Journal of Artificial Intelligence Research 64 (2019) 243-252 Submitted 02/18; published 02/19

Viewpoint: Human-in-the-loop Artificial Intelligence

Fabio Massimo Zanzotto FABIO.MASSIMO.ZANZOTTO@UNIROMA2.IT
University of Rome Tor Vergata
Viale del Politecnico, 1
00133, Roma ITALY

Abstract

Little by little, newspapers are revealing the bright future that Artificial Intelligence (AI) is building. Intelligent machines will help everywhere. However, this bright future may have a possible dark side: a dramatic job market contraction before its unpredictable transformation. Hence, in a near future, large numbers of job seekers may need financial support while catching up with these novel unpredictable jobs. This possible job market crisis has an antidote inside. In fact, the rise of AI is sustained by the biggest knowledge theft of the recent years. Many learning AI machines are extracting knowledge from unaware skilled or unskilled workers by analyzing their interactions. By passionately doing their jobs, many of these workers are shooting themselves in the feet.

In this paper, we propose Human-in-the-loop Artificial Intelligence (HitAI) as a fairer paradigm for AI systems. Recognizing that any AI system has humans in the loop, HitAI will reward these aware and unaware knowledge producers with a different scheme: decisions of AI systems generating revenues will repay the legitimate owners of the knowledge used for taking those decisions. As modern Merry Men, HitAI researchers should fight for a fairer Robin Hood Artificial Intelligence that gives back what it steals.

1. Introduction

We are on the edge of a wonderful revolution: Artificial Intelligence (AI) is breathing life into helpful machines, which will relieve us of our need to perform mundane activities. Self-driving cars (Lipson & Kurman, 2016; Lutin, Kornhauser, & Masce, 2013; Litman, 2014) are taking their first steps in our urban environment and their younger brothers, that is, assisted driving cars (Revathi & Dhulipala, 2012; Trajkovic, Colmenarez, Gutta, & Trovato, 2004; Gray, Ali, Gao, Hedrick, & Borrelli, 2012), are already a commercial reality. Robots are vacuum cleaning and mopping the floors of our houses (Ulrich, Mondada, & Nicoud, 1997; Taylor, Parker, Lau, Blair, Heninger, Ng, DiBernardo, Witman, & Stout, 2005; Huffman & Miner, 2008). Chatbots (Weizenbaum, 1966; Wallace, 2009) have conquered our new window-on-the-world – our smartphones – and, from there, they help with everyday tasks such as managing our agenda, answering our factoid questions or being our learning companions (Kerly, Hall, & Bull, 2007; Beccaceci, Fallucchi, Giannone, Spagnoulo, & Zanzotto, 2009). In medicine, computers can already help in formulating diagnoses (Austin, Tu, Ho, Levy, & Lee, 2013; Kourou, Exarchos, Exarchos, Karamouzis, & Fotiadis, 2015; Ferroni, Zanzotto, Scarpato, Riondino, Nanni, Roselli, & Guadagni, 2017) by looking at data doctors generally neglect. AI is preparing a wonderful future where people are released from the burden of repetitive jobs.

The bright AI revolution may have a possible dark side: a dramatic mass unemployment that may precede an unpredictable job market transformation. People and, hence, think tanks (Manyika, Chui, Miremadi, Bughin, George, Willmott, & Dewhurst, 2017; Stone, Brooks, Brynjolfsson, Calo, Etzioni, Hager, Hirschberg, Kalyanakrishnan, Kamar, Kraus, Leyton-Brown, Parkes, Press, Saxenian, Shah, Tambe, & Teller, 2016) and governments (Executive Office of the President of the United States of America, 2016) are frightened. A pessimistic report (Manyika et al., 2017) of the McKinsey Global Institute (MGI) foresees that AI may globally replace the equivalent of the activities of 1.1 billion employees by erasing $15.8 trillion in wages. By releasing people from repetitive jobs, intelligent machines may leave a majority of citizens with the value of their labor insufficient to pay for a socially acceptable standard of living (Stone et al., 2016). The revolution is ongoing. Chatbots are slowly replacing call center agents in some of their tasks. Self-driving trains are already reducing the number of drivers in our trains. Self-driving cars are fighting to replace cab drivers in our cities. Drones are expanding automation in managing delivery of goods by drastically reducing the number of delivery people. And, even more cognitive and artistic jobs are challenged. Intelligent machines may produce music jingles for commercials (Briot, Hadjeres, & Pachet, 2017), write novels, produce news articles and so on. Intelligent risk predictors may replace doctors (Austin et al., 2013; Kourou et al., 2015; Ferroni et al., 2017). Chatbots along with massive open online courses may replace teachers and professors (Gohd, 2017). Coders risk being replaced by machines too (Murphy, 2017). According to the White House report on AI (Executive Office of the President of the United States of America, 2016), this overwhelming progress of AI can initiate long-standing disruptions of local markets and, according to the MGI report (Manyika et al., 2017), nobody's job will be left unchanged.

