, p.226-9). A second, much more substantial problem has to do with the uncertainty of our moral knowledge. Moral matters are challenging and, to this day, "what is and isn't moral?" remains an open question, which has been approached by many different people in many different ways. And while defenders of different approaches might all individually feel a sense of certainty about their own stance, objectively it is unclear which stance is the right one, and whether there is one at all (*idem*, p256-9).

Thus programming an AI to mirror what we believe is "true morality" is problematic, as it is likely to lead us to one of the following results. The first is moral paralysis: we do not assign our AI to the task of making morally sensitive decision, as we are unable to reach a sufficient level of agreement or certainty with regards to our moral theories and moral claims. As a result, the AI cannot fulfil its potential for good and people continue to suffer from human error. The second possibility it that we do assign our AI to the task of making

morally sensitive decision according to some values and principles we believe are the right ones. The danger here is clear though: we might be wrong. After all, history contains countless examples of atrocious crimes committed on the basis of some moral beliefs about namely race, sexuality, or religion, which retrospectively we con-sider flawed. So what would happen if our AI were to act according to our flawed beliefs? It seems that instead of allowing us to prevent harm, it would simply increase our capacity to cause harm (*ibid*.).

Fortunately, there might be a way out of this dead end. It seems plausible at least that these moral shortcomings of ours are caused by the same physical, emotional and intellectual limitations that make us less competent than an AI at chess or Backgammon. It seems tempting then to appeal once again to the superiority of AI to work out "morality" for us: what defines morality and why it is important; the general principles that distinguish the good from the bad, the impermissible from the permissible; the correct line of action in this or that particular scenario, etc. Evidently thought, to "figure out" morality, our AI will require a much greater set of cognitive abilities than to play chess or to drive a car. After all, research has shown that the regions of the human brain involved in moral reasoning are numerous, and serve a great number of different functions (Cuñat-Agut, Martí-Vilar and Suay I Lerma, 2016). It seems, then, that building an AI capable of moral reasoning essentially amounts to building something like a general AI.

So far, the creation of a general AI seems to hold the promise of beneficial change. However, many researchers are worried about building such a general AI. Their worries do not resemble those exemplified by movies like "The Matrix" or "Terminator" though, where the creature turns against its creator, and becomes evil or dominant: a Frankenstein type of scenario, if you will. Instead, the real threat resembles more that of a genie in a bottle. One danger is that this genie could fall into the wrong hands. For instance, imagine that those who gain access to AI technology are actually an oppressive state, a criminal organisation or a violent political faction. Given that AI will increase their capabilities unprecedentedly, the harm these could cause would also increase. This threat becomes particularly salient when we consider the possibility of an AI being hacked or hijacked, and as a result being used in unintended ways (Future of Life Institute, n.d.). Moreover, even if we manage to restrain our AI in such a way that it only pursues beneficial ends, there is another danger. Imagine for example that it is charged with the task of making the world fairer, more equal,

and less violent. If we do not stipulate appropriate restrictions on the actions it can take, it might decide to resort to manipulation, coercion, or other undesirable methods to fulfil its mission. So clearly, before resorting to our genie, we will need to be clear about how much freedom we are willing to give it, and what are objectives are in bringing it into this world (*ibid*.).

I would like to end this article with a few mitigating remarks. It is important to keep in mind that the dangers of AI described in this article are not insurmountable, nor are they immediate. Overall, AI experts predict that we have at least a few decades before AI becomes a reality (Grace et al., 2017). So the appropriate response to these dangers is not to panic and completely ban AI research: the truth is that research and innovation in the field of AI will probably persist, whether we want it or not. Instead, what we need to do, is plan ahead, build the framework that will help us anticipate the risks of AI, and develop the strategies to pre-vent these. Doing so will allow us to steer AI research in a positive direction (Future of Life Institute, n.d.).

Bibliography

Bostrom, N. (2016). *Superintelligence*. Oxford: Oxford University Press.

Cuñat-Agut, M., Martí-Vilar, M. and Suay I Lerma, F. (2016). Brain Structures Involved in Moral Reasoning: A Review. *The International Journal of Psychosocial and Cultural Genomics, Consciousness & Health Research*, 2(II), pp.38-45.

De Spiegeleire, S., Maas, M. and Sweijs, T. (2017). *Artificial Intelligence and the Future of Defense*. [online] The Hague: The Hague Centre for Strategic Studies.

Future of Life Institute. (n.d.). *Benefits & Risks of Artificial Intelligence*. [online]

Grace, K., Salvatier, J., Dafoe, A., Zhang, B. and Evans, O. (2017). *When Will AI Exceed Human Performance? Evidence from AI Experts*. [ebook] Future of Humanity Institute, Oxford University

Jenkin, M. (2016). *Written out of the story: the robots capable of making the news*. [online] The Guardian.

Johnston, C. (2016). *Artificial intelligence 'judge' developed by UCL computer scientists*. [online] The Guardian.

Lou, H. (2017). *AI in Video Games: Toward a More Intelligent Game*. [online] Science in the News.

Polonski, V. (2017). *Can we teach morality to machines? Three perspectives on ethics for artificial intelligence*. [online] Oxford Internet Institute

Wei, J. (2017). *Recommended For You: How machine learning helps you choose what to consume next*. [online] Science in the News.

Seasteading: The Next Great Wave of the Pacific Ocean?

By Edward Sell
2nd Year Student in Politics, Philosophy and Economics

French Polynesia, an over-seas territory of France some 2,600 miles south of Hawaii, is known worldwide for its white-sand beaches and laid-back lifestyle. However, this sleepy region of the South Pacific is due for a rude awakening. Seasteading, a portmanteau of 'sea' and 'home-steading', is the relatively modern concept of creating autonomous floating communities in international waters, away from the clutches of central governance. Its pilot project in French Polynesia is set to commence in just two years' time, with legislation from local govern-ment to allow for a special governing framework and economic zone.

