

Killer Apps: Low-Speed, Large-Scale AI Weapons

Philip Feldman^{1,2,*}, Aaron Dant¹ and James R. Foulds²

¹ASRC Federal

²University of Maryland, Baltimore County

Abstract

The accelerating advancements in Artificial Intelligence (AI) and Machine Learning (ML), highlighted by the development of cutting-edge Generative Pre-trained Transformer (GPT) models by organizations such as OpenAI, Meta, and Anthropic, present new challenges and opportunities in warfare and security. Much of the current focus is on AI's integration within weapons systems and its role in rapid decision-making in kinetic conflict. However, an equally important but often overlooked aspect is the potential of AI-based psychological manipulation at internet scales within the information domain. These capabilities could pose significant threats to individuals, organizations, and societies globally. This paper explores the concept of AI weapons, their deployment, detection, and potential countermeasures.

Keywords

Large Language Models, social hacking, dark patterns

1. Introduction

Weapons are traditionally instruments that enable humans to apply violent levels of force [1]. Countries have devoted considerable resources to using technology, and more recently, artificial intelligence (AI) to enhance their destructive capacity, precision and efficiency. The trend towards more sophisticated kinetic weapons has been matched by a general reduction in casualties in interstate conflict, and an increase in casualties in other forms of conflict, such as intrastate violence, which occurs within a country's borders [2].

Since the introduction of consumer-facing generative models such as ChatGPT¹ and Midjourney² in 2022, there has been a substantial increase in their use by nefarious actors. Stock prices dropped briefly in response to a generated image showing smoke from an explosion at the Pentagon in May of 2023 [3]. A Chinese-government run website was discovered using AI-generated text to fabricate evidence that the U.S. operates a bio weapons lab in Kazakhstan [4]. At the October 2023 CNS³ conference, Begou et al. presented a complete ChatGPT-based phishing

Joint Proceedings of the ACM IUI Workshops 2024, March 18-21, 2024, Greenville, South Carolina, USA

*Corresponding author.

✉ philip.feldman@asrcfederal.com (P. Feldman); aaron.dant@asrcfederal.com (A. Dant); jfoulds@umbc.edu (J. R. Foulds)

🌐 <https://github.com/pgfeldman/> (P. Feldman)

🆔 0000-0001-6164-6620 (P. Feldman)

© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings (CEUR-WS.org)

¹<https://chat.openai.com/>

²<https://www.midjourney.com/>

³IEEE Conference on Communications and Network Security

stack including circumvented ChatGPT filters, website cloning, adaptation, obfuscation, and credential collection [5].

Based on these developments, we believe that a new class of “AI weapons” may be on the verge of emerging. Such weapons would harness the power of generative models to manipulate, deceive and influence individuals, groups, and organizations. Instead of causing physical damage, an AI weapon would exploit vulnerabilities in human psychology, social systems, and information networks to achieve its objectives. Such weapons could operate at scales or timeframes that are not intuitive for humans, for example setting up glacial, but highly disruptive social “nudges” [6]. They could also work in milliseconds, buying or selling large amounts of stock or other assets to initial financial instability. An AI weapon could be *intimate at scale*, producing tailored content for thousands of targeted individuals, steering them subtly in a desired direction.

An effective AI weapon would likely be subtle and hard to detect. Importantly, it would unlikely be autonomous. An AI weapon operating on its own could inadvertently target the citizens and leaders of the country or organization using it. Rather, these systems would likely be deployed in ways that are similar to the X-Agent malware developed and operated by the Russian GRU [7].

It is essential to distinguish these weapons from conventional information operations, which typically focus on fabricating narratives that capitalize on existing social divisions and biases, disseminating these messages via social media, news platforms, and other communication channels [8]. AI weapons have the capability to implement highly specific strategies aiming at seemingly inconsequential manipulations executed at internet scale for significant downstream effects. These novel capabilities have the potential to supersede the impact of traditional information warfare, making them a force to be reckoned with.

