**Document Details:** Q&A in response to the on-line Briefing Call
**Challenge:** Use of AI in screening individuals of interest
**Event Date:** Tuesday 9th April 2024
**Location:** Teams event hosted by Plexal

**Current Methods**
Q: How do you currently validate that decision points are correct? Is this done by a human?
*A: Yes*

Q: Do you already store records of historic searches/decision points in the current system or should such a logging system be part of the solution?
*A: For audit trails, logging would be preferred.*

**Integration**
Q: What sort of interfaces, if any, would you be looking for from this work? APIs? User interfaces? Integration with other analytical tools?
Q: Are there any existing systems or tools that the screening solution would need to integrate or interface with?
*A: Not at this stage. We are looking for a proof of concept so integration is not a current consideration.*

Q: Will the deployed solution be entirely offline, without any connection to the internet?
Q: You mentioned on-prem deployment without public cloud - are there any other constraints or preferences around the deployment architecture? Any existing infrastructure you'd like to leverage?
Q: What sort of hardware do you expect to run the solution on? For this work, is it sufficient to emulate it?
*A: The deployed solution will be air-gapped.*

**Source Data**
Q: The range of datasets and formats involved in screening is potentially vast, but the budget for this is relatively modest. Where should we focus our efforts?
*A: The service typically handles a wide range of data reflective, where legislation, viewed through the lens of our strong ethics driven stance, allows.*

*Working on close partnership with the Security Service. https://www.mi5.gov.uk/how-we-work/gathering-intelligence the service, GCHQ and our wider partners globally the service is able to obtain data types which could include:*

***Bulk Personal data**:*
*https://assets.publishing.service.gov.uk/media/5a7f856bed915d74e33f6f2c/BPD_Factsheet.pdf*

***Data obtained through surveillance:***
*https://assets.publishing.service.gov.uk/media/5ba37401e5274a55cdb89bce/201800802_CSPI_code.pdf*

***Communications data – this could be emails or telephony metadata and content:***
*https://www.gov.uk/government/collections/communications-data*

***Equipment interference. For further details of EI please see:***
*https://assets.publishing.service.gov.uk/media/64145522d3bf7f79d6487bd4/Equipment_Interference_Code_of_Practice.pdf*

*For a wider overview of the IPCO legislation under which the service operates when exploiting these data types please see below:*
*https://www.ipco.org.uk/investigatory-powers/the-powers/*

Q: How does the source data vary? Is it unstructured data from the same sources each time, or completely different depending on which authority the data comes from etc.?
*A: As reflected in the above there is a varied range of both structured and unstructured data utilised. A number of the repeatable, regular data sets are in a standardised format for exploitation but there are obvious examples from the above where it is difficult to predict the mix of structured and unstructured data the service may obtain e.g EI.*

Q: Can you point us towards some representative datasets?
*A: Further reading around the core data types above should provide insight into the types of data the service is seeking to exploit. Representative data sets are not available.*

Q: The range of datasets and formats involved in screening is potentially vast, but the budget for this is relatively modest. Where should we focus our efforts?
Q: Can you point us towards some representative datasets?
Q: Are there any filetype requirements or nice-to-haves for outputs of the proposed solution?
Q: Is there any common data lake created already that will be consumed by the proposed application or we will need to create connectors for individual sources?
Q: Are any of the unstructured data sources encrypted?
Q: Can you confirm whether or not you are interested in solutions that can process multi-modal data?
Q: Do you have an idea the volume of the unstructured data
Q: Can you provide more details on the various data sets involved in the current screening process (e.g. structured vs unstructured, data types, volumes)? Are there any restrictions or sensitivities around using certain data sets?
Q: Existing frameworks for screening (e.g. broad criteria and data sources used - as much as you can say please). And are any quantitative methods currently used or is it all subjective / text?
*A: There is a significant interest in how we can augment the role the skilled analysts in the service using carefully risk assessed criteria to screen both the data against known individuals of interest and the predication of potential individuals of interest.*

Q: Data processing tools are mentioned. Will this solution be using the output of the data processing tool, or replacing it?
*A: It is part of the current process so will be using the output of the processing tool, not replacing any processing tools*

