Artificial Intelligence in Bureaucracy and Government

Madi Turgunov
School of Engineering and Digital Sciences
Nazarbayev University
Astana, Kazakhstan
madi.turgunov@nu.edu.kz

***Abstract*—This survey paper explores the potential applications of Artificial Intelligence in government and bureaucracy. The paper examines various papers looking over the current, past and future developments in the usage of such models in the government apparatus. It also considers the main issues with said models, including ethical concerns, in terms of accountability, cost effectiveness and bias, and describes the research that is attempting to solve these issues.**

***Index Terms*—AI, Government, Automation, Neural Networks, Decision-making**

I. Introduction

A. Background Information

Artificial intelligence (AI) is, undoubtedly, a sector of computer science that, in the last several years, has drawn the attention of the entire world to itself. With the rapid progress done in the field of Natural Language Processing (NLP), the expansion of AI infrastructure, and the increasing efficiency of transformer models in general [1], the world has seemingly split into two categories of people: those who believe Artificial Intelligence is just a temporary fad, and those who think that AI will change the future as we know it – and the former camp is getting smaller and smaller with every new headline that appears on the news about the next significant achievement by researchers in this field.

There are, of course, doubtful opinions on both sides of the argument, which is expected to occur when any concept goes mainstream – doctored news, fearmongering, and utterly false information about AI's sentience, supremacy, or thirst for blood. In times like these, staying on track with reality and the capabilities of such powerful tools is essential, as it is still necessary to utilize them to their full potential.

While one must admit that the general population's consensus on AI seems to either be far-fetched or completely underestimated, one fact is clear – AI is not going anywhere, and those who find ways to implement it efficiently and sustainably will reap the benefits in the long game. Private companies and corporations have already jumped on the AI bandwagon, with predictive models dominating the advertisement space, detecting non-human internet clients, and executing other tasks that normal algorithms are unable to complete with reliable accuracy [2].

Nowadays, AI is already used in some aspects of government, most famously in street-level CCTV systems [3]. However, with the current generation of Artificial Intelligence models being exceptionally good at training on a large corpus of data [4], the idea of implementing it into one of the most ancient, convoluted, and confusing aspects of human society – bureaucracy – has, unsurprisingly, crossed a lot of people's minds. Bureaucracy is something that checks most of the marks for what AI is good at – it is well-defined, its processes done thousands of times a day, it has thousands of years of history and data to feed off of behind it, and most importantly – it is highly prone to human error.

This paper will study and evaluate the various ways in which AI has been used in government before the AI "Boom," how it is used now, and what the potential uses for it are in the future. Being a survey paper, it will also go over several publications and ideas about the potential future usage of Artificial Intelligence in many branches of the government, and give an abstract and filtered view to the reader about what the future of AI in government holds.

II. AI in the Government Sector

A. Application of Artificial Intelligence in Bureaucracy and Government

There are a multitude of different ways to utilize the power of Artificial Intelligence in the government sector. The main reason for that is that "Artificial Intelligence" is an extensive umbrella term covering many different subcategories, such as Natural Language Processing, Image Recognition, Generative AI, and others [4]. While simply piling all of these fundamentally different fields of advancement into one category might not show them justice, it is fitting for the scope of this survey paper.

The most basic and, seemingly, obvious usage of Artificial Intelligence in government, is already well-implemented on many non-governmental services. The usage of Generative AI for "Chatbots" has spread rapidly over many companies that find it cheaper to use it rather than pay a much larger amount of money to support personnel. Governments, however, have not yet fully embraced this, as many countries in the world are yet to digitize their services. Canada is one of the leading examples of the usage of Generative AI for its government services. All the way back in 2019, the government of Canada has made plans to utilize the potential of these chat models to improve the efficiency of their digital public services. [5]

Another strength of Artificial Intelligence is its unique ability to think probabilistically. Based on certain factors, AI can calculate the probability of something happening, given a large enough dataset for training. Such AI models are already in use by the Canadian immigration committee, which uses it to calculate the probability of fraudulent applications or potentially dangerous individuals. A similar system is in use in Poland. However, it is used for analyzing government job applications and streamlining employment. [5]