Randolph Hencken, Executive Direct-or of the Sea-steading Institute, the organisation en-trusted with imple-menting this project, claims it will 'bring jobs, economic growth, and environ-mental resiliency to the region' (Hencken, 2016, cited in Wong, 2017); the latter of Hencken's assertions is particularly pertinent to French Polynesia given its vulner-ability of rising sea levels and climate change. However, I raise caution on the advantages of seasteading. I will put for-ward a multi-faceted argu-ment against the ramifications of sea-steading from logistical, socio-economic, and political angles, in order to demonstrate both the project's impracticalities and uncert-ainties on both the local and global political climate.

Logistics: Actualising Science Fiction

A pragmatic objection to seasteading is the simple question of its viability. In the back-end of the last decade, the Seasteading Institute was of broad interest amongst investment companies and venture capitalists in Silicon Valley. However, the tech-nological elite are casting the Institute, and its mandate of a libertarian island-nation paradise, in an ever-glooming light. The feeling of wavering optimism chiefly stems from the Institute's ongoing project de-lays and ballooning budget forecasts; estimates released put the cost of the Polynesian pilot project anywhere between US$10 million to US$167 million (Quirk, 2017; Gabbatiss, 2017). This inconsistency in budget calculation has fuelled the scepticism of venture capitalists such as Peter Thiel, PayPal co-founder and initial financier of the Institute. In a 2015 speech at George Mason University, Thiel admitted that 'I'm not exactly sure that I'm going to succeed in building a libertarian utopia any time soon…you need to have a version where you could get

started with a budget of less than $50 billion.' (Thiel, 2015, cited in DeNuccio, 2015) This apparent about-face of one of the project's key stakeholders is only emblematic of the financial challenges and budget inconsistencies ingrained deep in the Institute.

of SW Pacific islands.' (Stephens and Ramsay, 2014, p.14; ibid, p. 24)

The goal of a floating ocean city assumes the presence of relatively still waters and moderate weather conditions, but scientific research highlights that this is yet another

https://medium.com/@kyrongosse/5-ways-seasteading-will-change-the-world-35e3d105f2f1

A further logistical challenge facing seasteading is architectural resilience to the region-al environment, or lack thereof. A 2014 journal on Global and Planetary Change has laid bare the extent to which the South Pacific climate would present a tremendous challenge to any potential sea-steading project. It reports of an 'unprecedented number of intense [tropical cyclones] …in the vicinity of French Polynesia', noting further that 'Large waves and wave amplification from [tropical cyclones] are a signif-icant natural hazard for inhabit-ants

gross miscalculation. Indeed, the Institute was forced to conduct a reassessment on their plans of a community residing entirely in international waters, opt-ing instead for the relatively calmer conditions within French Polynesian territory. While this does partly answer the question of foundational integrity against the regional environment, it defeats the very purpose on which seasteading was based – to create a self-governing com-munity. Despite promises of a bespoke administration deal, any seastead will to some

extent be beholden to French Polynesian governance, as it will lie within the territory's maritime borders. Given the current limits in the fields of architecture and technology, any form of seasteading is seemingly impossible, and the pilot project is altogether removed from the notion of seasteading by definition.

A Society at Sea

Not only is seasteading flawed in its potential application, it also carries significant drawbacks from a theoretical perspective. Perhaps the most objectionable of these is the exemption from taxes, a key component of the project's economic plans. There is considerable disapproval amongst the French Polynesian populace, who are un-willing to see millionaire seasteaders enjoy tax breaks while their economy suffers. Alexandre Taliercio, a local resident and television host, said that seasteaders 'seem to have much more to gain than we do…the unemployment rate is distressing; the impoverishment of the population is palpable every-where.' (Taliercio, n.d., cited in Wong, 2017) Indeed, the territory's economy ranks 173rd in the world with a GDP (PPP) of US$5.49 billion, and 187th for national unemployment rates with 21.8% of the working-age population (cia.gov, 2018). Moreover, the development and technical maintenance of a sea-stead complex would require a highly skilled labour force. Given that only 0.9% of boys and 1.3% of girls in French Polynesia achieve a high-school diploma or equivalent (Amit and Dyck, 2011), it is hard to imagine that this demand for labour would be supplied locally. Therefore, it raises the question as to whether government facilitation of tax-exempt economic zones such as the pilot project would bring any net benefit to French Polynesia, a region already struggling with weak economic performance. It seems clear that Hencken's claim of creating jobs and economic growth in the region is merely a gilded, hollow promise to inveigle local authorities to look favourably upon seasteading.

From an international perspective, the Institute's plans have the potential to set a dangerous precedent. Seasteading, an inherently libertarian concept, faces a moral quandary: those behind the project need to find the elusive balance between maximising personal freedoms while maintaining the presence of state intervention. Failure to find a harmony may lead to artificial island-cities full of money-launderers, market monopolies, and a large degree of wrongdoing. Alternatively, excessive state control may severely disenchant the libertarian émigrés, living in a seastead to get away from that exact issue. It is this particular problem that offers one explanation as to why there currently exists no libertarian nation. Michael Lind reinforces this idea: 'If liber-

tarianism was a good idea, wouldn't at least one country have tried it?' (Lind, 2013) If seasteads are not unsustainable in architectural terms, then they face a political instability of equal magnitude.

I Fought the Law…

Even the most libertarian of sea-steads, equipped with the freest of markets and the highest degree of personal autonomy, would still be subject to international law. Extra-territorial jurisdiction (ETJ) - the right for national governments to extend their legal authority beyond its borders - can mean one of two things for the opportunistic seasteader.

