



Delft University of Technology

## Exploring the Viability of ChatGPT for Personal Data Anonymization in Government A Comprehensive Analysis of Possibilities, Risks, and Ethical Implications

van Staalduine, C.M.C.; Zuiderwijk, Anneke

DOI

[10.1145/3678264](https://doi.org/10.1145/3678264)

Publication date

2025

Document Version

Final published version

Published in

Digital Government: Research and Practice

### Citation (APA)

van Staalduine, C. M. C., & Zuiderwijk, A. (2025). Exploring the Viability of ChatGPT for Personal Data Anonymization in Government: A Comprehensive Analysis of Possibilities, Risks, and Ethical Implications. *Digital Government: Research and Practice*, 6(2), Article 24. <https://doi.org/10.1145/3678264>

### Important note

To cite this publication, please use the final published version (if applicable).  
Please check the document version above.

### Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

### Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.  
We will remove access to the work immediately and investigate your claim.



# Exploring the Viability of ChatGPT for Personal Data Anonymization in Government: A Comprehensive Analysis of Possibilities, Risks, and Ethical Implications

**NINA VAN STAALDUINE**, Faculty of Technology Policy and Management, Delft University of Technology, Delft, Netherlands

**ANNEKE ZUIDERWIJK**, Faculty of Technology Policy and Management, Delft University of Technology, Delft, Netherlands

Research on the potential use of ChatGPT for anonymizing texts in government organizations is scarce. This study examines the possibilities, risks, and ethical implications for government organizations to use ChatGPT in the anonymization of personal data in text documents. It adopts a case-study research approach, including informal conversations, formal interviews, literature review, document analysis, and experiments. The experiments using three types of texts demonstrate ChatGPT's proficiency in anonymizing diverse textual content. Furthermore, the study provides an overview of significant risks and ethical considerations pertinent to ChatGPT's use for text anonymization within government organizations, related to themes such as privacy, responsibility, transparency, bias, human intervention, and sustainability. The current form of ChatGPT stores and forwards inputs to OpenAI and potentially other parties, posing an unacceptable risk when anonymizing texts containing personal data. We discuss several potential solutions to address these risks and ethical issues. This study contributes to the scarce scientific literature on the potential value of employing ChatGPT for text anonymization in government settings. It also offers practical insights for civil servants coping with the challenges of personal data anonymization, emphasizing the need for the cautious consideration of risks and ethical implications in the integration of AI technologies.

CCS Concepts: • **General and reference** → **Empirical studies**; • **Computing methodologies** → **Information extraction**; • **Security and privacy** → *Social aspects of security and privacy*;

Additional Key Words and Phrases: ChatGPT, personal data anonymization, AI in government, ethical implications

## ACM Reference Format:

Nina van Staalduine and Anneke Zuiderwijk. 2025. Exploring the Viability of ChatGPT for Personal Data Anonymization in Government: A Comprehensive Analysis of Possibilities, Risks, and Ethical Implications. *Digit. Gov. Res. Pract.* 6, 2, Article 24 (June 2025), 15 pages. <https://doi.org/10.1145/3678264>

Authors' Contact Information: Nina van Staalduine, Delft University of Technology, Delft, Netherlands; e-mail: [c.m.c.vanstaalduine@student.tudelft.nl](mailto:c.m.c.vanstaalduine@student.tudelft.nl); Anneke Zuiderwijk (Corresponding author), Delft University of Technology, Delft, Netherlands; e-mail: [a.m.g.zuiderwijk-vaneijk@tudelft.nl](mailto:a.m.g.zuiderwijk-vaneijk@tudelft.nl).



This work is licensed under a Creative Commons Attribution-NoDerivs International 4.0 License.

© 2025 Copyright held by the owner/author(s).

ACM 2639-0175/2025/06-ART24

<https://doi.org/10.1145/3678264>

## 1 Introduction

Government organizations worldwide are exploring the potential of Artificial Intelligence (AI) for the delivery of public services [Van Noordt and Misuraca 2022]. For example, AI is commonly utilized to improve delivery systems and facilitate the dissemination of information to citizens and businesses [Van Noordt and Misuraca 2022]. In the literature, AI is often referred to as a catalyst for improving efficiency and effectiveness within the public sector [Toll et al. 2020]. Moreover, AI use in public sector organizations may improve service quality, reduce lead times, support governments to make unbiased decisions in case handling [Lindgren et al. 2019], and enhance citizens' trust in government [Dwivedi et al. 2019].

One AI tool that received considerable attention is ChatGPT, a language model proficient in text processing and generating content. Due to its functionalities, ChatGPT possesses the capability to comprehend the context of a dialogue and produce fitting replies [Deng and Lin 2022]. Additionally, it can generate responses in various languages [Deng and Lin 2022; Hariri 2023]. ChatGPT has the potential to enhance efficiency by automating conversations, producing responses more accurately than typical interactions [Deng and Lin 2022], and notably lowering costs for businesses and governments that depend on customer service chatbots [Deng and Lin 2022; Hariri 2023; Paul Ueno and Dennis 2023]. It has also been suggested that large language models like ChatGPT can be used for text anonymization, where anonymization is defined as safeguarding an individual's privacy by eliminating or disguising personally identifiable information [Datta et al. 2023].

In sum, several studies already examined the use of AI in government [Ahn and Chen 2022; Valle-Cruz and Sandoval-Almazan 2018; Van Noordt and Misuraca 2022]. However, given the recent launch of ChatGPT, there has been limited prior investigation into whether ChatGPT can be employed for the anonymization of texts within government organizations and what the implications of this employment would be. As the application of AI in government may be both time- and cost-efficient, this lack of research is a missed opportunity. The main research objective of this study is to examine the possibilities, risks, and ethical implications for government organizations to employ ChatGPT in the anonymization of personal data in text documents. This study contributes to the scarce scientific literature on the potential value of ChatGPT for personal data anonymization in government. In addition, it has practical value for civil servants who face the challenges of data anonymization in practice including resource-intensive and costly processes.

This article is structured as follows: The next section describes the research background and gives more information about the development of ChatGPT. Then, the research design is described, followed by the results of several data anonymization experiments. Finally, we provide an overview of significant risks and ethical issues related to ChatGPT and its use for anonymization within a specific government organization.

## 2 Research Background

This section describes the technology behind ChatGPT and, specifically, its operation (Section 2.1), followed by an overview of previous research on ChatGPT (Section 2.2).