A much more advanced method of using the probabilistic features of Artificial Intelligence is to use it to predict and prevent unwanted events. For example, the Los Angeles Police Department is already utilizing such systems on a large scale in order to predict crime spikes based on environmental, economic, and political situations. Such systems can help the government contain or predict criminal or terrorist activities much more quickly. [5]

Conversely, Artificial Intelligence systems are also widely used for government surveillance. According to Carnegie Endowment for International Peace, at least 75 of 176 recognized countries in the world are already utilizing AI systems, particularly those with image recognition capabilities, for security purposes. These systems range from basic traffic control systems, like Kazakhstan's "Sergek," to social media crawlers, like those used by Huawei in China. [3]

B. Defining AI's Role in Government: Decision-Making vs Advisory

Whatever the specific usage of Artificial Intelligence in the government could be, it can be split into two main categories: advisory and autonomous. That is, the question is not whether Artificial Intelligence will be used in the future by the government – that has already become inevitable – but whether these algorithms will be making autonomous decisions, or act as advisory bodies to human government workers.

In reality, there can be no singular answer to that question. It is obvious that relying entirely on Artificial Intelligence for decision-making is bound to create countless ethical dilemmas ranging from the question of fault when it comes to erroneous decisions to some people questioning the sovereignty of humanity as a whole because of this. Several papers have been published regarding which aspects of public work are acceptable to be done autonomously, and which aspects are to be managed with minimal influence from AI. Caseworking, for example, is a type of work which can, technically, be fully automated, but the decisions that are made in this field of work can change a person's entire life and future. It is, therefore, a poor idea to fully automate this process, and it is better to instead utilize the processing capacity of Artificial Intelligence to assist existing caseworkers, providing them with more background information [6].

However, in many cases, algorithms can and have helped a lot more when they are left alone to do their job. Even today, massive amounts of data are processed automatically using basic pre-set algorithms. With the introduction of generative and analytical AI into the scene, such cases could be further automated, albeit with necessary transparency and human moderation. In system-level bureaucracies, that is, bureaucracies where most rules are already laid out, and ones that have massive amounts of data ready for analysis, the usage of AI can be done with minimal human intervention. One example is detection of fraudulent requests, saving time for people to do their actual jobs without having to deal with too much useless information. [7]

C. AI and Big Data: Optimization and Lawbreaking Detection

As mentioned previously, the main advantage of Artificial Intelligence over standard algorithms is its ability to analyze large amounts of data and extrapolate information from it. Notably, deep learning networks have the capacity to analyze huge amounts of data that are impractical for human minds to be able to comprehend. On top of that, the usage of feature extraction with representation learning, that is, the ability to identify important data using hidden patterns, and models' ability to continuously improve make them well-suited for efficiently processing and identifying patterns within massive historical datasets.

One such implementation, which would bring huge economic benefits, would be optimization. The inherent abilities of deep learning models to analyze logs of companies, factories or even entire cities' energy consumption would enable them to suggest optimizations into the schedule of most of the work done by humans nowadays. This sort of technology is already being implemented today in several cities across the world. They are showing results: a small city of Wellesley, Minnesota, in the United States of America has already implemented such a model, and its analyses have saved the small town over $132,000 within three years, despite having worse environmental conditions during the testing phase. [8]

A less cheerful, but just as important use of AI is optimizing law enforcement. The city of Chicago has been the center of attention when it comes to its usage of AI in police duties for a variety of reasons. Its use of AI's pattern recognition capabilities has optimzied restaurant inspections, improving the rate of detection of violators by over 20% [8]. A similar use of AI was experimented on in Belfast, United Kingdom. There, an improved deep learning model was able to identify 200% more businesses committing tax fraud as compared to their human counterparts. [9]

Apart from that, the Chicago Police Department (referred to as CPD from now on), through the usage of its own locally-trained SSL (Strategic Subject List) AI model, can identify, trace, and prevent potential security threats. Their model analyzes data from various sources, including historical crime records, arrest data, and social network analysis, to identify individuals at a high risk of being involved in violent crime. This information is then used to guide police interventions, such as increased patrols in high-risk areas and proactive engagement with individuals identified as potential victims or perpetrators. The citizen is scored from a risk value of 0 (no risk) to 500 (very high risk). [10]

All in all, deep learning has the potential to not only optimize human work but even completely change the way we operate due to its ability to efficiently analyze data. So far, many experiments with the usage of such models have had successful runs, and work into implementing them is already underway.