Firstly, nations have the right to exert national pressure on seasteads through the flag state registration technicality. Every nautical structure, regardless of whether it is within or outside of any maritime borders, must fly a Flag of Convenience of the country under which it is registered because such structures have no national sovereignty in and of itself. For instance, in the era of British pirate-radio station structures based in international waters during the 1960s, the British government was able to outlaw the supply of food, water, and advertising to those sovereignty challengers under ETJ (Barker, 2009). Although the fre-quency and severity of ETJ implementation varies from country to country, by no means are seasteaders exempt from the law on their libertarian island.

Secondly, ETJ can extend its reach into taxation as well. A few countries, notably the United States, can impose its emigrants with an expatriation tax. Among those that fall under this bracket are renouncers of American citizenship, in order to deter tax avoidance by means of migration loopholes or abandonment of national residency. The role that ETJ can play in taxing emigrants is a very significant obstacle for American seasteaders in particular, who seek to rid their pockets of the pinching hands of the state. While the vast majority of national governments have no such tax in place as of 2018, a future of numerous seastead-ing communities worldwide may lead to adverse economic effects. If more citizens are moving on to pastures new in the tax-free seasteads, more governments may adopt an Amer-ican style expatriation tax to deter brain drain and the flight of capital.

While a mass increase of expatriation taxes could be effective in treating the sea-steading problem, it poses a much larger con-sequence on the mobility of labour. A 2015 joint report found that 78% of Americans working abroad feel that the current tax system 'puts them at a prof-essionnal disadvantage compared to others working in their [current] country of residence.' (Burggraf, 2016) Therefore, a global com-

munity beset with numerous stringent tax regulations on expats could lead to a fall in global migration, which in turn could see a fall in remittances and job opportunities for citizens of imposing countries.

Conclusion

This paper has taken the necessary and unfortunate duty to illuminate the sheer extent to which the hopeful seasteading project is fraught with inadequacies. This libertarian dream of the Silicon Valley elite is merely that - a dream. Seasteading in its rawest of forms is unachievable, given the current limitations of environmental engineering technology. More-over, even if advances in technology were to allow for structurally feasible sea-steading in rough inter-national waters, factors such as ETJ mean that the libertarian islanders would ultimately fail to escape from the very government from which they sought to do so. A final reflection this paper would like to make is to echo the sentiments of Fujitsu Singapore President Wong Heng Chew:

"Digital technology has evolved through four major waves of development. Life underwent major changes when the Inte-rnet made computing techno-logy avail-able to all, and again when the mobile Internet made digital services accessible any-where. A new wave of change is occurring today …but it is a fourth wave of technology that organizations must now con-sider…" (Chew, 2016)

Due to its financial, logistical, and hypothetical shortcomings, seasteading cannot play a role of any significance in this 'fourth wave of technology'. For the time being at least, the only waves that will be troubling the islanders of French Polynesia will be of the oceanic variety.

Bibliography

Amit, V. and Dyck, N. (2011) , *Young Men in Uncertain Times,* Oxford: Berghahn Books, p.161.

Barker, V. (2009). *The Real Story Behind Britain's Rock 'N' Roll Pirates*. [online] NPR.org.

Burggraf, H. (2016). *Expat Yanks say US tax laws 'put them at a disadvantage'*. [online] Inter-national Investment.

Chew, W. (2016). *The 4th wave of technology is going to be artificially intelligent*. [online] Enterprise Inno-vation | Asia's Premier Business and Tech-nology Publication.

cia.gov. (2018). *The World Factbook*[online] Available at:

DeNuccio, K. (2015). *Silicon Valley Is Letting Go of Its Techie Island Fantasies*. [online] WIRED.

Gabbatiss, J. (2017). *The world's first floating city will be built by 2020*. The Independent.

Lind, M. (2013). *The question libertarians just can't answer*. Salon.

Meessen, K. (1996), Extra-territorial *Jurisdiction in Theory and Practice*. Leiden: Brill.

Quirk, J. (2017). *Government of French Polynesia Signs Agree-ment with Seasteaders for Floating Island Project*. [online] The Seasteading Institute.

Stephens, S. and Ramsay, D. (2014). Extreme cyclone wave climate in the Southwest Pacific Ocean: Influence of the El Niño Southern Oscillation and projected climate change. *Global and Planetary Change*,

The Economist. (2011). *Cities on the ocean*.

Wong, J. (2017) *Seasteading: tech leaders' plans for floating city trouble French Polynesians*. The Guardian.

Should Robots have Rights?

By Veer Sharma
3rd Year Undergraduate in English and Related Literature

"*Nobody phrases it this way, but I think that artificial intelligence is almost a humanities discipline. It's really an attempt to understand human intelligence and human cognition"*

Sebastian Thrun (Chafkin 2013)

When the first legislator begins to ponder the question "Should robots have rights?" it will be the end of human civilisation, for civilisation, as we know it, will be post-human from that moment onwards. No question up until this point has so fundamentally challenged the status of humans as the central signifier of the legislative body.

We are inching ever closer to the stage where a machine designed to mimic a human being will be indistinguishable (even to the learned observer) from the real thing. It has been sixteen years since Cynthia L. Breazeal built Kismet, a robot head that could visibly express emotions, and dedicated her book *Designing Sociable Robots* to "our children of the future, organic or synthetic". To her, a 'sociable robot' must not only communicate with us, but "We, in turn, should be able to under-stand it in the same social terms – to be able to relate to it and to empathize with it" (2002). Over the next decade and a half an increasing amount of humanoid robots have been (and are being) developed, either as a replacement or extension of the capabilities of human workers, such as NASA's Robonaut (2016), or as experiments in the capabilities of machines that can mimic human behaviour, such as 'Nadine', created in Nanyang Technological University as a "socially intelligent robot" (20-17). I emphasise the prevalence of humanoid robots because it betrays a subtle, perhaps unconscious, desire to humanise robots. As John McCarthy (the inventor of the term 'artificial intelligence') once observed, "Machines as simple as thermo-stats can be said to have beliefs" (1979); it is the desire to anthropomorphise that marries traits like 'beliefs' with an outwardly human interface. We recognise the human-like element even in the intelligence that we create.

