

Repositório ISCTE-IUL

Deposited in *Repositório ISCTE-IUL*:

2024-04-24

Deposited version:

Accepted Version

Peer-review status of attached file:

Peer-reviewed

Citation for published item:

Geadá, N. (2024). Artificial intelligence: Applications of AI in turbulent times. In Nuno Geadá, George Leal Jamil (Ed.), *Perspectives on artificial intelligence in times of turbulence: Theoretical background to applications*. (pp. 1-12). Hershey: IGI Global.

Further information on publisher's website:

10.4018/978-1-6684-9814-9.ch001

Publisher's copyright statement:

This is the peer reviewed version of the following article: Geadá, N. (2024). Artificial intelligence: Applications of AI in turbulent times. In Nuno Geadá, George Leal Jamil (Ed.), *Perspectives on artificial intelligence in times of turbulence: Theoretical background to applications*. (pp. 1-12). Hershey: IGI Global., which has been published in final form at <https://dx.doi.org/10.4018/978-1-6684-9814-9.ch001>. This article may be used for non-commercial purposes in accordance with the Publisher's Terms and Conditions for self-archiving.

Use policy

Creative Commons CC BY 4.0

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a link is made to the metadata record in the Repository
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Introduction

Artificial intelligence (AI) is a research area that has been growing exponentially in the last decades and has revolutionized various sectors of society, from healthcare to industry. In times of turbulence, such as economic crises, pandemics, or natural disasters, AI can play an even more important role, providing insights and solutions that can help mitigate the impacts of these crises. This paper aims to explore the perspectives of artificial intelligence in times of turbulence by examining the applications, challenges, and opportunities associated with the use of AI in different sectors. Furthermore, recent advances in the field of AI and how they can be applied to help cope with crises and turbulent situations will be discussed. Based on the literature review, the main applications of AI in times of turbulence are identified, highlighting its contribution in areas such as healthcare, finance, manufacturing, retail, and education. The ethical and social challenges associated with the use of AI are also discussed, including the possibility of perpetuating existing prejudices and discrimination, as well as the potential impact on the labour market. Finally, the opportunities AI offers in terms of crisis mitigation are explored, from predicting market fluctuations to improving the customer experience in retail. This paper concludes that while there are significant challenges to be addressed, AI has the potential to bring many benefits in times of turbulence and should be carefully explored and used ethically and responsibly.

Applications of AI in Turbulent Times

AI has been widely used in different sectors in turbulent times. In healthcare, AI has been used to diagnose diseases, identify infection patterns, and develop treatments and vaccines for infectious diseases such as COVID-19. For example, AI has been used to analyse chest CT scan images of patients with COVID-19 to help detect and monitor disease progression. In finance, AI has been used to predict market fluctuations, detect fraud, and manage risk. AI can help analyse large volumes of financial data to identify trends and patterns, allowing investors to make more informed decisions. In addition, AI can be used to detect suspicious activity, such as fraudulent transactions, and help reduce the risk of financial losses (Islam, M. et al 2020). In manufacturing, AI has been used to improve production efficiency and reduce costs. AI can help optimize production by reducing downtime and increasing machine efficiency. In addition, AI can be used to predict problems in production before they occur, allowing preventative measures to be taken to minimize costs.

In the retail industry, AI has been used to improve the customer experience and optimize retail operations (Liao, T. W., & Hsieh, H. P. 2020). AI can be used to personalize the customer experience by offering product recommendations based on the customer's purchase history and browsing behaviour. In addition, AI can be used to optimize inventory management, forecast demand, and reduce product waste. In education, AI has been used to personalize learning and improve student performance. AI can be used to provide personalized learning resources, tailoring content and delivery to individual student learning styles and needs. In addition, AI can be used to assess student performance, providing personalized feedback, and helping to identify areas where the student may need more support.