III. Concerns

Even considering all of the benefits that Artificial Intelligence will inevitably bring to the government and public sector, there are, undoubtedly, many concerns regarding its efficiency and usage. The majority of the issues that AI might bring coincide with the benefits it can provide, and it is essential to consider, analyze, and solve these issues before scaled adoption can begin.

A. Efficiency

The first issue with adopting such systems en masse arises from its intended purpose – computing a large amount of data for each sector where such a model would be usable requires a lot of computational power. Even normal servers that run AI tasks are more often than most incomparable to the computers that normal governmental institutions utilize [11]. With the current monetary costs of the entire AI infrastructure being around $76 billion, its mass implementation across the world would increase that number exponentially, which is a problem in itself [12]. Such spending might force the hand of the free market to lay off staff that is no longer needed due to AI optimizations and offset some of the savings obtained from those optimizations entirely.

B. Transparency and Bias

The second issue is the complexity of deep learning networks – more specifically, the fact that deep learning models can achieve such great results results in their decision-making processes not being transparent. Such lack of transparency will make it difficult for maintainers to see if the AI model has developed a bias, whether it be racial, gender, or otherwise. If the developers of said models were to be able to modify the contents of the training dataset, tune it, or affect it in any way, it would give rise to a noticeably unfair and opaque process, which would not be verifiable. For example, if existing police arrest records specify that men are more likely to commit crimes than women, then women who do commit crimes will be less likely to be caught by the model, thus reinforcing said bias in the eyes of the model itself [7].

C. Accountability

The third, and perhaps biggest issue about the usage of Artificial Intelligence, particularly in a decision-making scenario, is the question of accountability. In the cases of decisions made by AI going wrong, who is to take the blame? The people designing the AI? The person who used it? It is yet unknown, as no legal framework exists in such cases. The issues get even more complicated as humans are entirely removed from the equation: who would be accountable for an automatic fine for breaking the traffic rules assigned to the wrong person [13]? In the cases where humans are working alongside and assisted by AI, there were situations when the algorithm gave out an answer that was different from what the person would have judged. Would we discard the human decision and treat AI as an apolitical, unbiased, and entirely neutral being? Certainly not, since there is a certain possibility of there being a bias inside the layers of the model. However, most human beings have such biases as well, making it difficult on how exactly we treat Artificial Intelligence.

These situations are a clear indicator that AI is still not close to being the ultimate answer to everything. Almost every single benefit that AI brings to the table comes with its own set of disadvantages that we, as human beings, have to carefully consider before embracing it. [8]

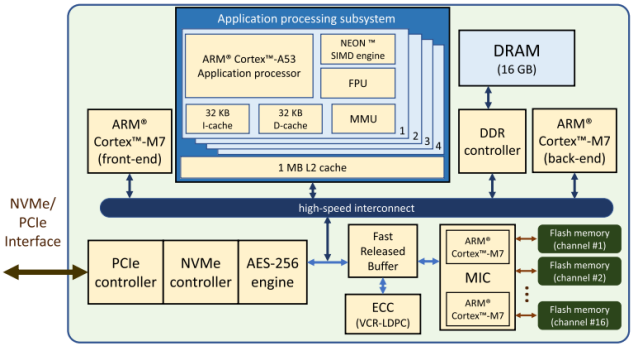

Fig. 1. Hardware Architecture of a Newport CSD

IV. Solutions

There has been a significant amount of research made in order to attempt to mitigate the issues described above. Numerous research articles were written in an attempt to help humanity get rid of them, which will be shortly discussed.