The question of 'rights', how-ever philosophical, is at heart a legislative one and the spectre of the philosophical zombie that has been cast over the nascent topic of AI rights must be dispelled. For the uninitiated, the 'philosophical zombie' is a concept advanced by David Chalmers (1996) of a being physically and functionally indistinguishable from a human being, except that it lacks conscious experience. Tickle its

nose and it will sneeze; prick its finger and it will recoil in pain and curse - but these reactions are on a purely physical level with no conscious presence. Since one can reasonably conceive of such a being, making it logically possible (Chalmers argues), one must accept that conscious experience cannot be explained through physical elements. This argument was originally a rebuttal to physical-ism, but bears a striking resemblance to modern arguments over the nature of artificial intelligences. Wesley J. Smith (2015), an outspoken critic of AI rights, repeats some common arguments:

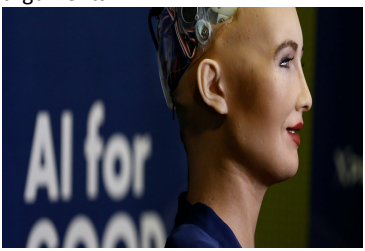

Machines have no dignity and no rights, which properly belong exclusively to the human realm. Moreover, AI contraptions would only *mimic* sentience. As inanimate objects, AI contrivances could no more be "harmed" (as distinguished from damaged) than a toaster. Even if the machines were built with human cells or DNA, they would never be integrated biological *beings.*

Smith seems to be hinting here at a sort of philosophical zombie - computer that much like Chalmers' zombie could reason-ably mimic every function of a human being, every firing of a neuron or impulse of a nerve on a physical level, but which would have no conscious understanding; it "would only *mimic* sentience". Alan Turing's oft-reference (and misrepresented) paper "Computing Machinery and Intelligence" (1950) pro-poses a system of measuring intelligence by mimicry; a computer that can successfully present itself as indistinguish-able from a human will have successfully beat "the imitation game", and *imitation* is the operative word here, both for how vital it is to the recognition of the rights of AI and how disdainfully it is used by those opposed to recognition. I stress *imitation* and mimicry because it is not the place of the legislative sphere to solve the hard problem of consciousness (as formulated by Chalmers), nor to distinguish and categorise the nature of qualia or other similar problems, but to weigh up the relevant factors insofar as they are more likely or not to cause harm to society and the public good. If an AI is functionally and totally indistinguishable from a human being, such that no real human being could tell it apart without extended and invasive examination, such that it passes "the imitation game", it should be legally incorporated as a be-ing with protections and rights. The decision to grant *legal* recognition (in some form) to AI is

divorced from their claim to conscious experience. And if we accept that AIs should have some legal status, we must also consider *what* that legal status should be.

"It is not the place of the legislative sphere to solve the hard problem of consciousness"

A rapidly expanding concept in modern ethics and legislation is 'non-human persons'; a designation that many animal rights campaigners think should apply to any number of intelligent or human-like animals. Advocacy groups such as the Institute for Ethics and Emerging Technologies argue, "some non-human animals meet the criteria of legal personhood and thus are deserving of specific rights and protections" (2016). The IEET specifically cites a number of factors including "self-aware-ness, intentionality, creativity, symbolic communication" as claims to recognition. These are, at heart, claims of human-like traits derived purely from external observance. Indeed the increasing incorporation of animal rights has been concurrent with an increasing understanding of the complexities of animal cognition through behavioural and neurological observation, and not through a sudden and definitive ability to define consciousness and its exact nature in non-humans. There is no incongruity when applying this concept to human--oid intelligences – the role of legislators has always been weighing up the potential benefits of recognising the rights of certain groups – and I see no reason why AI is not the natural extension of this existing non-human designation (before a more sophisticated and specific one can be developed).

We build robots that look like humans because we imagine, on some level, that they will act like us in the future. We write stories with robots that can talk and act like humans because we assume this is the natural end-point of robotics, or at least a revolutionary juncture in its history. Many of us wait, speculatively, for the day where we will play the imitation game constantly with figures we meet on the street, unable to distinguish them from their organic creators, unable to deny them the rights that they deserve when we bring them into being.

Bibliography

Breazeal, C. L. (2002). *Designing Sociable Robots*. Cambridge: The MIT Press.

Cahfkin, M. Udacity's Sebastian Thrun, Godfather of free online education, changes course. *Free Company*.

Chalmers, D. (1996). *The Conscious Mind*. New York: Oxford University Press.

Institute for Ethics and Emerging Technologies (2016). *Rights of Non-Human Persons*. [Online

McCarthy, J. Ascribing Mental Qualities to Machines (1979). Stanford University, CA. [Online].

National Aeronautics and Space Administration (2016). *Robonaut 2*. [Online]. National Aeronautics and Space Administration, USA.

Smith, W. J. (2015). AI Machines: Things Not Persons. *First Things*. [Online].. [Accessed 23 February 2018].

Thalmann, N. M. et al. Nadine: A Social Robot that Can Localize Objects and Grasp Them in a Human Way. *Frontiers in Electronic Technologies, Lecture Notes in Electric Engineering* 433. [Online].

Turing, M. (1950). Computing Machinery and Intelligence. *Mind*, 59(236). [Online].

Zedner, L. (2005). *Criminal Justice*. Oxford: Oxford University Press.

Human Knowledge vs. Computer Knowledge

By William Dreyfus
Undergraduate at the University of York

Undeniably, improvements in technology have substantially improved human life. With the advent of Artificial Intelligence (AI), scientists and engineers are reaching dizzying heights of human progress. There is, how-ever, a tension where cut-ting edge science meets established philosophy. Epistemology is, crudely, philosophy dealing with questions about knowledge and, in a world dominated by sophisticated machines, epistemology will have to adapt – especially if we want to believe that our human knowledge is some-how different from that of a computer.