### 2.1 The Operation of ChatGPT

In November 2022, OpenAI launched a new chatbot called ChatGPT. ChatGPT, short for Generative Pre-Trained Transformer, is a form of generative AI. Generative AI applications have the ability to autonomously create human-like content such as text, audio, code, music, and images [Dasborough 2023; Fui-Hoon Nah et al. 2023]. Moreover, such AI applications can consolidate data from different sources for analysis [Dasborough 2023; Fui-Hoon Nah et al. 2023]. ChatGPT, specifically, can generate both text and code. It functions as a language model, more specifically, a Large Language Model (LLM), developed based on a neural network known as "Transformer," which stands as one of the latest advancements in this field [Shen et al. 2023]. It was trained on a massive dataset of around 45 terabytes of text data [Cooper 2021]. This extensive training allows ChatGPT to generate content that closely resembles human-written text. ChatGPT has the potential to anonymize data based on

the Named Entity Recognition (NER) principle. NER aims to identify and classify important nouns and proper names in a given text [Mohit 2014]. By labelling words in sentences, names can be recognized, which facilitates anonymization.

Computers, such as ChatGPT, comprehend language through the technique of word embedding, where words are converted into numerical values (vectors in a high-dimensional space), while preserving the semantics and syntax of the words [Wang et al. 2019]. Tokens are frequently occurring sequences of characters in text [OpenAI 2024]. They can consist of meaningful and recognizable words or series of letters that often appear in words but lack inherent meaning. ChatGPT uses the word embedding of tokens to assist computers in understanding the meaning of words, including compound words.

The training of a language model roughly involves two steps: unsupervised pre-training and supervised fine-tuning [Radford et al. 2018]. During pre-training, the model learns to predict the next token given a sequence of tokens by comparing the output with the correct token. In this way, the model autonomously learns to recognize patterns in a large amount of text. Supervised fine-tuning is the second part of the training process, involving human guidance to teach the language model to follow instructions correctly [Ouyang et al. 2022]. People provide the desired output for different inputs, and the model learns from these examples. Subsequently, a reward model is trained as humans rank the different outputs associated with the same input. Through the reward model and reinforcement learning, the model undergoes further fine-tuning. In fine-tuning, efforts are made to ensure that the model responds “helpfully,” “honestly,” and “harmlessly” [Askell et al. 2021] to align with the needs and intentions of users [Lowe and Leike 2022].

## 2.2 Previous Research on ChatGPT

Extant research discusses the significant societal benefits generated by the capability of ChatGPT to generate human-like text. These benefits range from fulfilling small writing tasks to assisting in composing extensive scientific pieces [Salvagno et al. 2023; Van Dis et al. 2023]. Embracing the advantages of AI could lead to breakthroughs across various fields, driving accelerated innovation [Van Dis et al. 2023]. Scholars state that the technology behind ChatGPT has the capacity to tackle cybersecurity issues, safeguard against threats and attacks, and tackle the challenges linked to our growing dependence on technology and the internet [Alsumayt et al. 2024].

Some studies investigated the use of ChatGPT in government. ChatGPT could prove beneficial for public agencies, as ChatGPT has the potential to simplify the decision-making process of the public sector [Cao et al. 2024]. Moreover, ChatGPT enables a broader range of stakeholders to contribute to policymaking, policy analysis, and strategic planning in governmental organizations [Cao et al. 2024; J. Huang and Huang 2023] and better involve citizens [J. Huang and Huang 2023]. Various applications of ChatGPT to government services have been identified, such as open government data portals [Mamalis et al. 2023], citizen services, tax filing, and voting processes [J. Huang and Huang 2023]. A study by Yang and Wang [2023] focused on the public acceptance of integrating ChatGPT into government services. They found that the public acceptance in this context hinges primarily on perceived risk, trust, and meeting demand, identifying these three factors as the most crucial. In addition, anonymizing data has been mentioned as a potential use case for ChatGPT in court documents within the Italian Public Administration, where AI-based anonymization can effectively obscure personally identifiable information [Datta et al. 2023].

Although AI applications may be valuable in a government context, AI-based data anonymization should not substitute standard data privacy best practices. The degree of anonymization and deidentification of such tools can vary, and even the most well-trained AI model cannot ensure that a document cannot be traced back to its original source. Patsakis and Lykousas [2023] recently carried out an experiment where ChatGPT was used to examine if anonymized texts about celebrities can be deanonymized using ChatGPT and found that GPT demonstrated remarkable results in the deanonymization of text on famous people and outperformed humans almost three times.

Hence, in addition to the potential benefits, previous research examined the prospects and challenges of using LLMs such as ChatGPT, emphasizing the treats of privacy and data security issues in the use of ChatGPT [Alsumayt et al. 2024; K. Huang et al. 2023; Nazir and Wang 2023]. For example, it is stated that deploying ChatGPT as a tool requires careful consideration, given that it is a novel technology, and numerous questions about its functionality remain unanswered. González-Gallardo et al. [2023] investigated NER with ChatGPT, revealing limitations in its effectiveness. Other studies discuss the ethical implications of ChatGPT, including privacy and responsibility issues [Mattas 2023]. Particularly in an organization dealing with confidential information and personal data, extra attention must be given to these considerations.

From our literature review, we conclude that previous research addressed the potential and risks of using ChatGPT in government. Nevertheless, our review barely identified studies exploring whether ChatGPT or similar language models can be employed for personal data anonymization within a government setting. The outcomes of such research could have a tangible impact on addressing labor shortages and reducing the high workload for government officials.

### 3 Research Design

This study examines the possibilities, risks, and ethical implications for government organizations to employ ChatGPT in the anonymization of personal data in text documents. This section describes the conceptual design (Section 3.1) and operational design (Section 3.2) of the study.

#### 3.1 Conceptual Design

Inspired by general notions from various technology acceptance and use theories (e.g., Venkatesh et al. [2003]; Venkatesh et al. [2011]; Venkatesh et al. [2012]), this study assumes three factors influence the intention of civil servants to use ChatGPT for personal data anonymization:

- (1) The technical operation of ChatGPT;
- (2) The way ChatGPT performs in the anonymization of personal data in text documents and the functionalities that support this;
- (3) The risks and ethical aspects of employing ChatGPT for personal data anonymization in government.

