





Risk Identification in Public Contracts:

the role of Generative AI and
the Company Analysis Tool

January 2025





Credits

[Yash Yeola](#) is a consultant at Butterfly Data with a postgraduate degree in Scientific Computing and Data Analysis. He is recognised for his expertise in data modelling and machine learning, specifically Large Language Models, and developing AI applications. His insights helped to shape the logic behind the design and build of the Company Analysis Tool. Yash is the developer of the CAT, connecting the data sources and creating the LLM summaries.

[André Petheram](#) is research lead and principal consultant at Oxford Insights. As an anti-corruption specialist, he has overseen evaluation and research projects for the Open Contracting Partnership and Open Ownership. In 2024, André spoke at the Next Generation Procurement conference in Rome and delivered tailor-made training on using open data and AI in anti-corruption for the International Anti-Corruption Academy. He advised on this project and worked with Yash to come up with the concept of the Company Analysis Tool.

[Maja Stawinska](#) is a consultant at Butterfly Data, with a postgraduate degree in AI and Data Science and foundation in Criminology and Forensic Science. Passionate about AI's role in environmental and space applications, Maja explores machine learning for satellite image analysis and ecological modeling to drive sustainable innovation. Committed to ethical AI development, she advocates for responsible and transparent use of artificial intelligence in her projects to ensure equitable and sustainable outcomes. Maja conducted the baseline research into global developments of anti-corruption solutions and co-authored the white paper.

[Jasmine Kendall](#) is a principal consultant at Oxford Insights with a keen interest in open government, especially in procurement. She has led evaluations for several transparency focused NGOs, including the Open Contracting Partnership, Open Ownership, and the Extractive Industries Transparency Initiative. Jasmine was able to bring her experience to the fore in co-authoring the white paper.



Contents

1 Executive Summary

2 Introduction

3 The Company Analysis Tool

Overview

What data does the model use?

How does the tool flag companies for further investigation?

Generative AI

Limitations

4 Potential Application of the CAT

5 Next Steps

6 About Us

Executive Summary

The COVID-19 pandemic exposed major weaknesses in the UK's public procurement system, leading to billions being lost through fraud, inefficiency, and poorly executed contracts. In response, Rachel Reeves announced the appointment of [a COVID corruption tsar](#) and launched [investigations into £600m worth of COVID contracts](#).

At **Oxford Insights and Butterfly Data**, we recognise that artificial intelligence (AI) and machine learning hold significant potential to both identify and recover these losses. Indeed, seven EU countries are already employing AI to analyse procurement data or assist with tender processes, with 10 more planning to adopt similar tools. However, more needs to be done to create practical, actionable use cases in this space.

That's why we took action and partnered to develop the **Company Analysis Tool (CAT)** - an idea conceived and realised through Oxford Insights' and Butterfly Data's expertise. By combining UK contracting data from **Companies House and Contracts Finder** with proven 'red flag' methodologies, the CAT summarises any potential risks associated with contracting a given company. Enhanced with large-language-model (LLM) technology, it also synthesises contextual insights from news articles and Wikipedia. Importantly, the CAT has been designed as an ethical tool which is transparent about how data is sourced and respects copyright law.

In this white paper, we explain how the CAT works and outline its potential to help governments identify companies for manual investigation. Looking ahead, we are committed to further developing the CAT by integrating a broader range of data sources, enhancing its architecture to handle larger and more complex datasets, and exploring machine learning to detect suspicious patterns.

By working closely with government and civil society stakeholders, we aim to ensure the CAT becomes a practical solution to address real-world challenges in public procurement.

Ongoing feedback and testing will be essential to its success - **please get in touch if you would like to contribute!**

Introduction

During the COVID-19 pandemic, governments globally faced immense pressure to procure essential goods and services to save lives. In the UK, this urgency exposed critical flaws in the public procurement system, leading to billions of pounds lost through fraud, inefficiency, and poorly executed contracts. [A House of Commons report from 2023](#) found that 'since the beginning of the COVID-19 pandemic, the level of fraud against the taxpayer has increased fourfold' in the UK.

