

Narrative Detection Tracking: Call for Proposals

Closing date: 30 September 2022

Contents

Summary	2
Available Funding.....	2
Terms and conditions	2
Background.....	2
The Requirement	3
Eligibility.....	6
How to apply	7
Deliverables	8
Assessment and review	9
Key Dates.....	9
Post-award information	10
Queries	11

Summary

The Alan Turing Institute's Defence and Security programme is inviting proposals to investigate the application of AI techniques to suitable data sources to identify the emergence and track the spread of adversary narratives. This task will support a human analyst to identify narratives in a timely fashion and classify them as mis/disinformation when appropriate.

Once identified, an analyst should be able to track their spread – for example, by finding related material that either supports or refutes a specific narrative. We aim to develop both a repeatable method to apply to specific events, as well as considering longer-term, strategic narratives.

Available Funding

Funding will be available for one project over the initial duration of 8-months with an option to extend to a further 6-months. The proposal should be defined such that the initial project can be completed within 8-months. The proposal should also include an option for a 6-month follow on. The research will be funded at Full Economic Cost and VAT will apply.

Depending on the outcomes following the initial 8-month project and 6-month option phase our intention is to continue to support this research over several years.

Eligible costs include:

- Salary of personnel working directly on the project – this could include, for example, PIs, postdoctoral research associates, research assistants, data managers, data scientists or software engineers.
- travel and subsistence for project researchers (e.g., attending conferences, travelling to/from the Turing/other collaborators).
- conference or event attendance fees (where conference/event is directly applicable to the research project).
- Cloud computing or other high performance computing costs.
- Other costs which are specifically justified for the project e.g., books, meeting room or catering costs, specific laptops.
- Open access publications.

Terms and conditions

The funding will be made available under The Alan Turing Institute's Defence & Security programme Research Service Agreement terms and conditions. For copy of the terms please contact Alaric Williams by emailing dsprogramme@turing.ac.uk.

You will be required to confirm your university's acceptance of these terms as part of this application process.

Background

The Machine Speed Strategic Analysis (MSSA) project is tasked with applying AI for Intelligence, Surveillance & Reconnaissance (ISR) of the sub-threshold information environment, focussing on military aspects. An important aspect of this environment is the spread of accidental misinformation, or deliberate disinformation, which can affect military operations or open-source intelligence gathering.

Within the sub-threshold domain, narratives can be spread to deceive, obfuscate, deflect, and disrupt. To help counter the threats of misinformation and disinformation, we are

investigating how to detect and track narratives, focusing on those relevant to the sub-threshold military environment. To begin this effort, we will define what constitutes a narrative, develop techniques to learn the characteristics of their seeding & spreading, and investigate ways to use this information to help analysts spot emerging narratives.

The benefit of this task will be to identify relevant sub-threshold indicators & warnings at greater scale and speed than current techniques, helping human analysts focus on the highest priority and most complex tasks. Critically, it may enable getting ahead of the narrative, substantially broadening the response options available.

The Requirement

In order to detect and track narratives, we first need to define what is meant by these terms. Standard definitions may derive from academic publications, but they may be arrived at practically, through discussions between D&S Partners, Turing and researchers. Such a common set of definitions - properly scoped and reflectively developed and maintained - will avoid misunderstanding and confusion. The questions we seek to answer, and the work expected to derive from them, are described below:

- What is a **narrative**? Per O.E.D., a narrative is “a spoken or written account of connected events; a story.” This can be either fiction or non-fiction, and indeed in the sub-threshold domain the two may be combined (e.g., a false representation of true events). We will need to confirm that our shared definitions of narrative and narrative structures are suitable for enabling AI techniques to identify and characterise narratives. D&S Partners are working on producing a common terminology to frame narrative discussion in a Defence context. This is the start of a journey and not the end because the term Narrative carries different nuance in different contexts in Defence. What we can do is share our current view on terminology, narrative evolution and other factors we deem important along with why we currently hold this view. This will be presented in a workshop with the researchers at or near the beginning of agreed work; noting that we expect the definition to evolve with the researchers input. A consequent part of this work will then be to work through how to dock the emergent working definition with the usage then current amongst Defence practitioners for whom we appear to be developing a useful product. The point being that busy practitioners generally have limited capacity to engage with new views, so to achieve uptake of something we believe to be useful for their work, we will need to frame things in their terms, along with clear statements of what is new to them and why. An output from this work will in consequence be:
 - A working definition of narrative, supported by consideration of what is being held onto, what reflection has occurred, what developments have been made and what shift of stances has occurred (in the sense described in her work by [Boulton \(2012\)](#))
 - Mapping between the working definition and practitioner views, including what in common is brought into focus, if there are gaps in the coverage of one view (when perceived from the other), what can usefully be said of such gaps, and if part of a view has been occluded so as to be beyond discussion in another’s terms and how that part may be described (perhaps from a third stance)
 - Preferred mapping from the working definition to a particular practitioner context and why this mapping is preferred
- How do we **detect** a narrative? Is there benefit in monitoring circumstances in which a proto-narrative could emerge? Should we defer the application of resources until the moment that characterisations of a situation could become a proto narrative? Should

we await the emergence of a proto narrative in relation to a single event / assertion? Could we benefit from deferring action until a more complete 'story' is presented? Could it be considered proportionate to search for single actors as a primary source of a narrative, or do we need to wait until the narrative has spread (i.e., is there a threshold at which we 'care')?

Concerning the technical means required, we note that many are already in existence, but may require integration into a service-based architecture for this specific task (E.g., [Baleen](#)). The modalities of available data will also need to be considered and prioritised, both from the perspective of storage requirements and also in terms of the benefit to be derived from necessary approaches to interpretation across these modalities. To draw this part of the work together we require a product covering the following issues:

- What approaches to narrative detection are practical to take and proportionate to enact?
- What are the key options for achieving this as part of a service-based architecture and what are the advantages and potential issues associated with each of the options that you have identified? Please support this by the conduct of and reference to an appropriate literature review, along with the collation of relevant open source tools, looking beyond the Defence & Security community (and thus also considering for example; advertising, analysis of viral social media posts, etc.).
- How could such options be used to falsify a view as further information emerges, then flag how an assessment has changed and why (note that it is important that such a mechanism is candid, prompt and yet not prone to 'flip flop' wildly between widely divergent views). Here the key issue concerns how to make good use of a practitioner's time to supplement their understanding, rather than impede it; appropriately building trust as the quality of available insight improves.
- Desirable characteristics of a narrative tracking tool:
 - From a purely analytical perspective, a final narrative detection and tracking tool would include all data sources relevant to UK analysts, we acknowledge that this may not be feasible when developing an initial proof-of-concept.
 - We also note the need for some narratives to be pre-defined by a user (e.g., search for narratives around topic X).
 - In addition, we need an unsupervised means of identifying narratives that may be of interest to a practitioner, but have not so far been specifically asked for.
 - It is also desirable that such a system identifies topics of potential interest, in order that a practitioner may develop a broader view, without overburdening the practitioner presented with this information.

On the basis of the option or options selected and other prior outputs, please produce a product to:

- Outline what the minimum initial specification of the system would need to be in order to develop a proof of concept (Technology Readiness Level 3) that is useful to the work of a practitioner while avoiding information or system overload. In other words, a proof of concept that uses available resources (including the practitioner) in appropriate moderation, while still being useful to the work of the practitioner; rather than identifying all narratives on any subject, regardless of the potential utility of such a view or the cost of the effort to achieve this.
- As part of this work, consider how, once initial narratives have been identified, the system may learn - based on attribution to sources, wording, etc. - what to look for to find new narratives of interest. Human-flagged

failures may also be used to teach a system to avoid certain uninteresting / incidental topics. We do not wish to be proscriptive at this stage - any technical means with the potential to successfully identify narratives may be considered. If successful narratives have certain characteristics and / or structures, finding such structures may allow for automated detection of new ones.