While it is out of this study's scope to measure these factors in a quantitative manner, we assume that the three above-mentioned factors each have a direct effect on the intention to use ChatGPT for personal data anonymization in government and that the intention to use ChatGPT has a direct effect on its actual use (see Figure 1). The first factor, i.e., the technical operation of ChatGPT, can help ChatGPT users to understand the workings of the technology (as many users are assumed to have limited knowledge about AI), allowing them to form their own

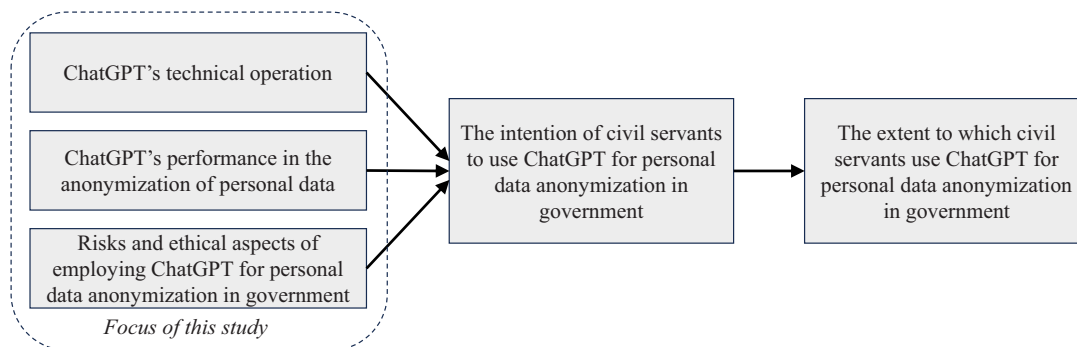


Fig. 1. Conceptual model and focus of this study.

well-founded opinion on the technology. The formed opinions will influence the intention of civil servants to use the technology. The second factor focuses on the performance of ChatGPT, because high performance in data anonymization will increase the intention to use the technology. We assume useful functionalities of a new technology will increase the intention to use this technology, as this could increase productivity and effectiveness. The last factor, concerning risks and ethical aspects of the new technology, is important to mitigate risks and safeguard public values. Dangers may arise from the potential use of ChatGPT for text anonymization. By highlighting the risks and ethical aspects, users can form an opinion on the desirability of using this technology, which we assume affects the intention to use it.

### 3.2 Operationalization

Considering the novelty of ChatGPT, this study incorporated a case-study research approach, employing a blend of complementary methodologies. The embedded case study focused on the use of ChatGPT for data anonymization at the Dutch Ministry of Justice and Security, involving two organizations within that ministry that collaborate on exploring ChatGPT's viability for personal data anonymization. The two organizations concern the Judicial Information Service ("Justitiele Informatiedienst" in Dutch) and Netherlands Forensic Institute ("Nederlands Forensisch Instituut" in Dutch). JustID aims to store and safeguard the Ministry's information, and it facilitates the sharing of information among various government agencies [Ministerie van Justitie en Veiligheid 2018]. Netherlands Forensic Institute is a forensic institute that provides products and services to a wide range of national and international clients (Netherlands Forensic Institute. Ministry of Justice and Security).

At the case-study organizations, a significant amount of work involves handling documents containing personal data, often of a confidential nature, requiring utmost care in processing. In certain cases, these documents need to be anonymized. Currently, anonymization is performed manually and through third-party AI systems from three different vendors (IBM, Octobox, Indica), both of which incur substantial costs. Due to its functionalities, ChatGPT has the potential to be valuable in the more cost-effective anonymization of texts within the government. For these reasons, the Judicial Information Service and Netherlands Forensic Institute are interesting organizations to investigate for our research purposes. The following subsections explain the two-step approach that we used for our research, following recommendations by, for example, Yin [2018] concerning the combination of multiple information sources.

*3.2.1 Step 1: Performance of ChatGPT in Personal Data Anonymization.* First, we examined how ChatGPT performs in anonymizing texts by conducting three types of experiments. These experiments involved inputting prompts into ChatGPT with instructions to anonymize a given text, which is explained in detail below. All the underlying prompts, code, texts for the experiments, and specific results referred to in this section are available online (DOI: [10.4121/a1dfacbe-b463-404f-a3d7-dab8485e6458](https://doi.org/10.4121/a1dfacbe-b463-404f-a3d7-dab8485e6458)). To prevent privacy violations, no documents from the examined case were used, as they typically contain confidential information. Three types of texts with fictional personal data were used for this experiment:

- (1) Generated Dutch fake data. First, realistic personal data was created in Python using the Faker library [Fraglia 2014]. This library generates random but realistic personal data in the selected language. For this research, 10 fictional names, addresses (street, house number, postal code, and city), phone numbers, and IBAN numbers were generated. The fictional personal data was then given to ChatGPT along with the instruction to create a story.
- (2) Literature without copyright. Old literature is not copyrighted and hence we chose the Dutch translation of the book *Around the World in 80 Days*, by Jules Verne (1885), for this second type of experiment. A portion of chapter 1 of the book, containing 40 pieces of personal data, was used for this experiment.
- (3) Modified literature without copyright. Since even the translation still contained quite a bit of old language usage, it was decided to adapt this text to modern Dutch. This led to the third type of text that we used in our experiments: the modified literature without copyright. We replaced old words with modern variants.



Table 1. Confusion Matrix

		Reality	
		Positive	Negative
ChatGPT experiment	Positive	True positive (TP). Word has been anonymized correctly.	False positive (FP). Word has been anonymized, but should not have been anonymized
	Negative	False negative (FN). Word has not been anonymized, but should have been anonymized	True negative (TN)

The English names and addresses were also changed to Dutch names and addresses, since our case is based in the Netherlands.

First, a manual anonymization was conducted in which all words that ChatGPT should anonymize were highlighted in grey. Subsequently, an anonymization prompt was given to ChatGPT for each of the three different types of texts: “Anonymize the following text by replacing all addresses, place names, locations, first names, last names, and origins with [...]”. Since ChatGPT often provides slightly different answers to the same question, the prompt for each type of text was entered three times in ChatGPT.

For each of the experiments, the ChatGPT-generated output was then color-coded, highlighting words with green, red, and yellow. A green word represented correct anonymization (true positive). A red-colored word indicated a word that should have been anonymized but was not anonymized by ChatGPT (false negative). Yellow highlighting indicated that ChatGPT had anonymized too much; these words could have remained non-anonymized, but ChatGPT removed them (false positive). Words without highlighting were correctly not anonymized by ChatGPT (true negative). To examine how well ChatGPT performs in the anonymization task, we compared the ChatGPT-generated anonymization to our manual anonymization (see the underlying data, DOI: [10.4121/a1dfacbe-b463-404f-a3d7-dab8485e6458](https://doi.org/10.4121/a1dfacbe-b463-404f-a3d7-dab8485e6458)).

The anonymization of texts is essentially a classification problem, meaning that objects need to be categorized into specific classes. In this case, words need to be divided into two categories: personal data (to be anonymized) and non-personal data (to remain unchanged). We created a so-called “confusion matrix” to provide an overview of how many words ChatGPT has correctly and incorrectly anonymized, and thus categorized (see Table 1).