At the same time, governments worldwide are recognising AI's potential for [combatting corruption](#) and inefficiency within public procurement. Internationally, tools such as Brazil's [ALICE](#) and the EU's [DATACROS II](#) leverage data and AI to identify cases of potential corruption cases and other financial crimes in the public sector. Seven EU countries are already using AI to analyse procurement data or assist with tender documentation, with many others planning to follow suit. In the UK, however, the use of AI in public sector procurement to identify red flags remains largely underdeveloped.

AI-driven accountability in the UK

Following analysis of public procurement data in the UK, [Transparency International](#) found corruption red flags in £15.3bn worth of COVID-19 contracts across 135 high-risk contracts. In an effort to begin to recoup these losses, the Labour government has announced the appointment of a [COVID Corruption Commissioner](#) Tom Hayhoe. The Commissioner will be responsible for examining potential fraud and flawed contracts, working closely with agencies such as HMRC, the Serious Fraud Office, and the National Crime Agency.


For some time, we have recognised the significant potential of artificial intelligence to analyse the vast amounts of data within government contracts and efficiently identify risks that warrant further investigation. Now, as the Commissioner begins their work, it feels like the right moment to move beyond theoretical discussions and start building a practical tool to show how AI might support these efforts.

This white paper introduces the logic behind the new **Company Analysis Tool**, outlining its purpose, limitations, and potential use cases. We also share our vision for the tool's development and plans for further testing, as we aim to refine it in line with user feedback.

The Company Analysis Tool

Overview

The Company Analysis Tool (CAT) enables UK Government users to identify corruption risks in public contracts as part of a more rigorous and effective procurement process. The tool is designed to identify red flags across awarded public contracts, and add any further red flags to a company profile according to historic company registration details. It aggregates data from multiple sources, including Companies House, Contracts Finder, Wikipedia and news outlets.



The tool then ranks the companies based on the total number of flags and their weightings. The data for the most suspicious companies is processed and summarised by a Large Language Model (LLM), an advanced AI system trained to understand and generate human-like text. The LLM identifies and highlights key insights from the different sources, making complex information more digestible. This enables users to assess whether a contract might have been awarded without proper due diligence – such as to a company with a conflict of interest or one inadequately equipped to meet the government’s requirements.

The tool has a modular design, allowing a lot of flexibility with the internal logic. For instance, the company selection logic can be updated, more relevant red flags specified, or a different LLM integrated. This allows the CAT to be easily tailored to various users, according to organisational priorities and insights.

The Company Analysis Tool

What data does the model use?

The model draws data from 4 sources:

Contracts Finder

Contracts Finder is a publicly available tool to search for information about contracts worth over £12,000 (including VAT) with the government and its agencies. It is commonly used to find past, present and future contract details across different sectors. Contracts Finder provides basic contract filtering capabilities such as on publication dates, deadlines, sector, and whether or not contracts were open to SMEs. It provides contract information such as who the contract was awarded to and when it was awarded.

Companies House

Companies House is the UK's official register of companies. It is responsible for maintaining the register for England, Wales, Scotland, and Northern Ireland. All companies must be registered on Companies House by law and make their details publicly accessible. Some examples of this information include:

- Company Name
- Company status (eg. active)
- Company creation date
- If the company has been liquidated
- Historic charges
- Insolvency history

The advantage of using Companies House is that it provides a variety of detailed company information on one platform*.

Wikipedia API

To provide extra context on a company's background, information is collected from the relevant Wikipedia page, if it exists.

Wikipedia is a publicly available source of information, created, edited, and verified by volunteers around the world. Content on Wikipedia must pass a notability test, whereby editors review verifiable information to decide whether a topic warrants a Wikipedia page. As such, in the context of the CAT, a company will only have an entry if multiple third parties have written about them, for example.