A. Efficiency

The issue with the high cost of Artificial Intelligence is bound to be mitigated as time goes by and as silicon becomes more efficient and less expensive. However, research into this area has proposed the usage of a new type of device: rather than relying on specialized processing units, the computational cost could be offloaded to the solid-state drives that store the data to be analyzed. As described by Jaeyoung Do and peers, who have worked on what they call a ”Newport CSD (Computational Storage Device).” The main benefit of such a device is that it enables in-storage processing. The architecture of the device can be seen in Figure 1, which showcases how exactly their proposed system is different from traditional storage solutions. [14]

These CSDs have shown to be efficient in minimizing power draw, as well as data efficiency by offloading some of the tasks usually done by the main AI processing chip to the drive itself. The drive is capable of caryring out **Similarity Search** using K-Nearest-Neighbor algorithms, as well as **Object Tracking**

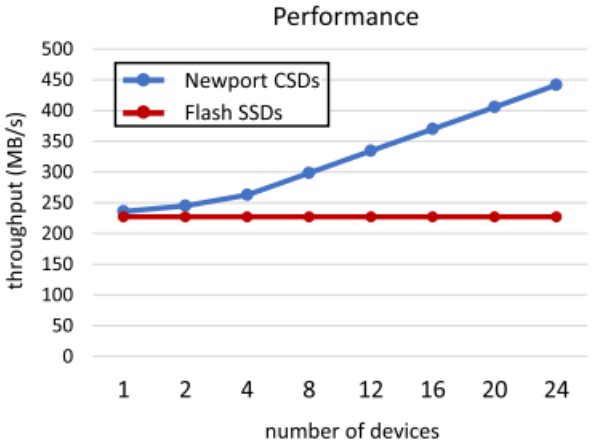

Fig. 2. Compression Throughput with multiple storage devices

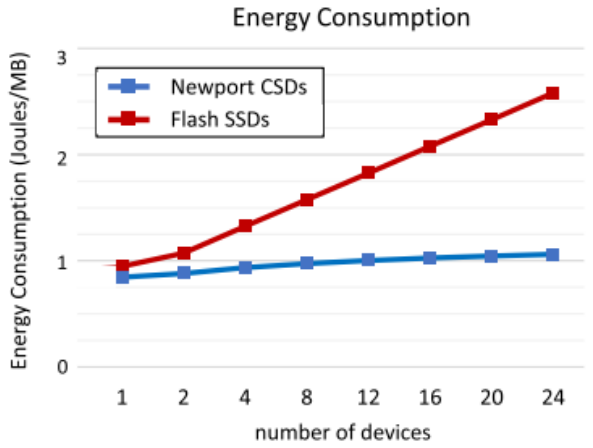

Fig. 3. Power draw from multiple storage devices

using famous algorithms such as YOLO, GOTURN, KCF, MOSSE and WiSARD.

The implementation of these CSD in building scalable AI datacenters could improve scalability, cost efficiency and power efficiency all at the same time. The price calculation done in the research suggested that the price of such a drive would not be marginally different from a standard storage drive but that they would dramatically improve performance and power draw as shown in Figures 2 and 3.

B. Bias

The removal of bias from datasets is a difficult task, which requires careful thought so as to preserve both the truthfulness of the data and not to cause the model that is to be trained by this dataset to develop discriminatory attitudes. One solution to this issue is synthetic datasets - that is, take existing datasets, such as video recordings of crimes, or crime logs and remove features that could lead to specific discrimination, such as gender, nationality or skin color. In one research paper, Ioannis Pastaltzidis et al. suggest an algorithm that would balance

Training Data	Mean Accuracy	std.
Original Videos	0.620	0.052
Synthetic Videos	0.642	0.028

TABLE I
RTFM ON RWF-2000 RACE BIAS 150 EPOCHS TRAINING

training datasets by recognizing and obfuscating the data [15]. Whilst this may sound controversial at first, since some might consider that these factors are critical to recognition of crime statistics, the analysis of the paper says otherwise: the accuracy of models trained on the obfuscated datasets turn out to be better than those with no such obfuscation, when trained on the RWF-2000 dataset (Table I). And while the paper does only consider video-based crime detection, this same technique could be applied to more models in order to make them unbiased, and to ensure that the decisions made by it are fair.

C. Accountability

The issue of accountability is still under heavy debate, however some frameworks have already been developed to attempt to solve the issue. One such framework is the SMACTR framework, proposed by researchers at Google, which streamlines such accountability questions [16]. The framework isn't build to attach blame directly, but to be proactively mitigating risks. SMACTR itself consists of four stages: *Scoping*, *Mapping*, *Artifact Collection*, *Testing* and *Reflection*.