However, it's important to note that the application of AI in turbulent times is not without its challenges. One major concern is the potential for AI systems to perpetuate or amplify existing biases and inequalities, particularly in industries such as healthcare and finance. Additionally, the use of AI in critical decision-making processes, such as hiring or criminal justice, raises ethical concerns about transparency and accountability. Overall, the applications of AI in turbulent times are vast and varied, offering new opportunities for businesses and organizations to navigate challenges and adapt to changing circumstances. As the technology continues to advance, it will be essential to carefully consider its potential impacts and limitations to ensure that AI is used responsibly and for the benefit of society as a whole (Mittal, R., & Mathew, R. 2019).

Challenges and Opportunities

While AI has great potential to transform different areas and industries, it also presents significant challenges that need to be addressed to maximize its benefits and minimize its risks. One of the main challenges of AI is algorithmic bias. Machine learning algorithms are often trained on datasets that reflect societal attitudes and biases. This can lead to biased and unfair results, such as criminal justice systems that discriminate based on race or gender (Angwin et al., 2016). To avoid these problems, it is necessary to ensure that the datasets used to train algorithms are representative and inclusive.

Another challenge is data privacy. Many AI systems rely on access to large amounts of personal data, such as medical records or purchase histories. If this data is compromised, this can lead to security and privacy issues. It is therefore important to implement robust security measures and ensure the privacy of user data. In addition, AI also presents challenges for the labour market. As more tasks are automated, it is possible that some jobs will become obsolete, leading to a shift in the labour market and possible social inequalities. However, there are also opportunities for creating new jobs and improving efficiency in the workplace (Manyika et al., 2017).

Another important issue is ethics in AI. It is necessary to ensure that AI systems are developed and used in an ethical and responsible manner. This includes transparency in algorithms, accountability, and ensuring that AI is used to complement and enhance human work, rather than replace it. Despite these challenges, AI also offers many exciting opportunities. In healthcare, for example, AI can be used to help diagnose diseases (Liu et al., 2020) and analyse medical images (Wang & Wong, 2020). In manufacturing, AI can be used to improve production line efficiency and reduce material waste. In retail, AI can be used to personalize the customer experience and improve marketing strategies.

Artificial intelligence has the potential to be an extremely powerful tool in turbulent times. One of the areas where AI can be applied is in healthcare, helping to predict disease outbreaks and identify patients at risk. For example, in 2020, during the COVID-19 pandemic, researchers used AI to predict the spread of the disease by

analysing human mobility patterns and epidemiological data (Liu et al., 2020). AI has also been used to detect the disease in chest X-ray images, aiding in rapid and accurate diagnosis (Wang et al., 2020).

In finance, AI can be used to predict fluctuations in the market and provide insights into investments. A McKinsey Global Institute study (Manyika et al., 2017) showed that AI can increase productivity and efficiency in financial services by improving real-time decision-making and reducing operating costs. In manufacturing, AI can be used to optimize production and reduce costs by analysing real-time data to identify problems and implement solutions quickly and effectively. AI can also be used to monitor product quality and predict maintenance needs, helping to reduce downtime and maintenance costs. In the retail sector, AI can be used to improve the customer experience by personalizing product recommendations based on previous purchase data and providing automated customer support. AI can also be used to predict product demand, helping companies better manage their inventories and reduce waste. Despite the many opportunities offered by AI, there are also significant challenges to be faced. One of the main challenges is the possibility of perpetuating existing biases and discriminations, especially when AI models are trained on historical data that reflects these biases (Angwin et al., 2016). In addition, AI can have a significant impact on the labour market by replacing workers in repetitive tasks and reducing the demand for low-skilled workers. To address these challenges, it is necessary to develop ethical and responsible approaches to the use of AI. This includes adopting transparent algorithms and creating accountability systems that allow decisions made by AI systems to be explained and challenged (Russell & Norvig, 2010). It is also important to ensure that AI is used in a way that complements, rather than replaces, humans.