*Some challenges of this are discussed in the [limitations section](#) of this white paper

The Company Analysis Tool

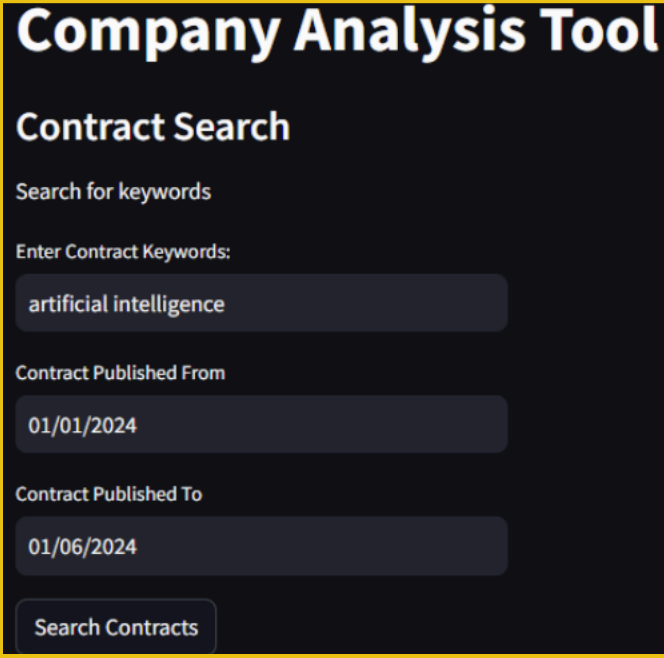
Therefore, smaller companies tend to not have Wikipedia pages. Wikipedia text is licensed under Creative Commons or GNU so it can be freely used without concerns over copyright.

News

News articles provide a treasure trove of recent information that might be absent from other sources such as Wikipedia. The CAT uses the [News API](#) to grab headlines and snippets from articles related to the query "<company name> corruption". The tool retrieves the top three articles for each company search. Furthermore, there is additional functionality to retrieve relevant articles from the internet from a diverse range of sources, providing more detailed information on companies. Copyright laws are respected during this process.

How does the tool flag companies for further investigation?

The user starts by providing keywords and a publication time period to filter for contracts. The keyword input is very versatile since it searches both tender names and descriptions, allowing the user to search for specific sectors or services, such as artificial intelligence in the example. Pressing the search button prompts the tool to process the contract information and flag companies based on the flags discussed on the following page.



The screenshot shows the 'Company Analysis Tool' interface with a 'Contract Search' section. It features three input fields: 'Search for keywords' (containing 'artificial intelligence'), 'Contract Published From' (containing '01/01/2024'), and 'Contract Published To' (containing '01/06/2024'). A 'Search Contracts' button is located at the bottom of the form.

Company Analysis Tool

Contract Search

Search for keywords

Enter Contract Keywords:

artificial intelligence

Contract Published From

01/01/2024

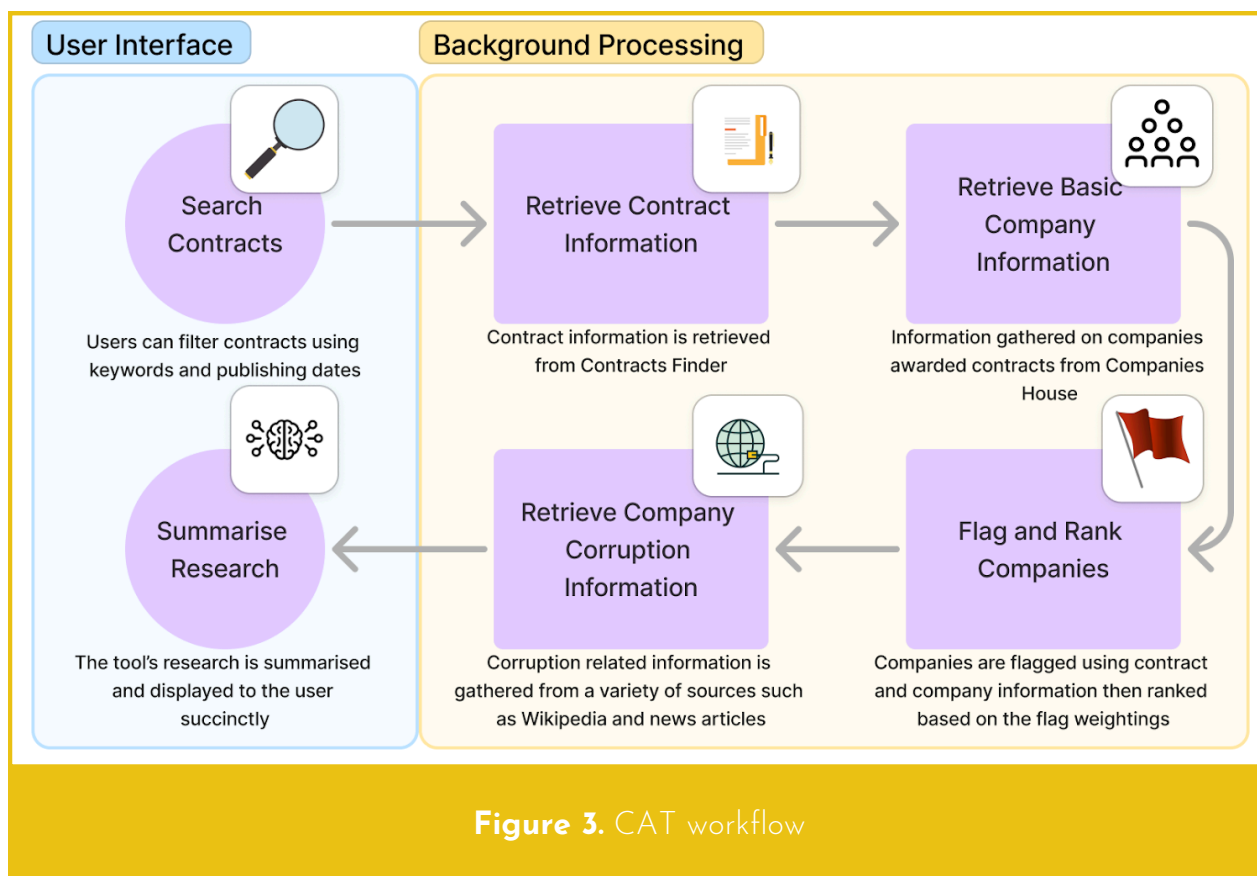
Contract Published To

01/06/2024

Search Contracts

Figure 2. Contract search interface

The Company Analysis Tool



The process by which the CAT returns flagged companies from this initial search is depicted in **Figure 3**. Relevant contracts which meet the search criteria are retrieved from Contracts Finder, and key details such as award dates and company names are extracted. These details are then processed for the purpose of identifying red flags, and company results ranked accordingly. In the current iteration of the CAT, companies are individually flagged for further investigation based on specific indicators. The companies are then ordered and the CAT displays only the top 5 companies which trigger the most red flags.

There are 2 categories of red flags which identify corruption risks; the first category of red flags identify risks within associated procurement processes, and the second, within historic data captured against the company themselves. Many different organisations have researched the use of red flags in the procurement space. For example, up to 2023 the [Red Flags project](#) supported by the European Commission aimed to provide a constant feed of contracts flagged by their set of red flags, [Transparency International explored COVID-19 contracts](#) in particular and the Open Contracting Partnership created a [guide](#) for common red flags. The red flags used in this project have been inspired by those detailed in these studies.