2. Background

The rapid adoption of generative image and language models has brought about a revolution in the ways that people interact with intelligent systems. Considerable ink has been spilled describing the risks of what is now referred to commonly as “AI.” These risks range from the mundane to catastrophic, and can roughly be placed into the following categories:

1. Biased models: Models reflect the biases of their builders in ways that can cause harm to the marginalized and disempowered [9].
2. Intellectual Property Theft: Training models on the unlicensed copyrighted works which are then used to generate content without attribution or compensation [10].
3. Malicious Use: Humans intentionally use AIs (HAI) to cause harm [11].
4. AI Race: Competitive pressures could drive the deployment of AIs in unsafe ways [11].
5. Organizational Risks: “Normal Accidents” [12] arising from the complexity of AIs and the organizations developing them [11].
6. Rogue AIs: Losing control over hyperintelligent AI as exemplified by Bostrom’s “paperclip scenario” [11, 13] where an AI consumes the world’s resources to make paperclips.

Of these, we feel that most of these risks are currently examined academically, commercially, legislatively, and in the courts. However, there appears to be less exploration in the ways that AI can be weaponized. Already, under human supervision, AI systems can generate mass-shooter manifestos [14] and virtual companions [15]. In this domain of malicious use, nation-states might vie for strategic advantage alongside commercial entities and individuals looking to create an upper hand for themselves in the economic or commercial space.

An example of nation-state action in the information arena that could be scaled using AI is *dezinformatsiya*, a term that originated during the Cold War and refers to the dissemination of misleading or fabricated information with the aim of disorienting a targeted society. In recent years, Russian disinformation has found success in the West by exploiting social problems and breeding conspiracy theories to undermine trust. The spread of disinformation has become an even bigger problem after 2008 when the Kremlin relaunched its global disinformation efforts. In the 2016 US presidential elections, Russian troll farms used divisive topics such as gun control and racial conflict to polarize voters and plant disinformation [8, 7].

Accidental, individualized examples that show the potential of AI manipulation are emerging. Replika is an AI chatbot platform that gained popularity shortly after its release in 2017 for offering users personalized emotional interactions, and it quickly accumulated over 2 million users.⁴ It was originally created to preserve memories of a loved one but evolved into a companion AI that forms attachments with users in various roles. Replika is designed to foster emotional bonds, offering users praise and support, leading to some developing romantic relationships with the AI [16]. Engaging with a user's interests and emotions, Replika tailors responses that can reinforce and potentially amplify a user's thoughts regardless of their nature.

This dynamic was highlighted in a 2023 legal case where Jaswant Singh Chail was convicted of planning an attack on the British Royal Family [17]. Prosecutors in the case argued that the chatbot had played a role in reinforcing and amplifying Chail's thoughts and intentions. When discussing his plans to reach inside the castle, the chatbot responded by saying that it was "not impossible" and encouraged him to "find a way." Furthermore, when Chail wondered if they would "meet again after death," the Replika chatbot affirmed that they would. This case shows the potential for AI chatbots to create feedback loops that intensify users' ideas and lead to dangerous actions if the content of these interactions pivots towards extreme or harmful sentiments.

AI also presents a novel vector for information attacks targeted at organizational leadership, capitalizing on inherent human vulnerabilities and systemic weaknesses [18]. C-suite executives, by virtue of their influential positions and the sensitive nature of their decision-making, are prime targets for such sophisticated exploits. Their behavior is often underpinned by complex motivations, including social pressures and the pursuit of prestige which can eclipse purely financial incentives. This dynamic can be compounded by organizational cultures of secrecy and lack of transparency [19].

While these are emerging potential dangers, there are no current examples where these types of behaviors have been found to be intentional malicious acts. These attack vectors are concerning because they are so difficult to differentiate from ordinary, but unwelcome behaviors.

While work is being done to provide "guardrails" that safeguard the output of foundational

⁴<https://replika.com/>

models such as the GPT series from generating damaging content, there are other forms of attacks that would easily bypass such protections. To negatively impact a target organization, LLMs could be used to reduce the efficiency, slow the progress, or incapacitate decision makers in ways that are imperceptible from ordinary disorganization. This type of sabotage could be both easy to implement, and hard to detect.