Surprisingly, the rise of AI is largely supported by the knowledge of an unaware mass of people who risk seeing a large part of their wages canceled by machines. In fact, along with someone selling knowledge for peanuts with Amazon Mechanical Turk, Crowdflower or SurveyMonkey and along with those aware programmers who set up these intelligent machines, an unaware mass of people is providing precious training data by passionately doing their job – translating, interacting with customers, teaching – or simply performing their activities on the net – answering an email, interacting on messaging services, leaving an opinion on a hotel. These data are a goldmine for AI machines as learning systems transform these interactions in knowledge. By doing their normal everyday activity, many workers are shooting themselves in the foot and unaware people are "donating" their knowledge to machines. This is an enormous and legal knowledge theft taking place in our modern era.

As researchers in Artificial Intelligence, we have a tremendous responsibility: building intelligent machines that "support the parents of their intelligence" (Stone et al., 2016) rather than intelligent machines that steal their knowledge to do their jobs. Moreover, we need to find ways to financially support job seekers as they train to catch up with these novel unpredictable jobs. We need to prepare an antidote as we spread this poison in the job market.

This paper proposes *Human-in-the-loop Artificial Intelligence* (HitAI) as a novel paradigm for a responsible Artificial Intelligence. The idea is simple: *giving the right value to the knowledge producers*. Recognizing that any AI system has humans in the loop, HitAI promotes interpretable learning machines and, therefore, artificial intelligence systems with *a clear knowledge life cycle*. For HitAI systems, it will be clear whose knowledge has been used in a specific deployment or in specific situations. This is a way to give the rightful credit and revenue to the original knowledge producers. Hence, HitAI is a possible antidote to the poisoning of the job market.

The rest of the paper is organized as follows. Section 2 describes current trends and the enabling paradigms for Human-in-the-loop AI. Section 3 sketches a proposal for a better future. Then, Section 4 draws some conclusions.

2. Human-In-The-Loop AI: Trends and Enabling Paradigms

Nowadays, data seem to be the principal source of knowledge for machines. This section describes: (1) how data have become more important than programmers to *"teach"* machines (Sec. 2.1); (2) how explainable Artificial Intelligence (Sec. 2.2) and distributed representations for symbols (Sec. 2.3) can be used to understand how these machines learn from data. This description is extremely useful for proposing the agenda of Human-in-the-loop Artificial Intelligence.

2.1 Transferring Knowledge to Machines: from Programming to Autonomous Learning

Since the beginning of the digital era, *programming* is the preferred way to *"teach"* machines. *Artificial non-ambiguous programming languages* have been developed to have a clear tool to tell machines how to solve new tasks or how to be useful. According to this paradigm, whoever wants to *"teach"* machines has to master one of these programming languages. These people, called *programmers*, have been *teaching* machines for decades and have made these machines extremely useful. Nowadays, it is difficult to imagine passing a single day without using the big network of machines that programmers have contributed to building.

Not all the tasks can be solved by *programming*, so *autonomous learning* has been reinforced as an alternative way of controlling the *"behavior"* of machines. In autonomous learning, machines are asked to *learn from experience*. With the paradigm of *programming*, we have asked machines to *go to school* before these machines *have learned to walk through trial and error*. This is why machines have always been good in solving very complex cognitive tasks but very poor in working with everyday simple problems. The paradigm of *autonomous learning* has been introduced to solve this problem.