The Company Analysis Tool

Procurement process risk flags

1. **Contract award published late:** The contracting authority published the award notice more than 30 days after the contract was signed, the awarded supplier is flagged.
2. **High priority contract:** the supplier was awarded the contract through a high priority process such as an open accelerated procedure.
3. **Limited competition:** The supplier was awarded the contract through a procurement procedure with limited competition. This includes cases such as restricted procurement where suppliers are invited to tender or need to meet specific criteria, rather than the process being open to all.
4. **No publication:** The supplier was awarded a contract which was not published before award or a direct award.

Company profile risk flags

1. **Young company:** The awarded supplier is less than 101 days old when winning the contract. The likelihood of young companies winning contracts is low especially if they win multiple contracts. Therefore, when combined with other flags, a young company can increase the risk rating.
2. **Trustworthiness score:** The following factors from the Companies House API are equally weighted to calculate a score*:
 - Overdue next accounts
 - Overdue confirmation statement
 - Liquidation status
 - Insolvency history
 - Disputed registered office
3. **Contextual information from the media:** The LLM element of the tool then uses the News API and Wikipedia API to summarise relevant news articles and Wikipedia pages about the company, identifying potential misconduct or controversies to complement the more quantitative indicators. It is then left to the user to make a more nuanced judgement about whether or not the contract merits further investigation.

*The more of these issues a company has, the higher its risk score. Weighting adjustments could be made to reflect the importance of each issue.

The Company Analysis Tool

The search results are displayed in a table indicating the number of times each company has earned a certain flag. Only 20 contracts were processed during this demo due to time and processing limitations though this number can be increased for a more comprehensive analysis. The companies are ranked based on the total number of flags and their weightings. For example the no publication flag has a higher weighting than the 'award published late' flag as it is a clearer indicator of suspicious activity. The company names in **Figure 4.** have been redacted for demo purposes. The user then has the ability to drill down on individual companies and get further information on their history.

Search Results ⇄

First 20 contracts published from 01/01/2024 to 01/06/2024 for the keyword artificial intelligence.

Company Name	Award Published Late	Limited Competition	Tender Not Published	Accelerated Procedure	Company Too Young	Trustworthiness Score
C	1	0	1	0	0	1
D	1	0	0	0	1	1
P	0	0	1	0	1	1
C	0	0	0	1	0	1
C	0	0	0	1	0	1

Figure 4. Red flags summary generated from the search query

The Company Analysis Tool

Generative AI

Generative AI, in this case a Large Language Model (LLM), is used to extract relevant information and summarise content. LLMs are Artificial Intelligence (AI) systems that can understand, manipulate and generate human language. They are proficient at a wide range of Natural Language Processing (NLP) tasks, pushing the boundaries of what machines can achieve. The use of LLMs in this scenario makes it easier for the user to understand the company background by creating quick-to-read summaries of content available online.

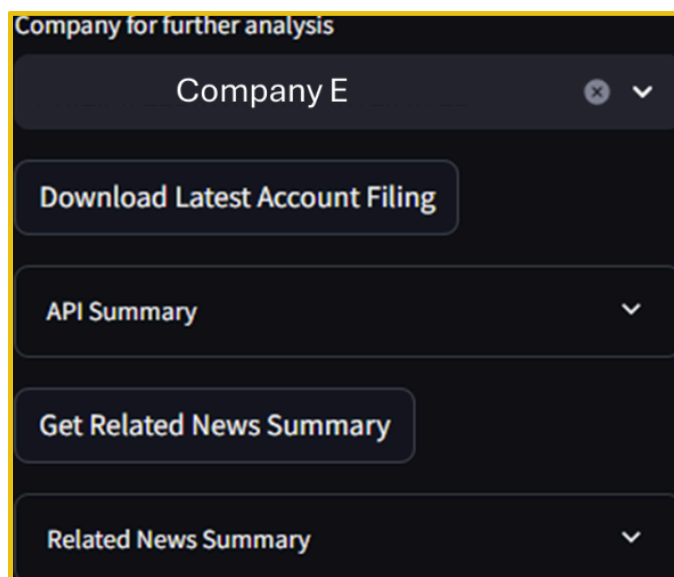


Figure 5. Generative AI content provided with company drill down feature

In addition, the user also has the option to download the latest account filing of the company. This is useful for analysing the financial state of a company.