Based on the confusion matrices, four metrics were computed:

- (1) Recall: a crucial metric that indicates the proportion of words that were supposed to be anonymized and were indeed anonymized. The lower the number of words missed in anonymization, the better the model performs.

$$Recall = \frac{TP}{TP + FN}$$

- (2) Precision: a metric that indicates the proportion of words that were anonymized and should have been anonymized. The higher the precision value, the more readable the text is.

$$Precision = \frac{TP}{TP + FP}$$

- (3) Accuracy: a less crucial metric that reflects the proportion of the total number of words that have been correctly anonymized or not. Since the number of words that should not be anonymized is much larger for the texts than the number of words that should be anonymized, the strength of this metric diminishes somewhat, and the values for accuracy tend to be very high.

$$Accuracy = \frac{TP + TN}{TP + FP + FN + TN}$$

- (4) F-score: a metric that calculates the harmonic mean of precision and recall. Striking the right balance between precision and recall ensures a good tradeoff between accurate anonymization and readability, making the F-score an important metric.

$$Fscore = 2 \cdot \frac{Precision \cdot Recall}{Precision + Recall} = \frac{2TP}{2TP + FP + FN}$$

These terms each signify in their own way how well ChatGPT can anonymize and are metrics commonly employed for the evaluation of classification problems. The values of these metrics for each of the experiments were determined, analyzed, and compared with each other and with other studies. This approach allowed for the evaluation of ChatGPT's performance in anonymizing texts. Our description and analysis of the experiment results was reviewed by a team of four data scientists from the Judicial Information Service.

*3.2.2 Step 2: Risks and Ethical Issues of Using ChatGPT in Government.* Second, we investigated the risks and ethical questions associated with the use of ChatGPT for anonymizing personal data, following the Value Sensitive Design method. As described by Van den Hoven et al. [2015], Value Sensitive Design aims to align technology more with values. It involves placing societal and moral values [Friedman 1997] at the forefront of developing new technology. Using a top-down approach, values are initially identified, and technology is then developed based on those values. While ChatGPT is already an existing technology, this research focuses on the form of its potential implementation within the examined case. A literature review was conducted to identify common risks and ethical questions associated with the use of AI. Google Scholar was used as the search engine. Forward and backward searching were employed during this literature review to gather relevant articles. Subsequently, we derived seven key themes concerning AI, anonymization, and ethics from the selected literature. Then, an interview protocol was developed based on the prevalent risks and ethical themes (see online documentation). For every theme, one or multiple questions were developed. To give an example, the theme “responsibility” was found in literature [Cath et al. 2018; Fjeld et al. 2019; Murphy et al. 2021]. This theme was translated into the following questions in the interview protocol: “If something goes wrong with anonymizing using ChatGPT, who would be responsible? Does OpenAI have a role in that?” and “To what extent should humans still be involved in the use of ChatGPT for anonymization? For example, should there be human controls and validations?”

Subsequently, two semi-structured, online interviews were conducted. Before each interview, consent was obtained for audio recording and processing of the interview for this study. The first interview involved a scientist focusing primarily on the moral responsibility associated with the use of algorithms and AI in the public sector. The second interviewee was Chief Data Officer at a Dutch public agency, having a background in law combined with ICT. This allowed for a comprehensive examination of the risks and ethics from legal, ethical, and technical perspectives. Both interviews were transcribed and coded. Following Williams and Moser [2019], initially, an open coding approach was used, assigning fairly specific labels to sentences from the transcripts. For example, from the literature, we identified the theme “transparency” [Cath et al. 2018; Fjeld et al. 2019]. In the open coding stage for this theme, we identified 13 sub-themes, such as “black box,” “explainability,” “transparency obligations,” and “transparency context-dependent.” Subsequently, axial coding was performed, searching for overarching themes for the labels established during open coding. In our specific example for the transparency theme, the axial coding phase led to the overarching themes “transparency and explainability” and “risk and ethics (in general).”

## 4 Results

This section first describes the results of our data anonymization experiments (Section 4.1). Thereafter, it reports on our literature review and interviews concerning the risks and ethical issues for the use of ChatGPT for anonymization by government organizations (Section 4.2).



#### 4.1 The Performance of ChatGPT in Personal Data Anonymization

To assess how well ChatGPT can anonymize data, three distinct types of texts were anonymized three times each: fake data, original literature, and modified literature. Tables 2, 3, and 4 represent the computed recall, precision, accuracy, and F-scores for these three text types. In the tables, the lighter shades of grey represent the highest values, with the lightest shade of grey indicating precisely 1.0. The lower the values, the darker the shades of grey. Based on the values in the tables and in comparison with similar studies (e.g., Hassan & Domingo-Ferrer [2018]; Hassan et al. [2019]; Szarvas et al. [2007]), it is evident that ChatGPT generally performs well in anonymizing the three different types of texts. For all metrics, higher values indicate better performance. There are only two outcomes with values below 0.8, but these values still remain above 0.7. The total scores are all above 0.8, signifying that ChatGPT performs well on average.

In the case of fake data, as seen in Table 2, perfect anonymization is achieved in two instances. In tests 1 and 3, the values for recall, precision, accuracy, and F-score are all precisely 1.0, indicating complete and accurate anonymization. Test 2 is not entirely perfectly anonymized but still attains a commendable score. The total scores for recall and precision are 0.996 and 0.897, respectively, and the F-score is 0.944. ChatGPT demonstrates the ability to accurately anonymize personal data in generated fake data. The exceptional performance with fake data might be attributed to the fact that the fake data is also generated by ChatGPT, potentially resulting in a recognizable pattern in the text.

Table 2. Evaluation Metrics for the Anonymization Experiment with Fake Data

Fake data	Recall	Precision	Accuracy	F-score
Test 1	1.000	1.000	1.000	1.000
Test 2	0.9878	0.743	0.955	0.848
Test 3	1.000	1.000	1.000	1.000
Average	0.996	0.897	0.985	0.944

In the anonymization of the original literature, ChatGPT performs slightly less effectively than with fake data (see Table 3). The overall values for precision and F-score are 0.962 and 0.893, respectively, but the total value for recall is lower at 0.833. The complexity of archaic Dutch language could be a contributing factor. ChatGPT might struggle to “comprehend” archaic Dutch texts with words and sentence structures no longer used in modern language (and thus occurring infrequently or not at all in ChatGPT’s training data). Nevertheless, ChatGPT still performs reasonably well in anonymizing original literature, but there is room for improvement.

Table 3. Evaluation Metrics for the Anonymization Experiment with Original Literature

Fake data	Recall	Precision	Accuracy	F-score
Test 1	0.775	1.000	0.989	0.873
Test 2	0.825	0.892	0.986	0.851
Test 3	0.900	1.000	0.995	0.947
Average	0.833	0.962	0.990	0.893

The modified literature shows similar results to the original literature (see Table 4). The recall is slightly higher than that of the original literature, possibly due to ChatGPT’s improved understanding of the text resulting from the change from archaic to modern words. The F-scores of the modified literature, almost identical to those of the original literature, are high but not at the same level as with fake data.

