

Global AI Risk Observatory

Background

This project builds on work done in the beginning of this year focused on analysing 9,821 annual reports from UK public companies, hosted at riskobservatory.ai and funded by the [UK AISI](#) societal resilience [team](#). That initial phase established a working methodology ([which can be found on the page](#)) but remained largely analytical - it didn't yet draw policy conclusions of findings - and remained limited to one data source (annual reports) and region (UK).

See [this paper](#) for more details on the work done thus far (request access needed)

Proposal

Through a partnership with our data provider, [FinancialReports](#), we now have access to all published data from public companies (annual reports, investor presentations, and earnings-call transcripts) in a machine-readable Markdown format that allows us to use LLM autograders to perform rigorous analysis. We want to expand the work to a much deeper and more nuanced taxonomy, and expand coverage to global.

The central goal is to scale the existing UK analysis into global coverage (beginning with the US) and to translate observed disclosure trends into policy-relevant findings. The agenda is deliberately exploratory: we have guiding questions, but we expect much of the value to emerge from the data itself. Indicative questions include:

- How do the AI-risk disclosure trends identified in the UK transfer to other jurisdictions, starting with the US?
- Can comparable signals be found in other classes of legal and regulatory filings beyond annual reports?
- Do disclosure patterns correlate with external data sources, for example, the [AI incidents database](#) or media coverage?
- What do differences in reporting (e.g. systematically less disclosure of certain risk categories i.e. in particular sector or country) imply for regulation and governance?

Approach

Process the 1,000,000 documents from public companies from around the world with LLM auto-graders, draw statistical analysis, define taxonomy of interest and cross-referencing disclosure patterns against external datasets to surface and validate trends.

Collaborations

The project is expected to be carried out with advice of / collaboration with the Oxford Martin School and the MIT Risk Initiative.

Fellowship roles

We are recruiting two fellows with complementary focuses.

1. Technical fellow — LLM data-analysis pipeline

- Experience building LLM applications
- Foundational computer-science knowledge
- Desirable: prior work on a similar application, or an interest in statistical analysis and LLM auto-graders
- A finance background or familiarity with financial markets and data is helpful but not required

2. Governance & policy fellow

- Interest in global AI regulation and cross-jurisdictional trends
- Keen to do empirically grounded, data-driven governance work
- Comfortable translating data patterns (such as regional gaps in risk reporting) into concrete policy implications