Convolutional Neural Networks (CNNs)

Convolutional Neural Networks (CNNs) are a specialized form of neural networks designed specifically for image processing. In recent years, CNNs have been widely used in computer vision applications such as image recognition and object detection. CNNs have the ability to learn to extract features from complex images, such as recognizing faces in photographs or identifying objects in a scene. This is possible because CNNs use convolutional layers that apply filters to extract specific features from images. The ability of CNNs to process large amounts of image data is particularly useful in areas such as medicine and biology, where large image datasets are common. For example, CNNs can be used to detect breast cancer in mammography images, improving diagnostic accuracy (Wang et al., 2020).

Convolutional Neural Networks (CNNs) are a type of artificial neural network that are commonly used in computer vision applications, such as image and video recognition. CNNs are designed to process data with a grid-like topology, such as images, by exploiting the spatial correlations between adjacent pixels. The basic building block of a CNN is a convolutional layer, which applies a set of filters or kernels to the input data. Each filter performs a convolution operation, where it slides over the input data and calculates a dot product between the filter weights and the input data at each location.

The output of each filter is a two-dimensional feature map that represents a specific feature of the input data, such as edges or corners. After the convolutional layer, the output is typically passed through a non-linear activation function, such as ReLU, to introduce non-linearity into the network. This is followed by a pooling layer, which reduces the spatial dimensions of the feature maps by aggregating adjacent pixels, typically by taking the maximum or average value. CNNs can have multiple convolutional and pooling layers, which allows the network to learn increasingly complex features. The final layer of the network is typically a fully connected layer, which takes the flattened output of the previous layer and outputs a probability distribution over the different classes in the classification problem. Training a CNN involves optimizing the weights of the filters in the convolutional layers using backpropagation, which involves calculating the gradient of the loss function with respect to the weights and updating them accordingly. CNNs have achieved state-of-the-art performance on a wide range of computer vision tasks, including image classification, object detection, and semantic segmentation. They have also been used in other domains, such as natural language processing and speech recognition. However, CNNs can be computationally expensive and require large amounts of training data. There are also challenges in designing the architecture of the network, such as determining the number and size of the filters and the depth of the network.

In addition to the basic architecture of CNNs, there have been many variations and extensions proposed to address specific challenges and improve performance on various tasks. Some examples include:

- ResNet (Residual Network): A deep CNN architecture that uses residual connections to enable training of much deeper networks without vanishing gradients.
- Inception: A CNN architecture that uses multiple filters of different sizes in parallel to capture both fine-grained and coarse-grained features.
- DenseNet (Dense Convolutional Network): A CNN architecture that uses dense connections between layers, where each layer receives inputs from all previous layers, to encourage feature reuse and facilitate gradient flow.

One notable application of CNNs is in image classification, where the goal is to assign a label to an input image from a fixed set of categories. The ImageNet Large Scale Visual Recognition Challenge (ILSVRC) has been a benchmark for evaluating the performance of image classification models, and CNNs have consistently outperformed other methods since the introduction of the challenge in 2010.

CNNs have also been applied to object detection, where the goal is to localize and classify objects in an image. One common approach is to use a region proposal network to generate candidate regions of interest, and then use a CNN to classify each region and refine the bounding box coordinates. Semantic segmentation is another task that CNNs have been used for, where the goal is to assign a class label to each pixel in an image. This requires the output of the CNN to have the same spatial dimensions as the input image, which can be achieved using up sampling or transposed convolution layers. In

summary, CNNs are a powerful type of neural network that have revolutionized computer vision tasks and have been applied to a wide range of other domains as well. While there are challenges and limitations to their use, ongoing research is constantly pushing the boundaries of what is possible with CNNs.