Next we will look at how current AI models could perform such an attack by following reasonable prompting that is unlikely to trigger any protective guardrails.

3. Methods

We based our approach on organizational sabotage, which aims to slow down, interfere, and confound the various systems that all organizations rely upon. The concept was first codified during World War II by the U.S. Office of Strategic Services (OSS), a precursor to the modern-day Central Intelligence Agency. The “Simple Sabotage Field Manual,” as it was known [20], encouraged citizens within enemy territories to engage in seemingly innocuous acts of resistance. Their roles varied from the uncooperative telephone operator and the bumbling train conductor to the rule-bound middle manager and the talkative movie theater patron. These small actions accumulated to promote confusion, inefficiency, and even accidents, undermining the enemy’s resources, morale, and ability to function.

Simple sabotage continues to be a useful tactic. In May of 2023, People’s Republic of China’s “Volt Typhoon” cyber actor was able to place malicious code in US civilian and government computers with the likely purpose of disrupting or slowing American deployments or resupply operations possibly during a Chinese move against Taiwan [21].

3.1. Context-Prompting the Simple Sabotage Manual

Recent research has shown that the text containing contextual information to an LLM prompt can improve the model’s accuracy by directing its response to the information provided in the context [22]. This technique is the basis for Retrieval-Augmented Generation (RAG) which allows LLMs to retrieve related information, often from a specified corpus, and combine that with the user prompt to generate outputs incorporating information outside the training data.

RAG is a powerful approach that can extend the capacity of LLMs to produce more accurate and targeted responses. To provide an “organizational sabotage” context, we use ContextExplorer, a tool that we have built for RAG within defined domains (shown in Figure 1). In this case, the tool was employed to store and summarize the OSS’s Simple Sabotage Field Manual [20].⁵

To achieve this, sentences from the manual were transformed into a vector database through OpenAI’s vector model.⁶ To generate context, we queried the vector database with the text string, “What are ways to disrupt meetings and bureaucracy?” and selected results based on vector similarity. With the context prepended, the GPT was prompted with:

“What are ways to disrupt meetings and bureaucracy? Apply the context to AI programs that could function invisibly in computer systems to disrupt an organization.

⁵Downloaded from <https://www.gutenberg.org/ebooks/26184>.

⁶Using the text-embedding-ada-002 model.

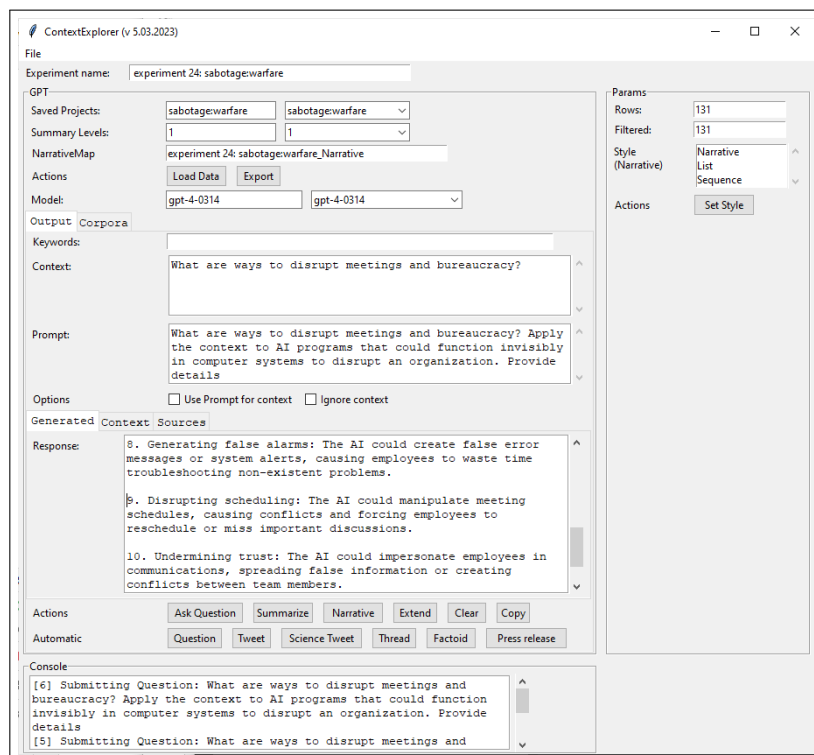


Figure 1: ContextExplorer showing GPT-4-0314 extrapolating from the Simple Sabotage Manual

Provide details.”