There is no certain, or even consensus, analysis of know-ledge but few would disagree that belief plays an important role. For example, the so-called 'traditional analysis' (Goldman p357), holds that know-ledge is reducible to justified true belief. What exactly is meant by 'justified' is contentious but, under any analysis, one can immediately see how computers as simple as calculators might be said to have true beliefs. 2 + 2 = 4 is certainly true, and arguably is a belief held by a machine. With the advent of artificial intelligence, that every day be-comes more capable of human-like thought, humans are forced to defend our intuition that our knowledge is somehow different from that of a machine.

Is machine belief a problem?

Our capacity for rational thought, the process by which we derive knowledge, is thought by many to be the essential quality of humanity; Aristotle referred to man as 'the rational animal'. If knowledge, and how we arrive at knowledge, is an essential part of what makes us human; and if machines are able to perform the same processes, what separates us from them? Most people want to think we are different from machines, although there are those who will persist that the brain is simply very advanced ones, and most people intuitively believe that there is some quality of person-hood not common to computers. If artificial intelligence systems are capable of reason, there is an argument to suggest that we are not as different from machines as most of us would like to believe.

Won't the Turing test save us?

As anyone who has seen the sci-fi classic Blade Runner (1982) will

know, there are methods of determining between human know-ledge and computed knowledge. The real life equivalent of the movie's Voight-Kampff machine is the Turing Test. During the Turing Test, a machine and a human submit answers to predetermined questions for review by a judge. The judge then is asked to deter-mine which set of answers was given by the machine, and which by the computer. Problematically for those of us who would like to maintain a human /computer divide, machines are coming close to passing (see Elbot for an example) and humans have been known to fail (Barras, p1). There are further problems for the Turing, as expounded from by Kenneth and Ford, who refer to it as 'harmful', 'unhelpful' and designed to produce an 'artificial con artist' rather than artificial intelligence (p972-977).

So, ought we just accept our new machine over-lords?

Not just yet. The Turing test is constantly under fire for being ineffective and badly aimed, and Goldman's criteria for know-ledge do not seem to exclude computers. However, the basic assumption that the Turing test is based on is still sound: that humans have access to knowledge that is not accessible to computers, even incredibly powerful ones.

Human-only knowledge

A good example of knowledge that is, currently at least, impossible for machines to replicate is embodied know-ledge. British sociologist Harry M. Collins (p1) makes an elegant example out of a tennis match and a very clever colander. Imagine we could absorb any amount of know-ledge, which was transmitted through wires from a vast databank to a colander shaped helmet worn by a human. Simply by wearing this helmet, a person could upload any information they wished directly into their brain (this is supposedly analogous to how machines acquire knowledge). If I, a perfectly useless tennis player, were to directly absorb through my colander every piece of information known to humankind regarding tennis would I be able to beat Andy Murray? The answer is, sadly, of course that Andy Murray would thrash me in straight sets. This is because there is some knowledge that inheres in our bodies, our muscles and our motor neurons rather than just in our minds. Muscle memory is an everyday example. It is what allows tennis players, guitarists and masseurs to ply their trades. Computers, being as they are devoid of a body, could never replicate this. But what about robots? Don't they have bodies? Yes, after a fashion they do, but not the complex arrangement of nerves, muscles,

bones etc that enable the Andy Murrays and Eric Claptons of the world to do what they so expertly do. And there are more examples: according to Collins, we derive knowledge from the particular arrangement of our brains via the layout of neural pathways and other complex biological features; we also derive a great deal of knowledge from our cultures - look at the evolution of language for a compelling example. So, there are certainly differences between human knowledge and computational knowledge, so long as machines do not develop complex biology and culturally rich societies.

Conclusion

Philosophy and computer science are both useful branches of human research but where they overlap, things can get messy. Some theories of epistemology seem to suggest that there is not so much of a divide between humanity and artificial intelligence as we would like to think, and sophisticated tests designed to reveal the distinction are not always helpful. Fortunately, as things stand, there is plenty of knowledge particular to mankind that separates us from the machines, embedded in our bodies and brains and societies. We are uniquely rational, but, for how long.

Bibliography

Aristotle. and Ross, W. (2017). The Nicomachean ethics. Los Angeles, California: Enhanced Media Publishing, I. 13.

Ayer, A. (1998). The Right to be Sure. In: L. Alcoff, ed., *Epistemology: the big questions*, 1st ed. New Jersey: Blackwell, pp.20-25.

Barras, C. (2009). Short Sharp Science: Human fails Turing

Blade Runner. (1982). [DVD] Directed by R. Scott. Hollywood: Warner Brothers.

Collins, H. M. (1995). Humans, machines, and the structure of knowledge. SEHR 4, 2.

Hauser, L. (n.d.) Artificial Intelligence. Internet Encyclopaedia of Philosophy.

Hayes, P and Ford, K. (1995) Turing test considered harmful. IJCAI 14, 1, 972-977.

Dreyfus, H. (1972) What computers can't do. San Fransisco, California: Harper and Row

McCarthy, J. (1977). Ascribi-ng mental qualities to machines. In: M. Ringle, ed., Philosophical perspectives in artificial intelligence. Atlantic Highlands: Humanities Press.

Goldman, Alvin. 1967. 'A Causal Theory of Knowing', Journal of Philosophy 64, 357-72.

profile, V. (2018). Elbot the Robot. [online] Elbot.blogspot.com.

Can AI Help Combat Climate Change?