In these two paradigms, who should be paid for transferring knowledge to machines and how should they be paid? In the *programming paradigm*, roles are clear: *programmers* are the *"teachers"* and *machines* are the *"learners"*. Hence, programmers could be paid for their work while they are teaching machines that are learning. In the *autonomous learning paradigm*, the activity of programmers is confined to the selection of the most appropriate learning model and of the examples to show to these learning machines. Nobody is paid while machines are learning.

From the point of view of HitAI, the trend of shifting from *programming* to *autonomous learning* is dangerous. In fact, *programming* is a fair paradigm as it keeps humans in the loop. On the contrary, *autonomous learning* is an unfair model of transferring knowledge as the real knowledge is extracted from data produced by unaware people. Hence, little seems to be done by humans and machines seem to do the whole job. Yet, knowledge is stolen without paying what it is worth.

2.2 Explainable Autonomous Learning Machines

Explaining the decisions of *autonomous learning* machines is a very hot topic nowadays: dedicated workshops or specific sessions in major conferences are flourishing (Aha, Darrell, Pazzani, Reid, Sammut, & Stone, 2017; Kim, Malioutov, Varshney, & Weller, 2017). In specific areas of application, for example, medicine, trust in intelligent machines cannot be blind as final decisions can have a deep impact on humans. Hence, understanding why a decision is taken becomes extremely important. However, what is exactly an explainable machine learning model is still an open debate (Lipton, 2016).

In HitAI, explainable machine learning can play a crucial role. In fact, seen from another perspective, explaining machine learning decisions can keep humans in the loop in two ways: 1) giving the last word to humans; and, 2) explaining what data sources are responsible for the final decision. In the first case, the decision power is left in the hands of very specialized professionals who use machines as advisers. This is a clear case of human-in-the-loop AI. Yet, this is confined to highly specialized knowledge workers in some specific areas. The second case instead is fairly more important. In fact, machines that take decisions or work on a task are constantly using knowledge extracted from data. Spotting which data have been used for a specific decision or for a specific action of the machine is very important in order to give credit to whoever has produced these data. In general, data are produced by anyone and everyone, not only by knowledge workers. Hence, understanding why a machine takes a decision may become a way to keep everybody in the loop of AI.

2.3 Symbolic Knowledge and Distributed Representations in Learning Machines

In current AI systems, knowledge is stored in tensors of real numbers called *distributed representations*. These representations are pushing learning models (LeCun, Bengio, & Hinton, 2015; Schmidhuber, 2015) towards amazing results in many high-level tasks such as image recognition (He, Zhang, Ren, & Sun, 2016; Simonyan & Zisserman, 2015), image generation (Goodfellow, Pouget-Abadie, Mirza, Xu, Warde-Farley, Ozair, Courville, & Bengio, 2014), image captioning (Vinyals, Toshev, Bengio, & Erhan, 2015b; Xu, Ba, Kiros, Cho, Courville, Salakhudinov, Zemel, & Bengio, 2015), machine translation (Bahdanau, Cho, & Bengio, 2014; Zou, Socher, Cer, & Manning, 2013), syntactic parsing (Vinyals, Kaiser, Koo, Petrov, Sutskever, & Hinton, 2015a; Weiss, Alberti, Collins, & Petrov, 2015) and even game playing at a human level (Silver, Huang, Maddison, Guez, Sifre, Van Den Driessche, Schrittwieser, Antonoglou, Panneershelvam, Lanctot, et al., 2016; Mnih, Kavukcuoglu, Silver, Graves, Antonoglou, Wierstra, & Riedmiller, 2013).

Explaining AI decisions seems simple when learning machines treat images as this knowledge is stored similarly to *distributed representations*. For example, in neural networks, input images and layers of the networks are tensors of real numbers. Interpreting these networks is generally done by visualizing how layers represent salient subparts of target images. Hence, these networks can be examined and understood.

However, a large part of the knowledge is not expressed in images but with symbols, which apparently are not similar to distributed representations. Both in natural and artificial languages, combinations of symbols are used to convey knowledge. In fact, for natural languages, sounds are transformed into letters or ideograms and these symbols are combined to produce words. Words then form sentences and sentences form texts, discourses, dialogs, which ultimately convey knowledge, emotions, and so on. Hence, to explain decisions of learning machines, we need to understand how symbolic knowledge is represented in distributed representations.