1) Scoping: Scoping is the initial phase of developing any AI system. It includes all the basics: project proposals, ethical reviews, use case reviews and calculating the social impact. This is the most important stage, as it helps reviewers identify their next goals.

2) Mapping: This phase emphasizes mapping the existing findings from the scoping phase to the affected demographics or key stakeholders. In this stage, priorities are set on the goals that are to be achieved with the AI model.

3) Artifact Collection: Artifacts are important documentation and insights. In this phase, datasheets for datasets, design checklists and model cards - high level descriptions of the intended use fields of the models - are collected.

4) Testing: In the testing phase, the AI model is actively probed in order to test its alignment on ethical issues. This includes adversarial testing, which aims to simulate a case of misuse of the AI model. In the end, an "Ethical Risk Analysis Chart" is produced, highlighting any risks that might arise from the usage of this model.

5) Reflection: This is the final and the most obvious stage. In the reflection stage, the development team as well as the stakeholders look over all the findings from all of the phases, and decide on whether the model is ready for deployment, as well as identify areas for improvement.

How exactly does it solve the issue of accountability? It creates a more transparent and traceable environment for the development of the AI system. Not only does the rigorous testing prevent a lot of the potential issues with the system, but anything that does go wrong can be traced back to one of

these sections. According to this framework, the developers of the model are responsible for their model's misbehaviour.

V. Conclusion

In conclusion, the usage of Artificial Intelligence in bureaucracy has been a topic of discussion for a long time, and has not reached its full potential. This, however, does not mean that ideas from the research into the sphere isn't in use today: on the contrary, more and more experiments show that it is inevitable that various AI algorithms will make their way into the government very soon. The question on the table is not whether we begin using Artificial Intelligence to optimize our bureaucracy, it's how we use it and whether or not we are able to solve all of it's issues before we start doing so.

From all of the publications mentioned in this survey paper, the majority of them have mentioned how these systems are still in their nascent stage: they still make mistakes, they are often inefficient and prone to pre-existing biases. It is also clear that our legal system is not yet ready for the full implementation of autonomous workers, as it lacks the means to judge their behavior. However, it is also clear that these systems show the potential to greatly improve the overall efficiency of our government apparatus, and automating even the least time-consuming and simplest processes can yield great results when looking at the big picture. AI has the potential to streamline repetitive work, enhance decision-making awareness and make government decisions fairer and less corrupt.

This paper also showed that despite the lingering problems of Artificial Intelligence, research is being done on mitigating them and making AI a more safe, accessible and efficient solution to many of the world's problems. Further research is, of course, required for AI to become truly applicable to the masses in the form of its use by the government.

VI. Ethical Considerations

The ethical considerations of the usage of such Big Data and AI projects are very broad.

1) Data Collection: To efficiently utilize Artificial Intelligence in areas such as bureaucracy, a large amount of data needs to be collected and analyzed. This data has to be analyzed ethically and consensually.

2) Historical Bias: If AI is used in decision-making processes, the data that it is trained on must be made out of just and correct decisions. A country with a historical past of oppression and injustice might have to fabricate a lot of training data to be just.

3) Personal Freedoms: Artificial Intelligence can give a lot of potential power to the government apparatus, and its usage must be overseen by independent organizations to minimize its potential abuse.

References

[1] D. Khurana, A. Koli, K. Khatter, and S. Singh, "Natural language processing: state of the art, current trends and challenges," *Multimedia Tools and Applications*, vol. 82, p. 3713–3744, July 2022.

[2] B. Gao, Y. Wang, H. Xie, Y. Hu, and Y. Hu, "Artificial intelligence in advertising: Advancements, challenges, and ethical considerations in targeting, personalization, content creation, and ad optimization," *SAGE Open*, vol. 13, no. 4, p. 21582440231210759, 2023.

[3] S. Feldstein, "The global expansion of ai surveillance," Sep 2019.