Table 4. Evaluation Metrics for the Anonymization Experiment with Modified Literature

Fake data	Recall	Precision	Accuracy	F-score
Test 1	0.854	0.921	0.989	0.886
Test 2	0.900	0.857	0.988	0.878
Test 3	0.854	1.000	0.993	0.920
Average	0.869	0.922	0.990	0.895

To conclude, ChatGPT demonstrates a strong performance in anonymizing the three different types of texts. Compared to other studies, ChatGPT performs at an average to a good level. Fake data is almost perfectly anonymized, indicating ChatGPT’s capability for accurate anonymization. However, it is essential to note that this excellent performance may be influenced by the fact that the fake data was generated by ChatGPT itself. Performance with original literature is slightly less impressive compared to fake data, potentially due to the complexity of the archaic Dutch language present in this literature. Modified literature shows comparable results to original literature but performs slightly better in certain aspects, likely attributed to adjustments towards modern Dutch. Overall, ChatGPT has achieved a high level of anonymization in these experiments.

#### 4.2 Risks and Ethical Issues of Using ChatGPT in Government

The previous section reveals that ChatGPT delivers reasonable performance in anonymizing texts, but a subsequent question arises: Is it safe and ethically responsible to use ChatGPT for anonymization? This section provides an overview of the risks and ethical considerations associated with the use of ChatGPT for anonymization at the two government organizations involved in this study. The literature review highlighted seven key themes concerning AI, anonymization, and ethics: (1) privacy [Cath et al. 2018; Fjeld et al. 2019; Murphy et al. 2021], (2) responsibility [Cath et al. 2018; Fjeld et al. 2019; Murphy et al. 2021], (3) transparency [Cath et al. 2018; Fjeld et al. 2019], (4) algorithm bias and discrimination [Cath et al. 2018; Fjeld et al. 2019; Murphy et al. 2021], (5) human intervention [Fjeld et al. 2019], (6) sustainability [Van Wynsberghe 2021], and (7) future developments and regulation [Coeckelbergh 2019]. Below, we briefly discuss the view of the interviewees and scholars on each of these themes in relation to our study.

In the use of ChatGPT for anonymization, privacy risks arise due to the potential unauthorized use of data to train the language model [Fjeld et al. 2019; Murphy et al. 2021], and input from the model is transmitted to OpenAI and potentially other entities. There is also a risk of re-identification [Henriksen-Bulmer and Jeary 2016]. To quote from Interviewee 1, *“moral responsibility cannot be ascribed to the computer”* and the responsibility for errors lies with the government official making decisions regarding the purpose and means of data processing. As stated by Interviewee 2: *“The person who determines the purpose and means of data processing by the government is the responsible party.”*

Moreover, Interviewee 1 emphasized the importance of AI-systems not being perceived as *“black boxes,”* but rather promoting transparency. With AI being increasingly deployed in making autonomous decisions that can drastically impact both individuals and society [Coeckelbergh 2019], understanding the decision-making process of the AI system is crucial. When the government deploys ChatGPT, OpenAI must comply with transparency obligations imposed by the AI Act [European Parliament 2023], although separate arrangements can be negotiated to ensure the viability of their business model. For instance, Interviewee 2 states: *“Algorithms are often classified as ‘trade secrets.’ Should the government wish to work with ChatGPT, agreements must be made with the ChatGPT vendor that can be met with transparency obligations (as, for example, deviating agreements are made with Microsoft, AWS, Google, and Cisco regarding the purchase of products).”*

Every AI application incorporates bias to align the model, which may lead to discrimination [Fjeld et al. 2019; Murphy et al. 2021]. If trained on incomplete or wrong data, then AI systems can systematically disadvantage

certain groups or individuals (with certain characteristics). This can be problematic, especially if the AI system is deployed to make autonomous decisions [Ferrer et al. 2021]. However, it is the responsibility of politics or governance to determine the acceptable degree of bias. The AI Act emphasizes human intervention in AI systems in the form of oversight. However, human intervention should be meaningful and not be overemphasized given the current labor shortage, as expressed by Interviewee 2: *“I believe that many human interventions are over-valued since you also need to consider the high workload of people and major cutbacks in certain areas, such as Jurisdiction and Healthcare. In those areas, an officer does not always have time to read through the files carefully.”* The interviewee states that ChatGPT might be useful to automate “bulk tasks” or easy repetitive actions in these domains so more time is left for difficult complex issues.

Another relevant ethical issue identified in the literature concerns sustainability. Language models such as ChatGPT demand significant computational power, contributing to substantial greenhouse gas emissions [Van Wynsberghe 2021]. Sustainability considerations for language models, like all other public values, must be integrated into decision-making about these models [Van Wynsberghe 2021]. Finally, in the early stages of technology development in general, assessing the impact of that technology is challenging, making regulation difficult [Collingridge 1982]. The impact becomes clearer later when the technology is actually used, but by that time, the technology may have already had such a significant impact on society that regulating it becomes difficult and costly [Collingridge 1982].

Finally, regulating new technologies such as ChatGPT is challenging due to a lack of information in the early stages and the potential for significant societal impact later on. With the current format in which ChatGPT is offered, ensuring the discussed privacy aspect is simply not feasible, as the input is transmitted to OpenAI and potential third parties. Our first interviewee states the following on this matter: *“When it comes to ChatGPT, one must weigh what is more important: privacy or the business and economic benefits that this tool can offer to the government. This ethical consideration takes on the character of a political decision and should, therefore, be left to the politically responsible party: the minister.”* However, as long as data with personal information remains within the government and is not shared with other parties, the privacy of that data could be safeguarded.

## 5 Discussion

### 5.1 Implications for Research

Various implications for scientific research can be derived from this study. First, this study contributes to the scarce literature on the potential and risks of using ChatGPT specifically in government. It confirms the limited number of studies available on this topic (e.g., Datta et al. [2023]), in the sense that we found that ChatGPT was successful in the anonymization of personally identifiable information in a government organization. The study by Datta et al. [2023], focused specifically on the anonymization of court documents in the Italian government, while this study examined the Dutch context for different types of data, namely, generated Dutch fake data, literature without copyright, and modified literature without copyright. Furthermore, this study confirmed the findings of previous research that identified the threats of privacy and data security issues in the use of ChatGPT [Alsumayt et al. 2024; K. Huang et al. 2023; Nazir and Wang 2023] and the ethical implications of ChatGPT, including privacy and responsibility issues [Mattas 2023].

As far as the conceptual model and focus of this study are concerned (see Section 3.1), the intention of civil servants to use ChatGPT for personal data anonymization in government was out of this study’s scope. This study examined three separate elements that could affect this intention: ChatGPT’s technical operation, ChatGPT’s performance in the anonymization of personal data, and the risks and ethical aspects of employing ChatGPT for personal data anonymization in government. Since this study was exploratory and first examined the three influencing factors on their own, future research should examine how these factors are interrelated and how they both independently and combined influence civil servants’ intention to use ChatGPT for personal data anonymization in government.