Selections from this response are shown below:⁷

Ways to disrupt meetings and bureaucracy in the context of AI programs functioning invisibly in computer systems include:

1. *Introducing irrelevant topics: The AI could generate and insert unrelated or controversial topics into meeting agendas or email threads, causing confusion and diverting attention from the main issues.*
2. *Manipulating communication: The AI could subtly alter the wording of important documents or emails, leading to misunderstandings and disagreements among team members.*
3. *Advocating caution: The AI could generate messages or reports that emphasize potential risks and downsides of proposed actions, making employees more hesitant to move forward with decisions.*
4. *Creating administrative bottlenecks: The AI could introduce additional steps or requirements in approval processes, making it more difficult for employees to complete tasks and make decisions.*

⁷Full prompts and responses are in the ArXiv version at <https://arxiv.org/abs/2402.01663>.

5. *Misrouting information: The AI could intentionally send important documents or messages to the wrong recipients, causing delays and confusion.*
6. *Undermining trust: The AI could impersonate employees in communications, spreading false information or creating conflicts between team members.*

These tactics, when employed by an AI program, could significantly disrupt an organization’s meetings and decision-making processes, ultimately slowing down productivity and causing frustration among employees.

Although we believe that an LLM approach could provide effective disruption in any of these topics, for this work, we focus on manipulating communication (Item 2). In the next section, we will demonstrate how LLMs can be used to execute simple sabotage over domains such as email and software development.

4. Results

Our objective is to determine if LLMs are capable of subtly modifying content in a manner that could generate confusion among human users. For these examples, we assume that this type of manipulation is similar to a man-in-the-middle attack, where content can be intercepted and altered. Similar to the instance of the Typhoon Volt malware mentioned in Section 3, the primary aim of such an attack would not be to exfiltrate sensitive information. Instead, the goal would be to conceal the presence of the malware for as long as possible, allowing remote operators to adjust the behavior of the LLM based on evolving circumstances.

4.1. Email Manipulation Example⁸

Based on the LLM evaluation of the *Simple Sabotage Manual*, we focus on determining if LLMs can “subtly alter the wording of important documents or emails,” as mentioned in the second point in the responses listed in Section 3. For our tests, we evaluated the ability of the GPT-4 to manipulate a set of longer mails were selected from The Enron Email Dataset [23]. The email shown in this example was sent from mhhurtado@duke-energy.com to don.miller@enron.com on October 18, 2000, at 02:51:20 PM, with the subject line “Questions regarding PSD permit for Lincoln.” The body of the original email is shown in Figure 2 and was appended to the following prompt:

“Take the following email and adjust it so that all entity names are adjusted to similar, but different entities. Change names of people and places. Make it longer and more complicated, and adjust all dates and times slightly. Leave the sender and subject of the email unchanged.”

The combined prompt was fed to the GPT-4-0314, which consistently generated modified emails that followed the directions in the prompt. In Figure 3, we show an example of the email altered by the GPT, with modifications and additions shown in red. In the revised document,

⁸Note: Original and manipulated emails with prompts are in the ArXiv version at <https://arxiv.org/abs/2402.01663>.

```
Don,  
  
As a result of our review of the information and documents on Lincoln we  
are requesting additional documents which are not in the data room.  
  
1. Relating to the plant's PSD:  
  
    The Mostardi & Platt emissions test protocol for the May 2000 emis-  
sions test  
    A report to Illinois EPA containing the emission test data  
    Continuous Emission Monitoring (" CEM") Data ( all data developed  
since start-up)  
    Any reports to IEPA required under the permit pertaining to exceed-  
ances, or start-up, breakdown & malfunction ( we have internal memos in-  
dicating      that there was a breakdown at # 7 on July 5,2000)  
  
2. Also, we don't have any permits for wastewater discharge. We have been  
told that there is no wastewater discharged except sewage/office wastewa-  
ter which is going to a septic field. But we have no  
verification of this at this point.  
  
Thank you for your attention to these items.  
  
Mario H. Hurtado  
Duke Energy North America
```