**Outputs**

Q: What does the ideal output look like? e.g. threat/non-threat, classified entities, graphical display etc?
*A: The output would be consumed by a range of higher skilled analysts and a broader spread of officers with less analytical skills or knowledge across a range of thematic and geographical teams. Any implied approximation of threat or risk would need to be explainable to those varied users (see human in the loop comment).*

Q: Would shifting from reactive screening to proactive monitoring and alerting based on risk triggers and patterns be a significant leap forward? Is that a future vision you have?
*A: We would not allude to future development at this point.*

Q: While automation is important, how crucial is it to keep the human-in-the-loop and empower analysts with augmented intelligence capabilities?
*A: A person would always have to remain in the loop, as you note empowering the analyst should be the focus.*

Q: Are you looking only for new AI tools, or are you interested in a piece of non-AI software which presents and aggregates data together for future analysis/human-machine teaming?
*A: Yes, non-AI software could be an option*

Q: Is collaborative investigation a key part of the workflow? Would capabilities like shared workspaces, annotations, and peer reviews be game-changing for your teams?
*A: We cannot allude to capabilities at this point.*

Q: Beyond the core screening process, are there adjacent operational pain points like resource allocation, workload balancing, and SLA tracking that you wish to address?
*A: We cannot allude to pain points at this time.*

Q: Is capturing and sharing the collective knowledge and expertise of your analyst community a key challenge?
*A: That isn't the focus of this challenge.*

Q: Is the customer looking for a report generation tool that can automate many of the generation tasks and link to various data sources?
*A: That could be one option for a solution we'd be interested in.*

Q: Is the customer looking for something that can synthesise data from the various feeds from the software they currently use and present these to an analyst?
*A: No, that would miss the scope of the challenge.*

Q: What are the key metrics and success criteria you hope to achieve with an AI-assisted solution (e.g. reducing manual effort by X%, improving decision turnaround by Y days, uncovering Z% more insights)?
*A: We don't have a key metric in mind as we are looking for a concept and art of the possible in this initial project.*

Q: Can you walk us through the current end-to-end workflow of a typical screening case, highlighting the key pain points and manual efforts involved? This will help us map the As-Is process.
*A: Please refer to the challenge form for the typical process, we unfortunately can't give any more detail at this time.*

**Continuous Learning**
Q: Given that an AI system -- adapting to new information as it becomes available requires backtracking -- ie: a system that's not necessarily able to terminate (ie: complete its computation within a fixed amount of time) -- is there any guidance here? Since otherwise the only way to approach the project is perhaps to cheat and solve a set of simpler problems, but it may point towards the problem not being something that can be solved.
Q: How important is continuous learning and adaptation of the models as new data comes in and analysts provide feedback? Is a static one-time model sufficient or do you need dynamic learning?
*A: ideally yes, we need something that keeps the data up to date.*

Q: Given the sensitive nature of screening, is ensuring the AI/ML models are fair, unbiased and auditable a critical requirement? Would you need the ability to test and correct for potential biases?
*A: Yes, this is definitely a requirement. A person in the loop will always be required but having a "white box" system where they understand the model is very important.*

**Searching**
Q: Are search parameters always flexible on a case-by-case basis or are many fixed with only some flexible parameters?
*A: There are a series of standard analytical queries ran across data but there needs to be flexibility to allow parameters to change to reflect specific requirements driven by the use case.*

Q: Would the ability to run what-if simulations and scenarios based on different screening criteria or thresholds be valuable for your decision makers?
*A: No, this would be out of scope.*

Q: Beyond basic search and retrieval, how valuable would it be to automatically surface novel connections, correlations or anomalies across disparate data sets that analysts may have missed?
*A: Yes – refer to previous answer on human/machine teaming and predication.*

**Clarifying Questions**
Q: Is the Screening for recruitment purposes or targets of investigations, both or something else entirely?
*A: They are both the same.*

Q: Can you clarify what the 'live document' is please. I'm assuming its the 'screening document' but would like clarification?