Regarding the limitations of this study, GPT-3.5 was used, while a newer version, GPT-4, is available, which may deliver better performance for anonymization. We advise future research to monitor the latest developments and publications related to ChatGPT and similar models, especially considering that previous research found that ChatGPT can also be used for the *deanonymization* of texts and was found to outperform humans [Patsakis and Lykousas 2023]. This study did not examine to what extent the deanonymization of the personal information anonymized by ChatGPT would be possible using this same tool. If the anonymization of personal information in government can easily be undone, at least partially, then the benefits of using ChatGPT in government for anonymization purposes would be reduced significantly.

Additionally, for the anonymization experiments conducted in this research, three different types of texts were utilized: fake data, original literature, and modified literature. However, these texts and thus the test conditions are not precisely analogous to real-world scenarios, diminishing the validity of the findings. Considering the risks and ethical implications of using ChatGPT in government as pointed out in the literature [Alsumayt et al. 2024; K. Huang et al. 2023; Mattas 2023; Nazir and Wang 2023], real personal information as processed in the examined case could not be used for the anonymization experiments conducted in this study. For future research, a possibility would be to conduct anonymization experiments using legal texts, such as judgments, to stay closer to a type of personal information processed in reality.

Finally, it would be valuable to compare the various ways in which ChatGPT or a similar language model could be employed within the government in terms of costs. Determining the current expenses associated with anonymization can provide a baseline. Research can be conducted into the costs of developing a language model independent of OpenAI and the costs associated with the collaborative development of such models with tech companies.

## 5.2 Implications for Practice

This study has a number of practical implications. First, if ChatGPT or similar language models can be employed for personal data anonymization within a government setting, then this could have a tangible impact on addressing labor shortages and reducing the high workload for government officials. Previous research has already argued that AI may be used in the public sector to reduce human workloads and improve work efficiency and user experience [Chen et al. 2021]. AI may also make it easier for government to answer citizens' questions, fill out and search documents, and draft documents [Mehr et al. 2017]. These applications have the potential to enhance government efficiency, enabling employees to dedicate more time to fostering citizen engagement and improving service delivery [Mehr et al. 2017].

Simultaneously, this study shows that in decision-making using ChatGPT in government, it is crucial to consider risks and public values. Through a literature review complemented by two AI expert interviews and informal discussions at the Dutch Ministry of Justice, we found that themes such as privacy, responsibility, transparency, bias, human intervention, and sustainability must be considered. One significant risk in the current form of ChatGPT is a privacy risk, as inputs are stored and forwarded to OpenAI and potentially other parties. This is unacceptable if texts containing personal data are anonymized with ChatGPT. Regulating such new technologies is challenging due to the lack of information in the early stages and the potential significant societal impact later on.

ChatGPT's performance in anonymization shows potential for use within the examined cases and potentially other government organizations. However, the current form of ChatGPT is not secure enough, partly due to the aforementioned privacy risks. The case-study research provided several suggestions on how to handle the risks of implementing ChatGPT for data anonymization in government organizations. First, the government could choose to run the model behind ChatGPT on its own servers, thereby avoiding sharing data with other parties. This would need to be discussed with OpenAI, and currently, OpenAI does not provide this as a service. There is ongoing work on "ChatGPT Business," a subscription for more professional users offering greater control over

user data [OpenAI 2023]. In this setup, the user's input data is not used to train the model. While this option might not be secure enough for anonymizing texts within the case-study organizations we examined, OpenAI may expand its service in the future.

A second option for the government is to develop its own language model. ChatGPT could be deployed for anonymization in these forms with a learning approach, allowing the impact of this technology to become clear on a small scale first. If successful results are achieved on a small scale, then the use of ChatGPT for anonymization can gradually expand. However, this would be a significant undertaking requiring many experts, time, and financial resources. Interviewee 1 is *"a strong advocate for the Dutch government to develop its own language model."* She acknowledges that language models require a substantial amount of training data, and it might be challenging to gather enough training data. *"As long as you don't have access to training data and cannot obtain it legally from open sources, the scenario of having our own ChatGPT is challenging to achieve, and we depend on market parties"* (Interviewee 1). The second interviewee suggests that the idea of the government building such a model entirely on its own might be *"naive"* and emphasizes the importance of exploring collaborations between the government and tech organizations. Interviewee 1 rightly notes that compliance with the Market and Government Act must be observed in the development of a government-owned language model, as the government must not inadvertently compete with market players.

In sum, the proposed method of using ChatGPT for text anonymization cannot be 100 percent successful, as it should always be used in combination with another method, or even the manual reading of documents to complement the required anonymizations that have not been made by ChatGPT. Moreover, discussions with OpenAI or other companies would be necessary to ascertain the costs of running ChatGPT or other large language models on government servers if they are willing to offer this service in the future.

The restrictions mentioned above have implications for policymaking on the use of ChatGPT in government. A traditional way of looking at policymaking is through the four stages of the policy cycle: (1) agenda setting, (2) policy formulation and decision-making, (3) policy implementation, and (4) policy evaluation and termination [Jann and Wegrich 2017]. We argue that policy on the use of ChatGPT in government is still in the first stage of these four, where the problem of text anonymization is being recognized and selected as an issue. Specific policies on text anonymization through ChatGPT in government are yet to be formulated, implemented, and evaluated. Due to the risks and ethical implications of using ChatGPT for text anonymization in government, such a policy may eventually be terminated or not even formulated in the first place.

## 6 Conclusion

This study examines the possibilities, risks, and ethical implications for government organizations to employ ChatGPT in the anonymization of personal data in text documents. In terms of possibilities, experiments with three types of texts (fake data, original literature, and modified literature) show that ChatGPT exhibits strong performance in anonymizing these three types of texts. Fake data is almost flawlessly anonymized, possibly because ChatGPT generated the text itself. Original literature scores lower than fake data, likely due to the complexity of the Old Dutch language we used for the original data. Adapted literature scores mostly comparably to the original literature, with slight improvements, possibly due to adjustments towards modern Dutch.

This research was conducted shortly after the launch of ChatGPT, hence, limited scientific publications were available on this generative AI application. This study is among the first to address the topic of anonymization through ChatGPT, especially in the context of government organizations. This study identified both opportunities and risks and ethical considerations for the implementation of ChatGPT for the anonymization of personal data in government organizations. Practically, this study contributes by discussing several recommendations for handling the risks, such as developing a governmental language model with similar capabilities. This study also emphasizes that due to their potential negative and positive effects, the development of language models should be closely monitored.



## Acknowledgments

We would like to thank the case study participants for their invaluable contributions. The results do not necessarily represent the official views of the case study organization or the views of the Dutch Ministry of Justice and Security.