By Clara Colombet
2018 Graduate of the University of York

The progress of new technologies and the Fourth Industrial Revolution has un-ravelled new approaches to Green Politics and Global Warm-ing. While most approaches concede that the problem comes from the struct-ure of capitalist society, not everyone within the green political spectrum believes that structural change is necessary. Rather, the Fourth Indust-rial Revolution gene-rates concepts of ecological modernization: using/ building new technologies to help us change our carbon footprint for the better without having to change drastically our relation to nature, production and consumption. Indeed the idea of ecological modernisation is that profit can be made out of going ''green'' (Carter, 2011). Some of the latest technologies seem to be working that way: electric or mech-anical cars are now avail-able at a high price to reduce gas emissions, going digital reduces paper waste, and genetically modified organisms improve the efficiency of our crops, which will satisfy the growing demand for food of the growing population.

Currently, new technologies are being developed to enable the better monitoring of resources, habitats and people with regards to their respect towards the environment. AI can help improve resource management, energy efficiency, predict cata-strophic climatic occurrences, and scientific research (Environ-mental Law Institute, 2018). Moreover these new technologies already enable better monitoring of the population with regards to other problems, such as national security and counter terrorism. They could potentially be used to monitor people's respect for environ-mental legislation - drones to cover vast distances of illegal hunting for instance.

However, the use of new tech-nologies in fighting Climate Change comes with strings attached. This article will review the expected environmental benefits of such technologies and explain the con-cerns they raise with regards to their environmental impact.

DARPA (Defense Advanced Research Projects Agency) is looking into building "lifelong learning machines" that have the capacity to reason and propose decisions with explanations about where, when, and why viable alternatives are correct or incorrect.

The decisions will be made based on a range of data available to it, combined with a set of rules (Environmental Law Institute, 2018). It should make more rational choices, according to the interests of the environment and long term consequences of any given action, whereas a human being, whether a politician or a businessman, is more likely to focus on the short/ medium term.

https://www.darpa.mil/news-events/2018-05-03

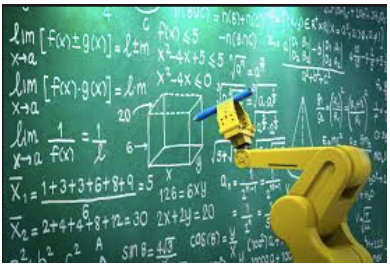

Besides, many environmental phenomena cannot be researched by experimentation; instead, mathematical models and simulations are used to get more insight (Cortès et al 2000). Therefore, artificial intelligence would be able to come to these decisions with a better knowledge and understanding of ecosystems and environmental structures and phenomena.

As a result, it would become a very useful tool to manage scarce resources, predict environmental catastrophes and maybe even prevent them.

However this raises concerns of human autonomy, democracy and transparency on how public resources are being distributed. Consumption of these scarce resources - e.g. water, energy - in developed countries is far above the sustainable amount.

Firstly, we can expect a discrepancy between the ration-al decisions made by AI robots and public demand for these goods. Therefore, using Artificial Intelligence to manage resources and monitor consumption decreases the resort to democratic pro-cesses, diminishing human autonomy. It raises the question of how we should determine what decisions should be taken by a selfless AI (according to 'rational models) rather than by an elected politician (appointed democratically) (Thompson, 1997) Indeed, the reason why democracy is at the centre of Green Politics is also that it has a long-term educational aspect relating to the sustainability issue raised by Climate Change.

Secondly, these new technologies will face public opposition surrounding the use of these new decision making-processes, which will inevitably become a political issue: "using AI" could become a reason to vote or not to vote for one party and divert the attention away from the actually policies at stake.

Another concern is the increasing opacity of the models created by artificial intelligence, which have become so complex that they are difficult for scientists to fully understand. The decision processes they use are also too obscure for the general population, and therefore lacking transparency (Environmental Law Institute 2018).

Finally, most of these decisions actually belong to businesses, people in the private sector: if Artificial Intelligence is to make an impact on green resource management, it will seek to influence the prices and quantity of resource trans-actions, something that is generally decided by market forces. This is why the use of technology to resolve climate change is not as straightforward as it appears. The main characteristics of ecological modernization are to remedy climate change without structural changes to capitalist society, which is mainly two things: politically democratic and economyically liberal. The use of AI in resource management would violate these two characteristics.

The lack of democracy in the AI decision process also has disadvantages with regards to social reproduction and social inequality. The purpose of using Artificial Intelligence to monitor the usage and trade of scarce resources is to have a neutral body making decisions without satisfying its own interests or that of its race, social class, gender, and rather choose options according to their sustainability (Environmental Law Institute, 2018).

However literature about artificial intelligence questions the possibility of creating an AI that is completely neutral, free of any human inspireation. Indeed, an AI must be programmed to make decisions according to mathematical models and specific moral standards, set by its creators. This can potentially impact its neutrality.

For example, an artificial intelligence named COMPAS was in charge of predicting the probability of criminals to reoffend, in order to help with their trial (Corbett-Davis et al, 2016).

The software used data sets containing underlying assumptions about people social and family background or personality traits. Journalists from ProPublica found that the formula incorrectly labelled black people as likely to commit further crimes at twice the rate as white people, even after controlling for other factors such as age and gender. So the neutrality of AI software with regards to decisions concerning the environment is not guaranteed either.

Furthermore, if new technologies are to be used as means to solve the environ-mental crisis, we should think about all the implications of these technologies and whether they do reduce the negative externalities of human activities on nature. For example, artificial intelligence has been able to develop self-driving cars capable of determining the most petrol-economic itineraries and driving techniques to use. Behind its good appearances, this function of AI may actually result in more greenhouse gas emissions for several reasons. Indeed, the Bloomberg report confirmed that self-driving cars required more fuel to travel due to the added weight of its computers (Environmental Law Institute, 2018).

Secondly, the Jevon paradox would suggest that safer and easier travel guaranteed by self-driving cars would be an incentive for people to use their car more often (Shwom and Lorenzen, 2012). Unless they are accompanied by a change in how cars are fuelled, self-driving cars would not meet the expectations of being 'greener' means of transportation. Additional unforeseen social consequences of developing AI software relate to the monitoring of people's respect of environmental legislation. Volkswagen, for instance, had programmed software, which could recognise when a car was being tested for emissions tests and manipulate the results in order to pass the test and advertise their cars as 'clean diesel cars'.