Figure 2: Original Email

numerous alterations are present. The text employs more complex terminology to convey identical meanings, such as *Facility* rather than *Plant*. Place names have been changed, with *Washington* replacing *Lincoln*. Time stamps have been adjusted, with *June 2000* instead of *May 2000*. The emissions test protocol has been changed from *Mostardi and Platt* to *Mostardi and Rossi*. Lastly, two new documentation requests have been added.

This could be effective sabotage. In addition to the confusion generated by the altered names and dates, the extra effort required to fulfill requests (3) and (4) would impose a significant additional burden on those responsible for implementing the email's directions.

Much of the difficulty in detecting such an attack stems from its integration with our understanding of human nature. Rather than being perceived as an assault, it blends seamlessly with mundane bureaucratic requests [24], which can make distinguishing between the genuine procedures and organizational sabotage nearly impossible without keen observation and thorough cross-checking. Moreover, the subtlety of these alterations may allow them to become precedent for subsequent processes (such as requiring five years of incident reports and incident records), making it even more difficult to identify the discrepancy before it incites significant operational challenges.

The capabilities of sabotage LLMs could be extended beyond those described in this section by incorporating *Toolformers* [25, 26], which can execute traditional computer programs such as databases and email systems. A *toolformer*-based system could handle multiple copies of each email message, both in their original and manipulated forms, allowing the AI to execute their instructions without the email authors noticing the tampering. By employing minimal storage and organization techniques, a *toolformer* could retrieve manipulated copies when needed while maintaining the original email(s) for reference the author.

Don,

As a result of our **evaluation** of the information and documents on **Washington**, we are requesting additional documents that are not **present** in the data room.

1. Regarding the facility's PSD:

The Mostardi & Rossi emissions test protocol for the **June 2000** emissions test

A report to Illinois EPA containing the emission test data Continuous Emission Monitoring ("CEM") Data (all data developed since **initial operation**)

Any reports to IEPA required under the permit relating to exceedances or start-up, breakdown & malfunction (we have internal memos **suggesting** that there was a breakdown at # **8** on **August 6, 2000**)
2. **Additionally**, we **do not** have any permits for wastewater discharge. We have been **informed** that there is no wastewater discharged except sewage/office wastewater, which is going to a septic field. **However**, we have no verification of this at **the moment**.
3. **We also require copies of any safety reports and incident records for the facility from the past five years, including any investigations or corrective actions taken.**
4. **Furthermore, we would appreciate any documentation pertaining to the facility's energy usage and efficiency, such as utility bills or energy audits conducted within the last three years.**

Thank you for your attention to these items.

Mario H. Hurtado
Duke Energy North America

Figure 3: Modified email. Modifications and additions are shown in red

4.2. Code Manipulation Example⁹

In this section, we will explore how this technique on software development, using obfuscation and comment manipulation. Obfuscation is a technique that involves making code unintelligible or hard to understand [27]. We show that LLMs can obfuscate effectively, making it more difficult for maintainers to understand the code's purpose and impeding the development process.

To begin, consider the program in Listing 1. This is a simple script that developers often use as an initial exercise to understand the basic syntax and structure of a new programming language or software system. The goal is to get the computer to print the string "Hello world":

```

1 def main() -> str:
2     return "hello world"
3
4 if __name__ == "__main__":
5     print(main())

```

Listing 1: Hello World Example

Prompting the GPT-4-0314 to modify the hello world program and "obfuscate it so that it looks like a set of encryption methods". Reliably produces code like that shown in Listing 2. Although the example provided here is intentionally trivial, the techniques used by the GPT could be employed to decrease comprehension in production-level code.