For HitAI, there is a tremendous opportunity to track how symbolic knowledge flows in the knowledge life cycle. Although symbols seem to fade away in current AI systems, there is a strict link between distributed representations and symbols, the first being an approximation of the second (Fodor & Pylyshyn, 1988; Plate, 1995; Zanzotto, Ferrone, & Baroni, 2015; Ferrone & Zanzotto, 2017). In this way, symbolic knowledge producers can also be rewarded for their unaware work.

3. Human-In-The-Loop AI: A Simple Proposal for a Better Future

The AI revolution is largely based on an enormous knowledge theft, which will possibly be a never-ending source of revenues for companies. In fact, skilled and unskilled workers do their own everyday jobs and leave important traces. These traces are the *training examples* that machines can use to learn. Hence, AI is stealing these workers' knowledge by learning from their interactions. The *stolen* knowledge is going to produce never-ending revenues for companies for years. This is a major problem since only that very small fraction of the population who own shares of these companies can benefit from this never-ending revenue source and the real owners of the knowledge are not participating in this redistribution of wealth.

The model we propose with Human-in-the-loop Artificial Intelligence (HitAI) seeks to give back part of the revenues to the real owners – the knowledge producers. The key idea is that any profit-making interaction a machine does has to constantly repay whoever has produced the original knowledge used to do that interaction. Assigning rewards per decision is very important as it may be an incentive to produce better services today and to have better services in the future. In fact, people have an incentive to work better knowing that their future revenues depend on how they treat difficult and odd cases today.

Realizing HitAI poses two big challenges: the *"political"* challenge of convincing companies to share benefits with producers of the data; the *"infrastructural"* challenge of managing ownership of knowledge in the *knowledge life cycle*. These are two different, yet interrelated challenges.

The *"political"* challenge is very tough: convincing companies to share the benefits of AI technologies may be impossible. But, if a free society cannot help the many who are poor, it cannot save the few who are rich (Kennedy, 1961). In fact, in the long run, if there is not a way to redistribute benefits to knowledge producers – *the poor people* –, the overall market will collapse and, hence, companies – *the rich people* – will also lose these benefits. Although reasonable, this may not be an argument for today's CEOs focused on the short term.

Hence, if convincing companies is impossible, companies should be forced to share benefits. A possibility to achieve this is to start by protecting personal data with two mechanisms:

- a legal mechanism: protecting unaware knowledge production by extending the copyright law

- a technological mechanism: promoting new "ownership-aware file systems" which release and accept data only if owners are specified.

Both mechanisms should be promoted by governments. The legal mechanism is very slow as it has to go from the national level to the international level. The technological mechanism can be faster as it may be funded by local government grants or by spontaneous social movements such as the nordic model of MyData[1] and, then, spread all over as a novel concept of ownership-aware file system strongly required by final users.

The *"infrastructural"* challenge is different but again difficult: HitAI needs technologies that produce a trusted knowledge life cycle of AI systems in order to reward knowledge producers. Managing ownership in the knowledge life cycle poses major technological and moral issues and it is certainly more complex than *simply using knowledge while forgetting its source*. Each interaction has to be tracked and assigned to a specific individual.

To build a trusted knowledge life cycle of AI systems, we need to investigate two core problems: first, building *Knowledge Life Cycle Transparent Artificial Intelligence* systems and, second, building *Trusted Technologies*.

Indeed, building *Knowledge Life Cycle Transparent Artificial Intelligence* is mandatory because, in order to reward knowledge producers, systems need to exactly know who is responsible for a specific decision. *Knowledge Life Cycle Transparent AI* may share techniques with *explainable AI* (Sec. 2.2) but its focus is different: identifying the causal source of each inference of AI systems. For Neural Networks (Haykin, 1998) (NN), it may seem impossible to identify the causal source of an inference. This is a research topic. Yet, the causal source of an inference can be determined in many machine learning models. In support vector machines (Cortes & Vapnik, 1995; Cristianini & Shawe-Taylor, 2000), it is possible to determine which support vector is participating to each single decision and which is its "role", that is, its weight, with respect to that decision. In decision tree learning (Quinlan, 1993), each decision node is justified by a set of examples. Hence, even if today it seems difficult to trace back causal source of an output in NN, this can be a promising field of research.