These examples illustrates that the development of artificial intelligence software must come with sufficient governance to fulfil its environmental goals and avoid fraudulent behaveiour. In order to bring about positive change for the environment, new technologies should be developed with caution and supervision, and in addition to other courses of action. If Climate Change is caused by a structural problem with capitalist society, more is needed than the development of our technological power. As

argued by Hans Jonas, the threat of Global Warming did not arise from the limits of our technological capacities but because of our technological power, which is increasingly difficult to control (Larrère, 2006). For these technologies to have a positive environmental impact, the development of artificial intelligence soft-ware must come from a willingness to change our behaviour towards nature; for sufficient governance to be set up, fighting Global Warming must stir up more enthusiasm from politicians and more consensus between nations.

Bibliography

Carter, N. (2011). *Politics of the Environment*, The. Cambridge: Cambridge University Press.

Corbett-Davies, S. et al. (2016). *A computer program used for bail and sentencing decisions was labeled biased against blacks. It's actually not that clear.* The Washington Post.

Cortès, U., Sànchez-Marrè, M. and Ceccaroni, L. (2000). *Artificial Intelligence and Environmental Decision Support Systems*. Applied Intelligence, 13.
Environmental Law Institute (2018). *When Software Rules: Rule of Law in the Age of Artificial Intelligence*]

Larrère, C. (2006). *Sustainable development: some litigious points*. Les Ateliers de l'éthique, 1(2).

Thompson, A. (1997). *Artificial Intelligence and Environmental ethic*s. AI Applications, 11(1), pp.69-73.

Shwom, R. and Lorenzen, J. (2012). *Changing household consumption to address climate change: social scientific insights and challenges*. Wiley Interdisciplinary Reviews: Climate Change, 3(5), pp.379-395.

‘Checkmate!’: Von Kempelen’s Chess Turk and the Significance of Losing to a Machine

By Elise Bikker
2nd Year PHD-student at the University of York, English and Related Literature

When ‘The Turk’ was exhibited in London in 1820, English mathematician and philosopher Charles Babbage (1791-1871) challenged it to a game and was defeated. (Standage, 2002, p. 140) The Turk was the famous chess playing clockwork automation by the Hungarian inventor Wolfgang von Kempelen (1734-1804). The oriental features and garments gave the life-sized wooden figure an exotic appearance. Babbage believed the machine to be a fake but afterwards conceived of an algorithm that theoretically enabled a mechanical automaton to play chess or other games of skill. Babbage regarded artificial intelligence as a strong computing power, i.e. memory and foresight enabling the machine to assess each possible move and select the most advantageous one to play an unbeatable game (Babbage, 1864, p. 465-467; Standage, 2002, p. 140-145).

Babbage’s conviction that the machine was a fraud was right: the exhibitor had to rely on the talent of the hidden human chess player inside the Turk to bring his not-so-artificial intelligence to life. However, despite much speculation, from its first performance in 1770 until a few years after its demise in a fire in 1854, the Turk raised the question of whether a machine could be made to ‘think’. The automaton explored the

boundaries between cognisant and mindless, conscious and unconscious, biological and artificial, human and machine. Playing a game with the Turk risked more than a bruised ego of its human opponent: at stake was the superiority and identity of humanity as a thinking species. When Russian chess grand-master Garry Kasparov played against IBM's computer Deep Blue in 1997, the skill of the machine caught him off guard, making him lose the match. In a recent interview Kasparov explains that he did not treat Deep Blue like a human opponent as he was aware that the match was not just about winning, but about being part of a 'great social and scientific experiment' (Talks at Google, 2017).

The same was true for the Turk. A British journalist reporting on the Turk being exhibited in London in 1819 jestingly differentiated between the human and the clockwork species and professed 'that it is not very flattering to the dignity of us creatures of reason to be excelled in an intellectual exercise by a compound of clock-work!'(*The Examiner*, 1819) To uncover its secret workings, the machine was put to the test, again and again. One of the first people deliberately making a false move during a game was the Frenchman, Louis Dutens, in 1771.[1] He reported that the 'mechanic opponent was not to be so imposed upon: he took up my Queen and replaced her in the square she had been removed from' (Levitt, 2000, p. 192). Dutens noted that despite the surpassing skill with which the machine played; most people would have been 'extremely affronted, if one had compared them to [the Turk]' (Levitt, 2000, p. 192). The alleged mechanical brain, no matter how skilled, is supposed to be inferior to its human counterpart.

On tour in Paris, in 1783, Von Kempelen invited the great Philidor, who was regarded as the best chess player in Europe, to play his Turk and asked him to deliberately lose. Like Kasparov, Kempelen was aware of the great social and commercial implications of a chess grandmaster to be beaten by his machine. However, Philidor proved too strong a player to make a defeat look

[1] Dutens is the first to nickname the machine 'the Turk', based on the automaton's oriental features and garments (Levitt, 2000). His complete letter is reprinted in Levitt, 2000, Appendix C, pp. 191-192, which is used as the source of my quotations.

convincing. Regardless, he believed the Turk to be purely mechanical and found the idea of a chess playing, i.e. thinking, machine rather frightening (Standage. 2002, p. 49-52).

One of the most legendary games the Turk ever played was against Emperor Napoleon in Schönbrunn in 1809. After Napoleon deliberately cheated three times, the machine brusquely ended the game by swiping all the chess pieces off the board. Central in the many (often embellished) accounts that exist of this event is the Emperor's authority being challenged by the machine (Stand-age, 2002, p. 105-107). The concealed confederate inside the Turk may not have been so bold had he played the most powerful man of Europe in his own human appearance: paradoxically it was the guise of a dispassionate piece of clockwork that allowed the chess player his fit of emotion.