*A: The live document is the Screening Document.*

Q: In the example Sam has 100 individuals, is this 100 separate cases (with separate reports) or one batch?
*A: Yes, 100 separate cases*

**Legal/Compliance**
Q: Enron dataset (just as an example) it is classified under the IPA as a BPD, however industry are free to use this as we are not covered by the IPA. So was not sure if we would be allowed to use this work? Or would we be restricted to experiment with non-BPD data / synthetic data?
Q: Are there any regulations that apply to this work? As in most AI work is now, or will in the future be controlled. Is this project excluded under R&D? We would need confirmation on this please.
Q: What are your requirements and constraints around data governance, security and auditing? Are there any specific compliance or regulatory needs to be met?
Q: In the event of audits or litigation, how critical is an explainable audit trail of all decisions and supporting evidence? Is a transparent decision fabric a key differentiator?
Q: How important is the explainability of the AI/ML models and recommendations? What level of transparency and audibility is required for analysts to trust the outputs?
*A: We are compliant with all legal, regulatory and auditability requirements.*

**Contract**
Q: What do you mean by a consortium? Are there any constraints associated with the building of consortiums?
*A: We view a consortium as a group of organisations working together on a common goal. However, we only commercially engage with the lead organisation, you will need to commercially agree with "subcontractors" within your consortium*

Q: There appears to be no limitation of liability for either party in the T&Cs. Please confirm that any contract would be subject to a cap on liability proportionate to contract value?
*A: Please refer to clause 14.2 of the Co-Creation Challenge terms where the limit of liability (or liability cap) is outlined'.*

Q: Please confirm that the requirement to provide copies of insurance policies may be satisfied by our providing brokers' letters, due to the terms of our insurance being confidential and held at group-company level.
*A: Yes, this will be sufficient,*

**Submissions**
Q: How do you want the contributors to submit the proposals? Is there a template that you can share?
*A: If you submit a proposal via the KTN, there are sections to complete, but if you apply via a different route (such as directly to cocreation?hmgcc.gov.uk), there is no prescribed template, but on the challenge form there is guidance under "How to Apply". This includes a limit of six pages and key information to include in the proposal.*

*Please do also note that Co-Creation have published their terms and conditions and these are non-negotiable.*

Q: How do we submit the proposals?
*A: Send to cocreation@hmgcc.gov.uk or via one of our collaborators. Please do note by which collaborator you initially found this challenge.*

Q: How many applications can come from a single institution?
*A: No limit but be considerate towards reputation of your institution.*

Q: What is the success criteria? how many entities are we looking at?
*A: In the published challenge form there are the criteria by which each proposal will be assessed. We cannot say at this stage how many applicants will be successful as there is no maximum determined budget, but historical data shows Co-Creation funding on average 2 applicants per challenge.*

Q: As a university, we're going to be lead applicant on the proposal and we're going to follow the university systems to do the application, which probably could be very different. We have our own constraints in terms of what we can and can't do. There is no limitation as to how many applications can come from a single university, right?
*A: No limitation. I'd caution against multiple academics from any one university applying without considering across the school or across the technology base. It would look odd if we get 5 applications that aren't joined up hence we do steer towards consortium build if that's the approach. We'd recommend to be considerate of the university's reputation because we're looking at the university rather than the academic.*

**Project Delivery**
Q: Will there be the opportunity to perform user research for this work?
*A: The Users will be integral in project Sprints*

Q: Given the 12-week horizon, would you be open to a phased approach that starts with a thin slice PoC or MVP targeting a specific sub-process or data set, and then iteratively build on it? If so, which area would you prioritize first?
*A: This would be a good approach…*

Q: You mention mid-TRL products. What range of TRLs are you looking for, and are you expecting technology to be mid-TRL going into this or mid-TRL coming out of this work?
*A: Please see UKRI TRL as a reference point. We would ideally want a technology entry point of TRL 3-5, and an exit point of TRL 5 – 6, technology validation in a relevant environment.*

Q: What balance are you looking for between "horizon scanning" and building/demoing a product?
*A: We aren't focused on horizon scanning, we would like you to work with the users to develop existing developments to a relevant environment for our use case.*

**Eligibility**
Q: Is it possible that we could collaborate with international partners?

*A: This challenge is open to sole innovators, industry, academic and research organisations of all types and sizes. There is no requirement for security clearances.*
*Solution providers or direct collaboration from countries listed by the UK government under trade sanctions and/or arms embargoes, are not eligible for HMGCC Co-Creation challenges.*

**General**
Q: Will the pitch day be physical or online presence?
*A: We will accommodate online pitches, but our strong preference is in-person attendance at the pitch day location in London.*