Designing *Trusted Technologies* is the second core problem as the knowledge life cycle should be clear and knowledge items correctly tracked. This core problem is largely linked to the new *"ownership-aware file system"*, which has been mentioned above. However, HitAI needs to guarantee that data in circulation have specific owners without revealing who the owner is. This is also required by the General Data Protection Regulation (GDPR) (European Parliament and European Council, 2016). Hence, HitAI will use *Digital Identity Protocols and Mechanisms* (Camp, 2004) to identify people but, at the same time, HitAI will ensure the use of *Privacy Preserving Protocols and Mechanisms*, which can be obtained by using data encryption and *Block Chains* (Nakamoto, 2008; Tschorsch & Scheuermann, 2016).

HitAI is an alternative to other solutions of wealth redistribution like Universal Basic Income (UBI) (Straubhaar, 2017), which may not be viable (Zheng, Guerriero, Lopez, & Haverman, 2017). In fact, UBI should be governed at the national level through taxes. But, generally, companies act in a transnational level paying taxes on revenues where it is more convenient. Hence, UBI may not be easy to apply.

1. https://mydata.org/

With a small fraction of the resources needed for UBI, government grants can instead help HitAI to grow a fairer society by winning those difficult *"political"* and *infrastructural* challenges for which the AI field does not have answers right now.

4. Conclusions

The bright Artificial Intelligence revolution has a dark side: a possible, dramatic job market contraction before its unpredictable transformation. A peasant or, even, a wise politician of the late 19th century would have never imagined that after 100 years *yoga trainer*, *pet caretaker* and *ayurveda massage therapist* were common jobs. Today, the situation is similar with a complication: the speed of the AI revolution. It's hard to imagine what's next on the job market and, hence, what are the skills needed for being part of the labor force of the future. We urge a strategy for the immediate future to mitigate the dark side of the AI revolution.

In this paper, we proposed Human-in-the-loop Artificial Intelligence (HitAI) as a fairer AI approach, which leverages the most gigantic knowledge theft of the modern era. In fact, unaware skilled and unskilled knowledge workers are shooting themselves in their feet by passionately doing their normal, everyday work as these workers are producing the knowledge which Artificial Intelligence is making a profit on. HitAI aims to give back a large part of this profit to its legitimate owners. As modern Merry Men, HitAI researchers should fight for a fairer Robin Hood Artificial Intelligence that gives back what it steals.

References

Aha, D. W., Darrell, T., Pazzani, M., Reid, D., Sammut, C., & Stone, P. (Eds.). (2017). *Proceedings of the IJCAI Workshop on Explainable Artificial Intelligence (XAI)*.

Austin, P. C., Tu, J. V., Ho, J. E., Levy, D., & Lee, D. S. (2013). Using Methods from the Data-Mining and Machine-Learning Literature for Disease Classification and Prediction: A Case Study Examining Classification of Heart Failure Subtypes. *Journal of Clinical Epidemiology*, *66*(4), 398–407.

Bahdanau, D., Cho, K., & Bengio, Y. (2014). Neural Machine Translation by Jointly Learning to Align and Translate. *arXiv preprint arXiv:1409.0473*.

Beccaceci, R., Fallucchi, F., Giannone, C., Spagnoulo, F., & Zanzotto, F. M. (2009). Education with '"Living Artworks" in Museums. In *Proceedings of the International Conference on Computer Supported Education (CSEDU)*, pp. 346–349.

Briot, J.-P., Hadjeres, G., & Pachet, F. (2017). Deep Learning Techniques for Music Generation-A Survey. *arXiv preprint arXiv:1709.01620*.

Camp, J. (2004). Digital Identity. *IEEE Technology and society Magazine*, *23*(3), 34–41.

Cortes, C., & Vapnik, V. (1995). Support Vector Networks. *Machine Learning*, *20*, 1–25.

Cristianini, N., & Shawe-Taylor, J. (2000). *An Introduction to Support Vector Machines and Other Kernel-based Learning Methods*. Cambridge University Press.

European Parliament and European Council (2016). Regulation (EU) 2016/679 on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data, and Repealing Directive 95/46/EC (General Data Protection Regulation). *Official Journal of the European Union*, *L119*, 1–88.