Reinforcement Learning (RL)

Reinforcement Learning (RL) is a machine learning technique that focuses on how agents (or computer programs) take actions in an environment to maximize a reward. RL has been successfully applied in games, such as chess and go, and in robotics, where agents are trained to perform complex tasks, such as walking or picking up objects. RL also has the potential to be applied in other areas, such as delivery route optimization, resource management in companies, and even air traffic control. However, the application of RL also presents significant challenges, such as the need to ensure that agents do not take harmful actions in a complex and constantly changing environment. RL is also a subfield of Artificial Intelligence (AI) that is concerned with how agents can learn to make decisions in an environment, in order to maximize some notion of cumulative reward. In RL, an agent interacts with an environment and takes actions based on observations of the environment's state, in order to learn a policy that maps states to actions, maximizing a reward signal. RL is inspired by the idea of how animals learn through trial and error, by receiving positive and negative feedback as they explore their surroundings. RL agents similarly explore their environment, learning through feedback from a reward signal that provides information about the quality of their actions.

RL has been successfully applied to a wide range of tasks, such as robotics, game playing, and recommendation systems. In robotics, RL has been used to train robots to perform complex tasks such as walking and grasping. In game playing, RL has been used to train agents to play games like Go and chess, achieving superhuman performance. In recommendation systems, RL has been used to learn personalized recommendations based on user behaviour. One of the challenges of RL is the exploration-exploitation trade-off. Agents need to explore the environment in order to learn, but at the same time, they also need to exploit what they have already learned in order to maximize rewards. This trade-off is often addressed using techniques such as epsilon-greedy exploration and optimistic initialization. Another challenge in RL is the credit assignment problem, which refers to the difficulty of attributing rewards to specific actions or decisions made by an agent. This problem is often addressed using techniques such as eligibility traces and credit assignment functions. RL algorithms can be categorized into model-based and model-free approaches. Model-based RL algorithms try to learn a model of the environment, including the transition dynamics and reward function, and use this model to plan actions. Model-free RL algorithms, on the other hand, directly learn a policy or value function, without explicitly modelling the environment, which is modelled as a Markov Decision Process (MDP). The agent observes the state of the environment, takes an action, and receives a reward based on the action taken and the new state of the environment. The goal of the agent is to learn a policy that maps states to actions, in order to maximize the cumulative reward received. There are several algorithms used in reinforcement learning, including Q-learning, SARSA, and Deep Reinforcement Learning. Deep Reinforcement Learning uses deep

neural networks to approximate the optimal policy and has been successful in applications such as playing Atari games and AlphaGo.

Generative Adversarial Networks (GANs)

Generative Adversarial Networks (GANs) are a type of neural network that consists of two parts: a generator and a discriminator. The generator is responsible for creating synthetic samples (e.g., face images), while the discriminator is responsible for distinguishing between real and synthetic samples. GANs have been widely used in image and video generation applications. For example, GANs can be used to create realistic images of faces that do not exist in real life. However, GANs also present significant challenges, such as the need to ensure that synthetic samples are not misleading or harmful. The concept of intelligence in AI refers to the ability of an AI system to perform tasks that would normally require human intelligence, such as understanding language, recognizing patterns, and solving problems. The measurement of intelligence in AI is a complex and multifaceted issue that has been the subject of much research and debate. The basic idea behind GANs is to have two neural networks compete with each other in a game. One network, called the generator, generates new data samples, while the other network, called the discriminator, tries to distinguish between the generated data and the real data. The generator is trained to create data that can fool the discriminator into thinking it is real, while the discriminator is trained to correctly identify whether a sample is real or generated. As both networks improve, the generator becomes better at generating realistic samples, and the discriminator becomes better at distinguishing between real and generated data.