The LLM is used to create two different summaries. One is using the information from the different APIs (Wikipedia and the News APIs) and the other uses internet news articles. As seen in the image below, the summaries use different sources providing a diverse range of information to help the user analyse a company. When creating these summaries, the LLM prompts are carefully engineered to ensure that it only uses the data provided, reducing the chance of hallucinations. This increases the reliability of the results and creates a succinct summary. An example of this can be seen in **Figure 6**. It contains real company information with modified names and dates. This LLM content should then be assessed by the user who can further investigate the results by visiting the original sources.

The Company Analysis Tool

API Summary	Related News Summary
<p>The provided text mentions two instances potentially related to corruption or unethical business practices:</p> <ol style="list-style-type: none">The Corrupt Cartel (1885): Company E joined with Company F to form a cartel that artificially limited the lifespan of plastic chairs to increase sales and profits. This is described as a form of planned obsolescence, an unethical business practice, though not necessarily illegal in that era. Source: The provided text under the "Environmental" section.Shopping Bag Recall (2024): Reports indicate Company E received thousands of reports of issues with Plastic Bags as far back as 2003, but did not disclose them to the FDA as required. This led to a large recall, lawsuits, and a significant settlement. Source: The provided text under the "Shopping Bag Recall" section. <p>Sources: https://en.wikipedia.org/wiki/CompanyE</p>	<p>In 2012-2019, Company E employees, distributors, or sub-dealers engaged in improper bidding practices in Country ABC, including providing funds to Shopping directors and tailoring bag specifications to favor Company E. This resulted in a settlement with the ABCDE in 2025 where Company E paid \$233 million in disgorgement and prejudgment interest. The ABCDE alleged violations of the books and records and internal accounting controls provisions of the Bags Exchange Act of 1888. Company E consented to the order without admitting or denying the findings.</p> <p>Sources</p> <p>https://www.bagtechdive.com/news/COMPANY-E-ABCDE</p> <p>https://www.bags.org.uk/news/news-stories/update-company-e-trust-corporations</p>

Figure 6. LLM summaries for news articles and API content

The information available on Companies House also provides background information on a company, including the components of the trustworthiness score and other important, commonly-used information, such as the previous and current officers of the company as seen in **Figure 7**.

All these pieces of information combine to form a well-rounded picture of the company, automating a significant part of the laborious research process. The whole interface can be seen in **Figure 8**.

Company Information	
company_name	Company E Ltd
company_number	
date_of_creation	19**-**-**
current_company_officer_details	
previous_company_officer_details	
next_account_overdue	False

Figure 7. Company background information

The Company Analysis Tool

Company Analysis Tool

Contract Search

Search for keywords

Enter Contract Keywords:

artificial intelligence

Contract Published From

01/01/2024

Contract Published To

01/06/2024

Search Contracts

Search Results

First 20 contracts published from 01/01/2024 to 01/06/2024 for the keyword artificial intelligence.

Company Name	te	Limited Competition	Tender Not Published
C	1	0	1
D	1	0	0
A	0	0	0
P	0	0	1

Company for further analysis

Company E

Download Latest Account Filing

API Summary

Get Related News Summary

Company Information

company_name

Company E Ltd

company_number

date_of_creation

19***.***

current_company_officer_details

previous_company_officer_details

next_account_overdue

False

company_status

Figure 8. Complete CAT interface

Limitations

Whilst the tool demonstrates significant potential to make it quicker and easier to identify red flags within contracts, there remain limitations in the approach so far. In particular, the absence of data on the number of bidders per contract, rollovers and overspends reduce the number of risk indicators that can be taken into account during the flagging process.