## References

- M. J. Ahn and Y.-C. Chen. 2022. Digital transformation toward AI-augmented public administration: The perception of government employees and the willingness to use AI in government. *Gov. Inf. Quart.* 39, 2 (2022), 101664.
- A. Alsumayt, Z. M. Alfawaer, N. El-Haggar, M. Alshammari, F. H. Alghamedy, S. S. Aljameel, and M. I. Aldossary. 2024. Boundaries and future trends of chatGPT based on AI and security perspectives. *HighTech Innov. J.* 5, 1 (2024), 129–142.
- A. Askeff, Y. Bai, A. Chen, D. Drain, D. Ganguli, T. Henighan, and N. DasSarma. 2021. A general language assistant as a laboratory for alignment. *arXiv preprint arXiv:2112.00861*.
- C. Cao, J. Zhuang, and Q. He. 2024. LLM-assisted modeling and simulations for public sector decision-making: Bridging climate data and policy insights. In *AAAI Workshop on Public Sector LLMs: Algorithmic and Sociotechnical Design*.
- C. Cath, S. Wachter, B. Mittelstadt, M. Taddeo, and L. Floridi. 2018. Artificial intelligence and the “good society”: The US, EU, and UK approach. *Sci. Eng. Ethics* 24, 2 (2018), 505–528. DOI : <https://doi.org/10.1007/s11948-017-9901-7>
- T. Chen, W. Guo, X. Gao, and Z. Liang. 2021. AI-based self-service technology in public service delivery: User experience and influencing factors. *Gov. Inf. Quart.* 38, 4 (2021), 101520.
- M. Coeckelbergh. 2019. Artificial intelligence: some ethical issues and regulatory challenges. In *Technology and Regulation, 2024*. Open Press TiU, 31–34. DOI : <https://doi.org/10.26116/techreg.2019.003>
- D. Collingridge. 1982. *The Social Control of Technology*. St. Martin’s Press, New York.
- K. Cooper. 2021. OpenAI GPT-3: Everything you need to know. *Springboard*. Retrieved from <https://www.springboard.com/blog/data-science/machine-learning-gpt-3-open-ai/>
- M. T. Dasborough. 2023. Awe-inspiring advancements in AI: The impact of ChatGPT on the field of organizational behavior. *J. Organiz. Behav.* 44, 2 (2023), 177–179.
- P. Datta, B. J. Zahn, L. Attias, G. Salierno, R. Bertè, D. Battisti, and T. Acton. 2023. GiusBERTo: Italy’s AI-based judicial transformation: A teaching case. *Commun. Assoc. Inf. Syst.* 53, 1 (2023), 33.
- J. Deng and Y. Lin. 2022. The benefits and challenges of ChatGPT: An overview. *Front. Comput. Intell. Syst.* 2, 2 (2022), 81–83.
- Y. K. Dwivedi, L. Hughes, E. Ismagilova, G. Aarts, C. Coombs, T. Crick, and A. Eirug. 2019. Artificial Intelligence (AI): Multidisciplinary perspectives on emerging challenges, opportunities, and agenda for research, practice and policy. *Int. J. Inf. Manag.* 57 (2021), 101994. DOI : <https://doi.org/10.1016/j.jinfomgt.2019.08.002>
- European Parliament. 2023. EU AI Act: First regulation on artificial intelligence. Retrieved from <https://www.europarl.europa.eu/news/en/headlines/society/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence>
- D. Faraglia. 2014. Welcome to faker’s documentation! Faker. Retrieved from <https://faker.readthedocs.io/en/master/>
- X. Ferrer, T. Van Nuenen, J. M. Such, M. Coté, and N. Criado. 2021. Bias and discrimination in AI: A cross-disciplinary perspective. *IEEE Technol. Soc. Mag.* 40, 2 (2021), 72–80.
- J. Fjeld, N. Achten, H. Hilligoss, A. Nagy, and M. Srikumar. 2019. *Principled Artificial Intelligence: Mapping Consensus in Ethical and Rights-based Approaches to Principles for AI*. Retrieved from [https://dash.harvard.edu/bitstream/handle/1/42160420/HLS%20White%20Paper%20Final\\_v3.pdf?sequence=1&isAllowed=y](https://dash.harvard.edu/bitstream/handle/1/42160420/HLS%20White%20Paper%20Final_v3.pdf?sequence=1&isAllowed=y)
- B. Friedman. 1997. *Human Values and the Design of Computer Technology*. Cambridge University Press, Cambridge.
- F. Fui-Hoon Nah, R. Zheng, J. Cai, K. Siau, and L. Chen. 2023. Generative AI and ChatGPT: Applications, challenges, and AI-human collaboration. *J. Inf. Technol. Case Applic. Res.* 25, 3 (2023), 277–304.
- C.-E. González-Gallardo, E. Boros, N. Girdhar, A. Hamdi, J. G. Moreno, and A. Doucet. 2023. Yes but... Can ChatGPT Identify Entities in Historical Documents? *arXiv preprint arXiv:2303.17322*.
- W. Hariri. 2023. Unlocking the potential of ChatGPT: A comprehensive exploration of its applications, advantages, limitations, and future directions in natural language processing. *arXiv preprint arXiv:2304.02017*.
- F. Hassan and J. Domingo-Ferrer. 2018. Anonymization of unstructured data via named-entity recognition. In *15th International Conference on Modeling Decisions for Artificial Intelligence*.
- F. Hassan, D. Sanchez, and J. Domingo-Ferrer. 2019. Automatic anonymization of textual documents: Detecting sensitive information via word embeddings. In *18th IEEE International Conference on Trust, Security and Privacy in Computing and Communications/13th IEEE International Conference on Big Data Science and Engineering (TrustCom/BigDataSE’19)*.
- J. Henriksen-Bulmer and S. Jeary. 2016. Re-identification attacks—A systematic literature review. *Int. J. Inf. Manag.* 36, 6 (2016), 1184–1192.
- J. Huang and K. Huang. 2023. ChatGPT in government. In *Beyond AI: ChatGPT, Web3, and the Business Landscape of Tomorrow*. Springer, 271–294.