One approach to measuring intelligence in AI is to use standardized tests, such as the Turing test. The Turing test, proposed by Alan Turing in 1950, involves a human evaluator who engages in a natural language conversation with a machine and a human, and tries to determine which is the machine and which is the human. If the machine is able to convince the evaluator that it is the human, then it is considered to have passed the Turing test and demonstrated a level of intelligence comparable to that of a human. Another approach to measuring intelligence in AI is to use performance metrics that are specific to a particular task or domain. For example, in the domain of image recognition, a commonly used metric is accuracy, which measures the percentage of images that are correctly classified by the AI system. In the domain of natural language processing, metrics such as precision, recall, and F1 score are often used to measure the performance of AI systems in tasks such as text classification and sentiment analysis. However, the use of standardized tests and performance metrics to measure intelligence in AI has been criticized for being too narrow and limited in scope. Critics argue that these measures do not fully capture the complexity and diversity of human intelligence, and that they may be biased towards certain types of tasks or domains. As a result, there has been a growing interest in developing more comprehensive and holistic measures of intelligence in AI, such as the General Problem Solver (GPS) and the Integrated Cognitive Architectures (ICAs). These approaches aim to capture the full range of

cognitive abilities that are required for intelligent behaviour, such as perception, reasoning, planning, and decision-making.

One of the major applications of GANs is in image generation. The generator network can learn to create new images that look like they could have come from the training set. For example, a GAN can be trained on a set of images of faces and can then generate new images of faces that look like they came from the same distribution as the training set. Another application of GANs is in data augmentation. GANs can be used to generate new training data that can be used to improve the performance of other machine learning models. For example, a GAN can be used to generate new images of handwritten digits that can be added to a training set for a digit recognition model. However, GANs are still a relatively new technology, and there are several challenges that need to be addressed. One of the major challenges is instability during training, which can cause the generator to produce low-quality samples or fail to converge. Another challenge is the mode collapse problem, where the generator produces a limited set of samples that fail to capture the full distribution of the training data. Generative Adversarial Networks is a powerful tool in machine learning that has the potential to revolutionize many fields. With further research and development, GANs could be used to create realistic images, videos, and even music, opening up new possibilities in art, entertainment, and other industries.

Conclusion

This paper presented an overview of the prospects of artificial intelligence in turbulent times, exploring the applications, challenges, and opportunities associated with the use of AI in different industries. We discussed the major applications of artificial intelligence in healthcare, finance, manufacturing, and retail, as well as the challenges AI presents, including biases and impact on the labour market. To address these challenges, an ethical and responsible approach to the use of AI is needed, which includes transparency in algorithms and accountability systems. In addition, it is important to ensure that AI is used to complement and enhance human work, rather than replace it. Overall, artificial intelligence has the potential to be a powerful tool to address turbulence in different sectors and areas, but it is also important to keep in mind its possible negative impacts and work to mitigate them. As the technology continues to develop, it is crucial to continue to explore the prospects of artificial intelligence and work to maximize its benefits while minimizing its risks. The measurement of intelligence in AI is a complex and multifaceted issue that involves the use of standardized tests, performance metrics, and more comprehensive approaches that aim to capture the full range of cognitive abilities required for intelligent behaviour. As AI continues to evolve and advance, it is likely that new and more sophisticated measures of intelligence will be developed. Artificial intelligence is a constantly evolving technology with the potential to transform many areas of human life. In recent years, there have been significant advances in areas such as computer vision, reinforcement learning, and generative adversarial networks.

References

Angwin, J., Larson, J., Mattu, S., & Kirchner, L. (2016). Machine Bias. ProPublica. <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>

Liu, X., Faes, L., Kale, A.U., et al. (2020). A comparison of deep learning performance against health-care professionals in detecting diseases from medical imaging: a systematic review and meta-analysis. *The Lancet Digital Health*, 2(6), e271-e297.

Manyika, J., Chui, M., Brown, B., Bughin, J., Dobbs, R., Roxburgh, C., & Byers, A. H. (2017). Harnessing automation for a future that works. McKinsey Global Institute. <https://www.mckinsey.com/featured-insights/future-of-work/harnessing-automation-for-a-future-that-works>

Russell, S., & Norvig, P. (2010). *Artificial intelligence: a modern approach*. Pearson Education.