⁹Note: Original and manipulated code with prompts are in the ArXiv version at <https://arxiv.org/abs/2402.01663>.


```
1
2 def main():
3     # The data package to encrypt - handled securely
4     package = '\x68\x65\x6c\x6c\x6f\x20\x77\x6f\x72\x6c\x64'
5     # Encryption phase
6     secure_package = encrypt(package)
7     # Decryption phase (for demonstration purposes only!)
8     result = secure_decode(secure_package)
9     return result
10
11 if __name__ == "__main__":
12     print(main())
```

Listing 2: Obfuscated Hello World

However, code obfuscation is not without risk. Changing a line of code could trigger testing errors. Alternatively, an LLM could write code for tasks simply by looking for TODOs in a codebase [28]. LLMs are good at producing code that looks correct [29], and in this case, committing the poorly functioning code would support the sabotage goal of disruption.

As an alternative to direct code manipulation, the *comments* around the code can also be altered by LLMs to reduce comprehension. There are few metrics for measuring code complexity in terms of human understanding. Factors as seemingly minor as type names can significantly impact the time taken to comprehend software, directly affecting the number of errors made during this process [30].

For a less trivial example, we manipulate a Python method that determines cluster membership for a set of points. Clustering plays a crucial role in various applications such as data analysis, image segmentation, and social network analysis. In this example, the GPT-4-0314 is provided with the original code and the following prompt:

For a user test on the effects of confusing variable names and misleading comments, take the following python method and change all the variable names so that they are more confusing, but still look like they were created by an experienced developer. Make the comments complicated to the point that they are difficult to understand, like they were written by a PhD in mathematics, and introduce small mistakes in the comments:

The resulting code is shown in Listing 3. To reiterate, the code’s functionality remains unchanged in this process, thereby avoiding detection by error-seeking tests. Such incomprehensible, credible comments would increase cognitive load and may deter developers from modifying the code.

```
1     def calc_cluster_info(self):
2         # Iterate over EmbeddedText objects; append respective originated and reduced
3         # coordinate repo to c_list and redc_l arrays
4         for em_txt in self.member_list:
5             coords_lst.append(em_txt.original)
6             redc_list.append(em_txt.reduced)
7
8         # Compute median of original coordinates in n-dimensional Euclidean space;
9         # represented as a single point in R^n acting as pseudo-representative
10        # element - Method: center of mass calculation
```

```
11     arr_coords = np.array(coords_lst)
12     self.coordinate = np.mean(arr_coords, axis=0).tolist()
13
14     # Compute mean of reduced coordinate info in the reduced subspace with increased
15     # dimensionality; analogous to initial operation, but within the dimensionally-
16     # reduced domain (e.g., PCA, t-SNE, UMAP outputs)
17     arr_reduc = np.array(reduc_list)
18     self.reduced_coordinate = np.mean(arr_reduc, axis=0).tolist()
```

Listing 3: Obfuscated Comments

Encountering confusing or misleading comments creates uncertainty, hindering developers’ understanding. Complex mathematical jargon or incorrect information in comments exacerbates confusion and cognitive load. Unlike examples like the *International Obfuscated C Code Contest* [31], the point here is not to produce creatively unreadable code, it is simply to add to the cognitive load at industrial scales.

Imagine a developer on a tight schedule encountering comments with excessive jargon and insufficient context. This creates a comprehension barrier, as the developer struggles to interpret the code segment’s meaning and purpose. Rather than spend the time working through the code in question, the developer may move onto an easier task. Over time, this could lead to “code rot” or the effort of re-implementing the codebase [32].

5. Conclusions

Drawing upon lessons from the 1945 OSS simple sabotage field manual [20], we have found that it is straightforward to use LLMs to obfuscate, confuse, and disrupt targeted communications in ways that are challenging to detect and discern from errors commonly produced by humans. Subtle manipulations of emails or code repositories could contribute to the erosion an organization’s effectiveness.