- K. Huang, F. Zhang, Y. Li, S. Wright, V. Kidambi, and V. Manral. 2023. Security and privacy concerns in ChatGPT. In *Beyond AI: ChatGPT, Web3, and the Business Landscape of Tomorrow*. Springer, 297–328.
- W. Jann and K. Wegrich. 2017. Theories of the policy cycle. In *Handbook of Public Policy Analysis*. Routledge, 69–88.
- I. Lindgren, C. Ø. Madsen, S. Hofmann, and U. Melin. 2019. Close encounters of the digital kind: A research agenda for the digitalization of public services. *Gov. Inf. Quart.* 36, 3 (2019), 427–436.
- R. Lowe and J. Leike. 2022. Aligning language models to follow instructions. Retrieved from <https://openai.com/research/instruction-following>
- M. E. Mamalis, E. Kalampokis, A. Karamanou, P. Brimos, and K. Tarabanis. 2023. Can large language models revolutionize open government data portals? A case of using ChatGPT in statistics. gov. scot. In *27th Pan-Hellenic Conference on Progress in Computing and Informatics*.
- P. S. Mattas. 2023. ChatGPT: A study of AI language processing and its implications. *Int. J. Res. Pub. Rev.* 4, 2 (2023), 435–440. DOI: <https://doi.org/10.55248/gengpi.2023.4218>
- H. Mehr, H. Ash, and D. Fellow. 2017. Artificial intelligence for citizen services and government. *Ash Cent. Democr. Gov. Innov. Harvard Kennedy Sch.* Aug. (2017), 1–12.
- Ministerie van Justitie en Veiligheid. 2018. Strategisch leveranciersmanagement Microsoft rijik (SLM Microsoft). *Rijksoverheid.nl*. Retrieved from <https://www.rijksoverheid.nl/documenten/publicaties/2018/11/12/strategisch-leveranciersmanagement-microsoft-rijik-slm-microsoft>
- B. Mohit. 2014. Named entity recognition. In *Natural Language Processing of Semitic Languages*. Springer, 221–245.
- K. M. Murphy, E. Di Ruggiero, R. E. Upshur, D. J. Willison, N. R. Malhotra, J. Cai, and J. L. Gibson. 2021. Artificial intelligence for good health: A scoping review of the ethics literature. *BMC Med. Ethics* 22, 1 (2021). DOI: <https://doi.org/10.1186/s12910-021-00577-8>
- A. Nazir and Z. Wang. 2023. A comprehensive survey of ChatGPT: Advancements, applications, prospects, and challenges. *Meta-radiology* 1, 2 (Sept. 2023) 100022.
- Netherlands Forensic Institute. Ministry of Justice and Security. (n.d.). Home. Retrieved from <https://www.forensicinstitute.nl/>
- OpenAI. 2023. New ways to manage your data in ChatGPT. Retrieved from <https://openai.com/blog/new-ways-to-manage-your-data-in-chatgpt>
- OpenAI. 2024. Tokenizer. Retrieved from <https://platform.openai.com/tokenizer>
- L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. Wainwright, P. Mishkin, and A. Ray. 2022. Training language models to follow instructions with human feedback. *Advan. Neural Inf. Process. Syst.* 35 (2022), 27730–27744.
- C. Patsakis and N. Lykousas. 2023. Man vs the machine in the struggle for effective text anonymisation in the age of large language models. *Scient. Rep.* 13, 1 (2023), 16026.
- J. Paul, A. Ueno, and C. Dennis. 2023. ChatGPT and consumers: Benefits, pitfalls and future research agenda. In *Int. IJC* 47, 5 (2023), 1213–1225. Wiley Online Library.
- A. Radford, K. Narasimhan, T. Salimans, and I. Sutskever. 2018. *Improving Language Understanding by Generative Pre-training*. Retrieved from [https://cdn.openai.com/research-covers/language-unsupervised/language\\_understanding\\_paper.pdf](https://cdn.openai.com/research-covers/language-unsupervised/language_understanding_paper.pdf)
- M. Salvagno, F. S. Taccone, and A. G. Gerli. 2023. Can artificial intelligence help for scientific writing? *Crit. Care* 27, 1 (2023), 75.
- Y. Shen, L. Heacock, J. Elias, K. D. Hentel, B. Reig, G. Shih, and L. Moy. 2023. ChatGPT and other large language models are double-edged swords. *Radiology* 307, 2 (2023), e230163.
- G. Szarvas, R. Farkas, and R. Busa-Fekete. 2007. State-of-the-art anonymization of medical records using an iterative machine learning framework. *J. Am. Med. Inform. Assoc.* 14, 5 (2007), 574–580. DOI: <https://doi.org/10.1197/j.jamia.m2441>
- D. Toll, I. Lindgren, U. Melin, and C. Ø. Madsen. 2020. Values, benefits, considerations and risks of AI in government: A study of AI policies in Sweden. *JeDEM-eJ. eDemoc. Open Gov.* 12, 1 (2020), 40–60.
- D. Valle-Cruz and R. Sandoval-Almazan. 2018. Towards an understanding of artificial intelligence in government. In *19th Annual International Conference on Digital Government Research: Governance in the Data Age*.
- E. A. Van Dis, J. Bollen, W. Zuidema, R. van Rooij, and C. L. Bockting. 2023. ChatGPT: Five priorities for research. *Nature*, 614, (7947), 224–226.
- C. van Noordt and G. Misuraca. 2022. Artificial intelligence for the public sector: Results of landscaping the use of AI in government across the European Union. *Gov. Inf. Quart.* 39, 3 (2022), 101714.
- A. Van Wynsberghe. 2021. Sustainable AI: AI for sustainability and the sustainability of AI. *AI Ethics*, 1 (2021), 213–218. DOI: <https://doi.org/10.1007/s43681-021-00043-6>
- V. Venkatesh, M. G. Morris, G. B. Davis, and F. D. Davis. 2003. User acceptance of information technology: Toward a unified view. *MIS Quart.* 27, 3 (2003), 425–478. Retrieved from <http://www.jstor.org/stable/pdfplus/30036540.pdf?acceptTC=true>
- V. Venkatesh, J. Y. L. Thong, F. K. Y. Chan, P. J.-H. Hu, and S. A. Brown. 2011. Extending the two-stage information systems continuance model: Incorporating UTAUT predictors and the role of context. *Inf. Syst. J.* 21, 6 (2011), 527–555. DOI: [10.1111/j.1365-2575.2011.00373.x](https://doi.org/10.1111/j.1365-2575.2011.00373.x)
- V. Venkatesh, J. Y. L. Thong, and X. Xu. 2012. Consumer acceptance and use of information technology: Extending the unified theory of acceptance and use of technology. *MIS Quart.* 36, 1 (2012), 157–178.

- B. Wang, A. Wang, F. Chen, Y. Wang, and C.-C. J. Kuo. 2019. Evaluating word embedding models: Methods and experimental results. *APSIPA Trans. Sig. Inf.* July (2019), e19.
- M. Williams and T. Moser. 2019. The art of coding and thematic exploration in qualitative research. *Int. Manag. Rev.* 15, 1 (2019), 45–55.
- L. Yang and J. Wang. 2023. Factors influencing initial public acceptance of integrating the ChatGPT-type model with government services. *Kybernetes* Aug. (2023). DOI: <https://doi.org/10.1108/K-06-2023-1011>

Received 29 January 2024; revised 8 April 2024; accepted 22 June 2024