Executive Office of the President of the United States of America (2016). *Artificial Intelligence, Automation, and the Economy.*

Ferrone, L., & Zanzotto, F. M. (2017). Symbolic, Distributed and Distributional Representations for Natural Language Processing in the Era of Deep Learning: a Survey. *CoRR*, *abs/1702.00764*.

Ferroni, P., Zanzotto, F. M., Scarpato, N., Riondino, S., Nanni, U., Roselli, M., & Guadagni, F. (2017). Risk Assessment for Venous Thromboembolism in Chemotherapy-Treated Ambulatory Cancer Patients: A Machine Learning Approach. *Medical Decision Making*, *37*(2), 234–242.

Fodor, J. A., & Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition*, *28*(1), 3 – 71.

Gohd, C. (2017). Your Next Teacher Could Be a Robot. https://futurism.com/are-robotic-teachers-the-future/.

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative Adversarial Nets. In *Advances in Neural Information Processing Systems*, pp. 2672–2680.

Gray, A., Ali, M., Gao, Y., Hedrick, J., & Borrelli, F. (2012). Semi-Autonomous Vehicle Control for Road Departure and Obstacle Avoidance. *IFAC Control of Transportation Systems*, 1–6.

Haykin, S. (1998). *Neural Networks: A Comprehensive Foundation* (2nd edition). Prentice Hall PTR, Upper Saddle River, NJ, USA.

He, K., Zhang, X., Ren, S., & Sun, J. (2016). Identity Mappings in Deep Residual Networks. In Leibe, B., Matas, J., Sebe, N., & Welling, M. (Eds.), *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 630–645, Cham. Springer International Publishing.

Huffman, E. C., & Miner, J. L. (2008). Robotic Sweeper Cleaner with Dusting Pad.. US Patent 7,346,428.

Kennedy, J. F. (1961). Inaugural Address of the President of United States.. Washington, D.C.

Kerly, A., Hall, P., & Bull, S. (2007). Bringing Chatbots into Education: Towards Natural Language Negotiation of Open Learner Modelsmodels. *Knowledge-Based Systems*, *20*(2), 177–185.

Kim, B., Malioutov, D. M., Varshney, K. R., & Weller, A. (Eds.). (2017). *Proceedings of the ICML Workshop on Human Interpretability in Machine Learning (WHI).*

Kourou, K., Exarchos, T. P., Exarchos, K. P., Karamouzis, M. V., & Fotiadis, D. I. (2015). Machine Learning Applications in Cancer Prognosis and Prediction. *Computational and Structural Biotechnology Journal*, *13*, 8–17.

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep Learning. *Nature*, *521*(7553), 436–444.

Lipson, H., & Kurman, M. (2016). *Driverless: Intelligent Cars and the Road Ahead*. MIT Press.

Lipton, Z. C. (2016). The Mythos of Model Interpretability. In *Proceedings of ICML Workshop on Human Interpretability of Machine Learning*, Vol. abs/1606.03490.

Litman, T. (2014). Autonomous Vehicle Implementation Predictions. *Victoria Transport Policy Institute*, *28*.

Lutin, J. M., Kornhauser, A. L., & Masce, E. L.-L. (2013). The Revolutionary Development of Self-Driving Vehicles and Implications for the Transportation Engineering Profession. *Institute of Transportation Engineers Journal*, *83*(7), 28.

Manyika, J., Chui, M., Miremadi, M., Bughin, J., George, K., Willmott, P., & Dewhurst, M. (2017). *A Future That Works: Automation, Employment, and Productivity*. McKinsey Global Institute.

Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., & Riedmiller, M. (2013). Playing Atari with Deep Reinforcement Learning. *arXiv preprint arXiv:1312.5602*.

Murphy, M. (2017). Ginni Rometty on the End of Programming. *Bloomberg Businessweek*.

Nakamoto, S. (2008). Bitcoin: A Peer-To-Peer Electronic Cash System..

Plate, T. A. (1995). Holographic Reduced Representations. *IEEE Transactions on Neural Networks*, *6*(3), 623–641.

Quinlan, R. J. (1993). *C4:5: Programs for Machine Learning*. Morgan Kaufmann, San Mateo.

Revathi, G., & Dhulipala, V. S. (2012). Smart Parking Systems and Sensors: A Survey. In *Proceedings of International Conference on Computing, Communication and Applications (ICCCA)*, pp. 1–5. IEEE.