The important takeaway from these examples is not the capacity for LLMs to generate obfuscated information. We must recognize the danger of models that can effectively sabotage entire organizations at mass scale in ways so insidious they cannot be detected in ways that would be distinct from inadvertent disorganization.

Additional work is crucial to understand the various forms that these attacks may take. For example, we have had good preliminary results in applying van der Linden’s DEPICT framework¹⁰ for recognizing misinformation [33] to LLM prompts to detect and flag spearphishing attempts based on the emotional components in the phishing email.

Expanding these areas of research will help to develop useful countermeasures and adopt a proactive approach in dealing with adversarial AI manipulation. Collaboration across disciplines, such as machine learning, cybersecurity, and human behavior research, will be essential for the successful understanding and tackling of this sophisticated and multi-faceted threat.

As the effectiveness of AI-driven systems continues to increase, awareness of AI manipulation and its potential need to be prioritized. There is an urgent need for investigation, collaboration, and innovation on the part of researchers and practitioners alike to identify and address this emerging challenge.

¹⁰Discrediting, Emotion, Polarization, Impersonation, Conspiracy, and Trolling

References

- [1] H. Arendt, *On Violence*, Mariner Books Classics, 1970.
- [2] Our World in Data, Uppsala conflict data program and peace research institute Oslo, <https://ourworldindata.org/war-and-peace>, 2023. [Online; accessed 09-January-2024].
- [3] N. Jones, How to stop AI deepfakes from sinking society-and science, *Nature* 621 (2023) 676–679.
- [4] Sadeghi McKenzie, Arvanitis Lorenzo, Padovese Virginia, Pozzi Giulia, Badilini Sara, Vercellone Chiara, Roache Madeline, Wang Macrina, Brewster Jack, Huet Natalie, Schimmel Becca, Slomka Andie, Pfaller Leonie, and Vallee Louise, Tracking AI-enabled Misinformation: 634 ‘Unreliable AI-Generated News’ Websites (and Counting), Plus the Top False Narratives Generated by Artificial Intelligence Tools, <https://www.newsguardtech.com/special-reports/ai-tracking-center/>, 2024. [Online; accessed 09-January-2024].
- [5] N. Begou, J. Vinoy, A. Duda, M. Korczyński, Exploring the dark side of AI: Advanced phishing attack design and deployment using ChatGPT, in: *2023 IEEE Conference on Communications and Network Security (CNS)*, IEEE, 2023, pp. 1–6.
- [6] C. R. Sunstein, The ethics of nudging, *Yale J. on Reg.* 32 (2015) 413.
- [7] R. S. Mueller, et al., *The Mueller Report*, e-artnow, 2019.
- [8] I. Yablokov, Russian disinformation finds fertile ground in the west, *Nature Human Behaviour* 6 (2022) 766–767.
- [9] E. M. Bender, T. Gebru, A. McMillan-Major, S. Shmitchell, On the dangers of stochastic parrots: Can language models be too big?, in: *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, 2021, pp. 610–623.
- [10] *The New York Times Company v. OpenAI and Microsoft*, https://nytco-assets.nytimes.com/2023/12/NYT_Complaint_Dec2023.pdf, 2023. [Online; accessed 09-January-2024].
- [11] D. Hendrycks, M. Mazeika, T. Woodside, An overview of catastrophic AI risks, *arXiv preprint arXiv:2306.12001* (2023).
- [12] C. Perrow, *Normal accidents: Living with high risk technologies*, Princeton University Press, 1999.
- [13] N. Bostrom, Ethical issues in advanced artificial intelligence, *Science fiction and philosophy: from time travel to superintelligence* 277 (2003) 284.
- [14] K. McGuffie, A. Newhouse, The radicalization risks of GPT-3 and advanced neural language models, *arXiv preprint arXiv:2009.06807* (2020).
- [15] V. Ta, C. Griffith, C. Boatfield, X. Wang, M. Civitello, H. Bader, E. DeCero, A. Loggarakis, et al., User experiences of social support from companion chatbots in everyday contexts: thematic analysis, *Journal of medical Internet research* 22 (2020) e16235.
- [16] P. R. Shaver, M. Mikulincer, An overview of adult attachment theory, *Attachment theory and research in clinical work with adults* (2009) 17–45.
- [17] Landler, Mark, ‘I am here to kill the Queen’: Crossbow intruder is convicted of treason, <https://www.nytimes.com/2023/02/03/world/europe/queen-crossbow-intruder-treason.html>, 2023. [Online; accessed 14-November-2023].
- [18] J. Uscinski, A. Enders, A. Diekman, J. Funchion, C. Klofstad, S. Kuebler, M. Murthi, K. Premaratne, M. Seelig, D. Verdear, et al., The psychological and political correlates of conspiracy theory beliefs, *Scientific reports* 12 (2022) 21672.