A number of challenges relate specifically to the capabilities of the Companies House API and its search functionality. For example, a common discrepancy is that certain companies are not returned in the search as their name listed on Companies House is not the same as the name used in contract awards. In cases like this, it is difficult to identify the red flags and provide a comprehensive company summary. A lot of time has been spent improving the hit rate of the search tool by using different string processing methods. However, in cases where even this fails, the companies are not included in the results table. Companies House is also not currently required to verify the accuracy of the filed information, which may leave room for some inaccuracies. It is widely reported that there are significant data quality problems with Companies House. That said, there are improvements, most notably [an identity verification process](#), being implemented in February 2025 which aims to improve reliability of the data. In the future, a possible work around for these limitations might be to create a file with common errors and inaccuracies or retrieve company numbers through external sources.

Potential Application of the CAT

The modular design of the Company Analysis Tool means it can be customised for different use cases across government, where departments would benefit from an efficient way of identifying the risks associated with companies. The following examples are just some of the potential use cases to which the CAT could be applied:

- The CAT could support the work of the **Covid Counter-Fraud Commissioner** in addressing the £8.7 billion lost to poor contracting practices during the pandemic. The CAT helps to identify instances where due diligence may have been overlooked - such as awarding contracts to companies with conflicts of interest or questionable credentials. Although the tool has limitations and cannot replace thorough human-led investigations, the CAT provides a valuable first step by flagging contracts for further review, saving time and resources.
- The **Department for Work and Pensions (DWP)** could use a tailored version of the CAT to identify fraudulent companies that have benefited from welfare payments. The tool could help DWP efficiently identify companies that are at high risk of being used for fraudulent purposes. For example, a company established just before the pandemic, with a disputed registered office and no online presence, would be identified by the CAT for being a likely recipient of fraudulent welfare benefits.
- The **Public Sector Fraud Agency** could use the tool in combination with their profiling and networks tool for the purposes of identifying new patterns in corporate fraud as it emerges. Clusters of companies are being set up using the same agents, addresses or even set up on the same day with the same directors. The CAT would be able to assist in more timely detection of fraudulent activity, and so enable future fraud investigations to be more robust and informed.
- The CAT has the potential to be applied at **His Majesty's Revenue and Customs (HMRC)** to help recuperate fraudulent grants bestowed on organisations as part of the various Covid support schemes that were implemented. The tool could be tailored to rank and prioritise companies which are due to be investigated as part of HMRC's post payment compliance investigations. For example, companies who had claimed furlough or restaurants that had claimed through the 'Eat out to Help out' scheme.

Next Steps

Looking ahead, we are committed to expanding the technical capabilities of the Company Analysis Tool by broadening the range of data sources integrated into the platform. This will include both structured data, such as financial metrics and operational KPIs, as well as unstructured data, market reports, and social media sentiment.

To support this enhanced integration and meet growing demand, we plan to upgrade the platform's architecture to process larger datasets and handle more complex analytical queries. Additionally, we are exploring the potential integration of machine learning (ML) to uncover patterns which might point to suspicious activity within the data.


At the core of our approach is ensuring that the CAT evolves into a solution that truly meets the diverse needs of its users. Collaborating with stakeholders across government and civil society will be essential to this process. Through ongoing testing and iteration, we aim to continuously develop features that address the real-world challenges faced by government procurement officials and civil society organisations monitoring public contracts.

If you have feedback on the tool or if you are interested in participating in testing its utility, please do get in touch.




About Butterfly Data

A consultancy specialising in data-driven solutions, leveraging advanced analytics and AI to deliver innovative strategies for businesses and organisations. Committed to sustainability and technological advancement, Butterfly Data empowers clients to make informed, impactful decisions backed up by data.

 hello@butterflydata.co.uk

 butterflydata.co.uk

 029 212 00140



About Oxford Insights

A consultancy that combines new thinking on technology and leadership with experience getting things done in government. With years of experience creating public sector transformation, both as implementors and advisers, Oxford Insights helps governments use new technologies and innovative approaches to make citizens' lives better.

 info@oxfordinsights.com

 oxfordinsights.com