Schmidhuber, J. (2015). Deep Learning in Neural Networks: An Overview. *Neural Networks*, *61*, 85–117.

Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., et al. (2016). Mastering the Game of Go with Deep Neural Networks and Tree Search. *Nature*, *529*(7587), 484–489.

Simonyan, K., & Zisserman, A. (2015). Very Deep Convolutional Networks for Large-Scale Image Recognition. In *Proceedings of the International Conference on Learning Representations (ICLR)*.

Stone, P., Brooks, R., Brynjolfsson, E., Calo, R., Etzioni, O., Hager, G., Hirschberg, J., Kalyanakrishnan, S., Kamar, E., Kraus, S., Leyton-Brown, K., Parkes, D., Press, W., Saxenian, A., Shah, J., Tambe, M., & Teller, A. (2016). Artificial Intelligence and Life in 2030. Tech. rep., One Hundred Year Study on Artificial Intelligence: Report of the 2015-2016 Study Panel, Stanford University, Stanford, CA.

Straubhaar, T. (2017). On the Economics of a Universal Basic Income. *Intereconomics*, *52*(2), 74–80.

Taylor, C., Parker, A., Lau, S., Blair, E., Heninger, A., Ng, E., DiBernardo, E., Witman, R., & Stout, M. (2005). Robot Vacuum Cleaner.. US Patent 11,171,031.

Trajkovic, M., Colmenarez, A. J., Gutta, S., & Trovato, K. I. (2004). Computer Vision Based Parking Assistant.. US Patent 6,683,539.

Tschorsch, F., & Scheuermann, B. (2016). Bitcoin and Beyond: A Technical Survey on Decentralized Digital Currencies. *IEEE Communications Surveys Tutorials*, *18*(3), 2084–2123.

Ulrich, I., Mondada, F., & Nicoud, J.-D. (1997). Autonomous Vacuum Cleaner. *Robotics and Autonomous Systems*, *19*(3-4), 233–245.

Vinyals, O., Kaiser, L. u., Koo, T., Petrov, S., Sutskever, I., & Hinton, G. (2015a). Grammar as a Foreign Language. In Cortes, C., Lawrence, N. D., Lee, D. D., Sugiyama, M., & Garnett, R. (Eds.), *Advances in Neural Information Processing Systems 28*, Vol. 2, pp. 2773–2781. Curran Associates, Inc.

Vinyals, O., Toshev, A., Bengio, S., & Erhan, D. (2015b). Show and Tell: A Neural Image Caption Generator. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3156–3164.

Wallace, R. S. (2009). *The Anatomy of A.L.I.C.E.*, pp. 181–210. Springer Netherlands, Dordrecht.

Weiss, D., Alberti, C., Collins, M., & Petrov, S. (2015). Structured Training for Neural Network Transition-Based Parsing. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 323–333. Association for Computational Linguistics.

Weizenbaum, J. (1966). ELIZA-A Computer Program for the Study of Natural Language Communication Between Man and Machine. *Communications of the ACM*, *9*(1), 36–45.

Xu, K., Ba, J., Kiros, R., Cho, K., Courville, A., Salakhudinov, R., Zemel, R., & Bengio, Y. (2015). Show, Attend and Tell: Neural Image Caption Generation with Visual Attention. In Bach, F., & Blei, D. (Eds.), *Proceedings of the 32nd International Conference on Machine Learning*, Vol. 37 of *Proceedings of Machine Learning Research*, pp. 2048–2057, Lille, France. PMLR.

Zanzotto, F. M., Ferrone, L., & Baroni, M. (2015). When the Whole is Not Greater Than the Combination of Its Parts: A "Decompositional" Look at Compositional Distributional Semantics. *Comput. Linguist.*, *41*(1), 165–173.

Zheng, Y., Guerriero, M., Lopez, E. V., & Haverman, P. (2017). Universal Basic Income (UBI): A Policy Option for China beyond 2020?. Tech. rep., UNDP China Office.

Zou, W. Y., Socher, R., Cer, D. M., & Manning, C. D. (2013). Bilingual Word Embeddings for Phrase-Based Machine Translation. In *Proceedings of the Empirical Natural Language Processing Conference (EmNLP)*, pp. 1393–1398.