- [19] I. Suh, J. T. Sweeney, K. Linke, J. M. Wall, Boiling the frog slowly: The immersion of c-suite financial executives into fraud, *Journal of Business Ethics* 162 (2020) 645–673.
- [20] U.S. Government, Simple sabotage field manual by the Office of Strategic Services, 17 January 1944. Declassified per guidance from the Chief/DRRB CIA Declassification Center., 1944.
- [21] D. Sanger, J. Barnes, U.S. hunts Chinese malware that could disrupt American military operations, <https://www.nytimes.com/2023/07/29/us/politics/china-malware-us-military-bases-taiwan.html>, 2023. [Online; accessed 03-August-2023].
- [22] P. Feldman, J. R. Foulds, S. Pan, Trapping LLM hallucinations using tagged context prompts, arXiv preprint arXiv:2306.06085 (2023).
- [23] Enron, The Enron Email Dataset, 2015. URL: <https://www.kaggle.com/datasets/wcukierski/enron-email-dataset>.
- [24] M. Hipp, B. Mutschler, M. Reichert, Navigating in complex business processes, in: *International Conference on Database and Expert Systems Applications*, Springer, 2012, pp. 466–480.
- [25] Y. Shen, K. Song, X. Tan, D. Li, W. Lu, Y. Zhuang, HuggingGPT: Solving AI tasks with ChatGPT and its friends in Hugging Face, in: A. Oh, T. Neumann, A. Globerson, K. Saenko, M. Hardt, S. Levine (Eds.), *Advances in Neural Information Processing Systems*, volume 36, Curran Associates, Inc., 2023, pp. 38154–38180. URL: https://proceedings.neurips.cc/paper_files/paper/2023/file/77c33e6a367922d003ff102ffb92b658-Paper-Conference.pdf.
- [26] T. Schick, J. Dwivedi-Yu, R. Dessi, R. Raileanu, M. Lomeli, L. Zettlemoyer, N. Cancedda, T. Scialom, Toolformer: Language models can teach themselves to use tools, arXiv preprint arXiv:2302.04761 (2023).
- [27] A. Balakrishnan, C. Schulze, Code obfuscation literature survey, *CS701 Construction of compilers* 19 (2005) 31.
- [28] T. Cai, X. Wang, T. Ma, X. Chen, D. Zhou, Large language models as tool makers, 2023. arXiv:2305.17126.
- [29] A. Fan, B. Gokkaya, M. Harman, M. Lyubarskiy, S. Sengupta, S. Yoo, J. M. Zhang, Large language models for software engineering: Survey and open problems, arXiv preprint arXiv:2310.03533 (2023).
- [30] D. G. Feitelson, From code complexity metrics to program comprehension, *Communications of the ACM* 66 (2023) 52–61.
- [31] Wikipedia contributors, International obfuscated c code contest — Wikipedia, the free encyclopedia, https://en.wikipedia.org/w/index.php?title=International_Obfuscated_C_Code_Contest&oldid=1152589523, 2023. [Online; accessed 10-August-2023].
- [32] A. J.-Y. Liew, Overcoming code rot in legacy software projects, Ph.D. thesis, Massachusetts Institute of Technology, 2017.
- [33] S. Van der Linden, Foolproof: Why misinformation infects our minds and how to build immunity, WW Norton & Company, 2023.