[4] Y. Xu, X. Liu, X. Cao, C. Huang, E. Liu, S. Qian, X. Liu, Y. Wu, F. Dong, C.-W. Qiu, J. Qiu, K. Hua, W. Su, J. Wu, H. Xu, Y. Han, C. Fu, Z. Yin, M. Liu, R. Roepman, S. Dietmann, M. Virta, F. Kengara, Z. Zhang, L. Zhang, T. Zhao, J. Dai, J. Yang, L. Lan, M. Luo, Z. Liu, T. An, B. Zhang, X. He, S. Cong, X. Liu, W. Zhang, J. P. Lewis, J. M. Tiedje, Q. Wang, Z. An, F. Wang, L. Zhang, T. Huang, C. Lu, Z. Cai, F. Wang, and J. Zhang, "Artificial intelligence: A powerful paradigm for scientific research," *The Innovation*, vol. 2, no. 4, p. 100179, 2021.

[5] M. Kuziemski and G. Misuraca, "Ai governance in the public sector: Three tales from the frontiers of automated decision-making in democratic settings," *Telecommunications Policy*, vol. 44, no. 6, p. 101976, 2020.

[6] A. Ammitzbøll Flügge, T. Hildebrandt, and N. H. Møller, "Street-level algorithms and ai in bureaucratic decision-making," *Proceedings of the ACM on Human-Computer Interaction*, vol. 5, no. CSCW1, pp. 1–23, 2021.

[7] J. Bullock, "Artificial intelligence, bureaucratic form, and discretion in public service," *Information Polity*, pp. 1–15, 2020.

[8] M. J. Ahn and Y.-C. Chen, "Artificial intelligence in government: Potentials, challenges, and the future," in *The 21st Annual International Conference on Digital Government Research*, dg.o '20, (New York, NY, USA), p. 243–252, Association for Computing Machinery, 2020.

[9] T. Vogl, C. Seidelin, B. Ganesh, and J. Bright, "Algorithmic bureaucracy," in *Proceedings of the 20th Annual International Conference on Digital Government Research*, dg.o 2019, (New York, NY, USA), p. 148–153, Association for Computing Machinery, 2019.

[10] G. N. Kouziokas, "The application of artificial intelligence in public administration for forecasting high crime risk transportation areas in urban environment," *Transportation Research Procedia*, vol. 24, pp. 467–473, 2017. 3rd Conference on Sustainable Urban Mobility, 3rd CSUM 2016, 26 – 27 May 2016, Volos, Greece.

[11] J. Welser, J. W. Pitera, and C. Goldberg, "Future computing hardware for ai," in *2018 IEEE International Electron Devices Meeting (IEDM)*, pp. 1.3.1–1.3.6, 2018.

[12] J. McGregor, "Generative ai breaks the data center: Data center infrastructure and operating costs projected to increase to over $76 billion by 2028," *Forbes Innovation*, vol. ..., May 2023.

[13] F. Gualdi and A. Cordella, "Artificial intelligence and decision-making: The question of accountability," in *Proceedings of the Annual Hawaii International Conference on System Sciences*, 2021.

[14] J. Do, V. C. Ferreira, H. Bobarshad, M. Torabzadehkashi, S. Rezaei, A. Heydarigorji, D. Souza, B. F. Goldstein, L. Santiago, M. S. Kim, P. M. V. Lima, F. M. G. França, and V. Alves, "Cost-effective, energy-efficient, and scalable storage computing for large-scale ai applications," *ACM Trans. Storage*, vol. 16, oct 2020.

[15] I. Pastaltzidis, N. Dimitriou, K. Quezada-Tavarez, S. Aidinlis, T. Marquenie, A. Gurzawska, and D. Tzovaras, "Data augmentation for fairness-aware machine learning: Preventing algorithmic bias in law enforcement systems," in *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '22, (New York, NY, USA), p. 2302–2314, Association for Computing Machinery, 2022.

[16] I. D. Raji, A. Smart, R. N. White, M. Mitchell, T. Gebru, B. Hutchinson, J. Smith-Loud, D. Theron, and P. Barnes, "Closing the ai accountability gap: defining an end-to-end framework for internal algorithmic auditing," in *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, FAT* '20, (New York, NY, USA), p. 33–44, Association for Computing Machinery, 2020.