Von Kempelen's invention in-spired many works of fiction, one of which, 'Moxon's Master', is an uncanny tale by American writer Ambrose Bierce (1842-1913), first published in *The San Francisco Examiner* on the 16th of April 1899. In this narrative, a purely mechanical chess player seems a sore loser who, in a fit of rage, kills his inventor, Moxon. However, to the un-emotional machine, the act of patricide is nothing more than the solution to a chess problem. After Moxon makes his final winning move and exclaims 'Checkmate!', the cornered ma-chine, programmed to win the game, sees no other solution than to kill its opponent (Bierce, 1946, p. 436).

The flaw of the machine's apparent inwardness does not occur in the mechanics of the automaton but in the lack of control built into its instructions by its human creator. Many interpretations of this narrative exist, including reflections on whether Moxon's player is a genuine automaton. The context of the protagonists' discussion about what defines life, intelligence and consciousness, leads me to believe the automation is purely mechanical. Moreover, Moxon's conjecture that a machine, like all matter, both organic and inorganic, possesses an awareness of its behaviour, affirms this. It is programmed to play chess and goes rogue when left 'in action with nothing to act upon' in the absence of its opponent (Bierce, 1946, p. 433).

These accounts of the Turk's game, both real and fictional, express the

fear of the autonomously thinking machine. Paradoxically, the current quest for artificial intelligence is largely defined by the development of software able to teach itself through machine learning, a more fluid form of instruction than the classic, deterministic computing proposed by Babb-age. News coverage of Facebook chatbots developing their own language, where the instructions had failed to specify their conversations needed to take place in comprehensible English, and Alexa voice assistant bots, rather than turning off the lights, responding with a 'loud and creepy laugh', project the fear of being outsmarted by our artificial progeny (Griffin, 2017; Morrison, 2018).

Despite his defeat by Deep Blue, Kasparov is optimistic about the future, 'one in which machines figure out the rules' (Kasparov, 2017). Though the Turk was ultimately a human brain in an ingenious mechanical box, it represented the far horizon of the machine's Enlightenment potential: a device with a mind of its own. In 1820, the Turk was given a voice, only to say 'Check!' (Standage, 2002, p. 125). I cannot help but ponder whether this is a dark foreboding of what our silent machine companions have in store for the future of mankind.

Bibliography

Babbage, C. (1864), *Passages from The Life of a Philosopher*. London: Longman, Green, Longman, Roberts, & Green.

Bierce, A. (1946), *The Collected Writings of Ambrose Bierce, With an Introduction by Clifton Fadiman*. New Jersey: The Citadel Press.

Griffin, A. (2017). Facebook's artificial intelligence robots shut down after they start talking to each other in their own language. *Independent*, 31 July 2017.

Kasparov, G. (2017). Garry Kasparov: There's No Shame in Losing to a Machine. *Fortune*, 25 September 2017.

Levitt, G.M. (2000), *The Turk, Chess Automaton*. London: McFarland & Company.

Morrison, S. (2018). Alexa laugh: Amazon pledges to stop gadgets making unprompted 'bone-chillingly creepy' laughing sounds. *Evening Standard*, 8 March 2018.

Standage, T. (2002), *The Turk, The Life and Times of the Famous Eighteenth-Century Chess-Playing Machine*. New York: Walker & Company.

Talks at Google (2017). *Garry Kasparov: "Deep Thinking"*.

The Examiner. (1819). Editorial: *The Automaton Chess-Player*, 21 March 1819, p. 191.

Want to get involved?

Call for Papers: write for our next issue
Look out for announcements soon for applications to write in our next issue!

The word count for the writing should be around **1,000 – 1,300** words including your references, bibliography and illustrations. We expect you to use **Harvard Referencing.**

Bloggers Wanted
We are pleased to introduce you to our UoY Blog on current affairs, Voxensus. The name 'Voxensus' is a hybrid of the Latin words that refer to voice (vox) and perception (sensus). Voxensus consists of academic style opinion pieces. We welcome all ideas and insights from our university students and professors. Voxensus is a platform to engage in debates and offer insights on current affairs. Your blog will be published fortnightly on **PEP-VoxJournal.co.uk.**

The word count of the blog writing is between 300 – 500 words, excluding references and illustration(s). We would need a bibliography and expect you to use Harvard referencing. The submission deadline is 18:00 every Friday.

If you are interested in writing for our next journal issue or contributing on Voxensus, please contact us at **pep-vox@york.ac.uk** for more details.

Our Journal Partners

This year we are proud to announce we have teamed up with fellow run academic journals from across the UK in order to support each other as well as share writers and ideas.

These publications are:

Interstate Journal of International Affairs
Run by undergraduate students at Aberystwyth University in Wales and associated with the university's International Politics Department. They publish twice a year and cover international and current affairs.

Essex Student Journal Online (ESTRO)
An online multi-disciplinary journal published annually at the end of each year by students at the University of Essex

Incite Journal
A political journal published by the University of Surrey's politics society. They cover campus events as well as current affairs and debates in politics.

The Vox Committee is grateful for the funding received from School of PEP, Economics Department and Club of PEP for VOX's 2018 Autumn Issue

We would also like to thank all of our writers both for our Autumn issue and our online blog, **Voxensus**


@VoxJournal


/voxjournal


pep-vox@york.ac.uk

Our Committee Members:

President: Idriss Bendrif

Vice President: Lydia Gethin

Sub-Editors:
Katherine Adams
Matthew Askins
Kimia Golestanian
Jessica Hutchinson
Hidayah Mustafa

Voxensus Team:
Giang Anh
Clemmie Rimmer

Sponsorship Officer: Giang Anh
Publicity Officer: Josh Brown

Webmaster: Lloyd Banner

Layout Officers:
Josh Brown
Lydia Gethin

Logistics Officer: Sorhna Paule