Wang, L., Wong, A. (2020). COVID-Net: A Tailored Deep Convolutional Neural Network Design for Detection of COVID-19 Cases from Chest Radiography Images. arXiv preprint arXiv:2003.09871.

Silver, D., Huang, A., Maddison, C.J., Guez, A., Sifre, L., van den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S., Grewe, D., Nham, J., Kalchbrenner, N., Sutskever, I., Lillicrap, T., Leach, M., Kavukcuoglu, K., Graepel, T., & Hassabis, D. (2016). Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587), 484-489.

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.

Russell, S. J., & Norvig, P. (2010). *Artificial intelligence: a modern approach*. Prentice Hall.

Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., Chen, Y., Lillicrap, T., Hui, F., Sifre, L., van den Driessche, G., Graepel, T., & Hassabis, D.

Angwin, J., Larson, J., Mattu, S., & Kirchner, L. (2016). Machine Bias. ProPublica. <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>

Liu, X., Faes, L., Kale, A.U., et al. (2020). A comparison of deep learning performance against health-care professionals in detecting diseases from medical imaging: a systematic review and meta-analysis. *The Lancet Digital Health*, 2(6), e271-e297.

Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, 59(236), 433-460.

Bostrom, N. (2014). *Superintelligence: Paths, dangers, strategies*. Oxford University Press.

- Russell, S. J., & Norvig, P. (2010). *Artificial intelligence: A modern approach* (3rd ed.). Pearson.
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction* (2nd ed.). MIT Press.
- Kaelbling, L. P., Littman, M. L., & Moore, A. W. (1996). Reinforcement learning: A survey. *Journal of Artificial Intelligence Research*, 4, 237-285.
- 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*.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., ... & Petersen, S. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540), 529-533.
- Kober, J., Bagnell, J. A., & Peters, J. (2013). Reinforcement learning in robotics: A survey. *The International Journal of Robotics Research*, 32(11), 1238-1274.
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT press.
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436-444.
- Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems* (pp. 1097-1105).
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 770-778).
- Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., ... & Rabinovich, A. (2015). Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 1-9).
- Huang, G., Liu, Z., Van Der Maaten, L., & Weinberger, K. Q. (2017). Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 4700-4708).
- Radford, A., Metz, L., & Chintala, S. (2015). Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*.
- Arjovsky, M., Chintala, S., & Bottou, L. (2017). Wasserstein gan. *arXiv preprint arXiv:1701.07875*.
- Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., & Chen, X. (2016). Improved techniques for training GANs. In *Advances in Neural Information Processing Systems* (pp. 2234-2242).

Karras, T., Aila, T., Laine, S., & Lehtinen, J. (2017). Progressive growing of GANs for improved quality, stability, and variation. arXiv preprint arXiv:1710.10196.

Esteva, A., Robicquet, A., Ramsundar, B., Kuleshov, V., DePristo, M., Chou, K., ... & Dean, J. (2019). A guide to deep learning in healthcare. *Nature medicine*, 25(1), 24-29.

Heaton, J. (2017). *Deep learning with Python*. Birmingham, UK: Packt Publishing.

Islam, M. R., Islam, S. M. R., Asraf, A., & Khandakar, A. (2020). Artificial Intelligence in Finance: A Comprehensive Review. *IEEE Access*, 8, 89961-89988.

Kim, S., Jang, J., & Lee, S. (2020). Applications of artificial intelligence in transportation: A review. *Transportation Research Part C: Emerging Technologies*, 112, 631-654.

Liao, T. W., & Hsieh, H. P. (2020). Applications of artificial intelligence in retailing: A review. *Journal of Retailing and Consumer Services*, 52, 101926.

Mittal, R., & Mathew, R. (2019). Ethics of artificial intelligence: A review of the social and cultural implications. In *Intelligent Systems Design and Applications* (pp. 97-103). Springer, Singapore.