cBRAIN

AI for Government

How to Unlock the Potential of Context-Specific AI by Leveraging a Sandbox Approach

Recent advancements in the field of artificial intelligence (AI) present significant opportunities for government organizations to heighten efficiency, elevate quality standards, and provide better citizen services. However, these opportunities are accompanied by challenges such as compliance adherence, data security, and enabling AI models to understand the specific context of public sector organizations. Due to these challenges, government bodies are often restricted from effectively using publicly available AI technologies such as ChatGPT. To overcome these challenges, cBrain has developed a "sandbox approach" that solves the problems of compliance and security, while enabling government organizations to gain the benefits of context-specific AI models.

By Christoffer Tejs Knudsen & Frejdie Søndergård-Gudmandsen

September 2023

Generative AI offers governments a wide range of opportunities for increased productivity, elevated quality standards, and delivering better citizen services. As an example, the emergence of large language models (LLMs) has provided advanced artificial intelligence models that can be trained on large datasets to generate "human-like" textual outputs based on requests. By applying both generative AI models and machine learning algorithms, government organizations are offered a pathway for increased productivity, as case workers can streamline tasks such as making decisions, drafting letters to citizens, and preparing reports. Some government leaders estimate that such AI-powered case worker support can offer organizational productivity gains of up to 20%, while citizens can experience faster response time from their local governments. Higher quality standards can also be achieved as specifically trained AI models can help ensure consistency when responding to citizen inquiries and making decisions. Finally, as modern AI models can process extensive datasets, decision-making can be based on significantly larger amounts of data, and thereby citizens can receive more accurate and detailed information.

However, the main challenges for governments are safeguarding data and upholding regulatory requirements such as GDPR compliance. The risk of leaking confidential or personal information to external entities remains a problem, as the tasks of formulating responses or making decisions often are based on such sensitive data. This hinders the use of cloud-based natural language processing models, such as ChatGPT, where data-sharing is unavoidable. Thus, government organizations are often required to use on-premise solutions to ensure data sovereignty.

By leveraging a sandbox approach during training of AI models, government organizations can solve the challenges of safeguarding data and upholding regulatory requirements. In a sandbox approach, AI models are trained within a "sandbox", which refers to an isolated and controlled software environment running on-premise and separated from external environments. This ensures secure and isolated training of AI models, while avoiding the sharing or leakage of data. As a result, governments can train AI models on confidential and private data, with the assurance that data remains confined within the isolated environment.

The next challenge is training AI models to understand the specific context within government organizations. Each of these organizations has responsibilities to perform e.g., administrative, regulatory, legislative, or executive functions either domestically or internationally. These functions include specific tasks that have specific workflows, formats, language styles, and considerations that a general-purpose AI foundation model would not adequately understand. Thus, the use of general-purpose AI models, such as ChatGPT, will often result in outcomes that are too generic to be useful without further finetuning.

By applying a sandbox approach, governments can train AI models to understand the domain-specific context within government organizations, referred to as a "domain-specific model". By using domain-specific AI models as opposed to general-purpose models, government case workers can receive information that is both more detailed and relevant. The information will be more detailed as the domain-specific model will not only leverage publicly available information, but also have access to internal information within the government organization if needed for a specific task.

The information provided will also be more relevant as the domain-specific model will understand the context of the task at hand, and thus understand which information is more important to provide. Such models can effectively be re-used across similar types of organizations, such as different ministries, as the shared domain of these organizations means that they often work in very similar ways.

As an example of a use case, an AI model can understand how government ministries work, and how ministries meet to discuss political initiatives with other government bodies internationally. If a government official is going to the EU for political discussions, the domain-specific model can provide briefs about what the government official should know before entering the political discussions. Such information can include an overview of the political stances of other nations as well as statements made earlier within the official's own organization. The domain-specific model will know which information is most important because the model understands the context of the task at hand. Such a model can be re-used across different ministries within a government that participate in similar political discussions.

A sandbox approach can also be used to build AI models that have deeper insight into a specific organization, referred to as an "organization-specific model". An organization-specific model will have deeper insights into the workings and information of a specific government organization. As these capabilities will be very specific to one organization, the AI model will not be equally re-usable across other organizations, when compared to domain-specific models. Yet this trade-off might provide the deeper insights required for very detailed applications.

As an example of a use case, an organization-specific AI model can help predict delays in large scale planning of energy grids within authorities dealing with environmental protection. During large scale planning of energy grids, construction firms must receive approvals from their local government. Construction firms often risk not forecasting blockages in environmentally safeguarded zones such as protected water areas. Once the construction firms apply for approval from the local government safeguarding the areas, a construction project might already be long underway, and thus significant delays are incurred if plans must be changed. By leveraging AI-powered recommendations for decision-making, government case workers can foresee delays by leveraging large geo-data sets, and consequently advise construction firms on how to avoid potential delays. Such an AI model will be organization-specific, as the specific data and use case will be relevant to the specific organization only, and thus not be relevant to re-use across other organizations within the government.

When choosing to build either domain-specific or organization-specific models, government bodies should compare their needs to the advantages and disadvantages of each approach. However, the shared denominator of a sandbox approach is the ability to enable AI models to understand the specific context within government organizations, while simultaneously maintaining data sovereignty and regulatory compliance. This unlocks tremendous potential for increased productivity, higher quality standards, and better citizen services.

For more information, feel free to contact the authors at cBrain.

Christoffer Tejs Knudsen
ckn@cbrain.com

Frejdie Søndergård-Gudmandsen
fsg@cbrain.com