- Then consider what will be required to **track** a narrative? One possibility for tracking narratives could be to use graph analytics, with nodes & edges corresponding to entities and their relationships. Such a system would also be extensible. Dynamic topic models ([Blei & Lafferty, 2006](#)) could for example be used to characterise the evolution of prevalent narratives in online media over time. When tracking, we need to consider what information will need to be stored (e.g., website origin, author, etc.)?
- As part of this product, also consider how narratives can evolve over time, and when large semantic divergences from the original narrative may occur. At what point in time could it be appropriate to acknowledge that a 'new' narrative now exists, if at all? How should any such new narrative be viewed in relation to the 'old narrative'; that could be seen as its 'jumping off point'. How should the 'old' version of the narrative be thought of, in this situation? What should be done when one narrative splits into two separate 'paths'? In consequence of these additional considerations, would you propose a different initial specification and if so then in what ways, or where would you place these additional considerations in a 'development road map' for actioning. Please also share any 'road map' that you propose, along with supporting narrative concerning your underlying reasoning, then identify how an initial specification could be credibly demonstrated to a practitioner if it lacked these features.
- If a graph methodology is proposed, is there a risk of it becoming overly cumbersome and if so what simplifications could potentially be introduced to enable its practical use at scale? This may be important for a 'final' system required to track potentially dozens of narratives in relation to a single event.
- Then consider how will the information be presented to the user? Graph approaches will allow hierarchical visualisation, with links to primary sources. Peaks and troughs in activity could be plotted graphically, allowing analysts to observe trends, but what would you propose and why?

The purpose of the above products is to make initial inroads into the many questions that we have, and we believe that the earlier stages of the task should primarily focus on defining and detecting narratives - tracking them could be the subject of investigation later into the project. As such, part of this work (set out above) needs to set out the case by which we achieve an appropriate balance between practicality, ambition, and development of practitioner buy-in.

Re-shaping the task, to better achieve its objectives, can also be discussed between D&S Partners, Turing and researchers. There will be regular technical update meetings between D&S Partners, Turing and research staff, in addition to the outputs described above and the deliverables specified below.

Once an evidence based approach is agreed, the question of both software artefacts working towards a service based architecture, as demonstration of proof of concept, along with datasets to demonstrate emergent capability will need to be addressed. An initial proof of

concept will probably be built around a single scenario, which could be a known, and completed, past event. In this case, issues of acquisition, policy compliance (if relevant) and realism will need to be considered. In addition, advice and suggestions for future action would be useful output for the initial phase of work, with the potential to inform future effort.

Once the landscape has been assessed, and suitable data identified, an initial proof-of-concept demonstrator for narrative detection shall be agreed. It is then envisaged that on reaching suitable agreement, it shall then be produced. This is expected to highlight gaps, issues, and areas for future work, as much as it produces a functioning tool. Future phases of work - pending review and assent from D&S Partners technical staff - can then begin to address the tracking problem as well.

Follow on Options: Further Development of Narrative Tracking Capability

The purpose of these options is to provide greater understanding of narrative tracking and supporting visualisation, beyond that achieved through actioning the work described above. In the case of either option, communication between the supplier and the authority will be used - closer to the time of implementation - to determine the full scope and details of the chosen options.

Option 1

This option is to break new ground beyond the emergence of potentially interesting ideas, previously surfaced by this work. It is intended to build on observations, posits and more speculative thoughts concerning the things that could provide an improved capability to detect, track and transparently explain the spread of narrative. It could for example examine the interplay between our own and others narrative, to see where they are reaching and provide at least an indication of their rate of spread and of their effect. It is envisaged that a distinct sub-option shall be taken, of the form sub-option 1.1 ... 1.n, for each key aspect agreed between the authority and the Turing Institute as worthy of our effort, and important enough to look at now.

Option 2

Re-examine the implementation option space for ideas that are failing to work at sufficient scale to deliver their intended purpose in the intended way. These ideas may be improved upon, or replaced with alternative techniques. It is again envisaged that a distinct sub-option shall be taken, of the form sub-option 2.1 ... 2.n, for each key aspect agreed between the authority and the Turing institute as worthy of our effort and important enough to look at now.

Eligibility

To be eligible to apply you must:

- Be part of a UK university or research institute. Commercial organisations or overseas universities or institutes are not eligible.
- Have permission from your organisation to apply, i.e., ensure your organisation agrees to the Terms and Conditions provided and that you submit an approval of submission letter from your research/finance office stating this. An example of a letter is available on request.

How to apply

Applications must be submitted via the online portal at <https://ati.flexigrant.com/>. If you have not already done so, all applicants must first register on the system and provide basic details to create a profile. If you have any questions regarding the application form or using the online system, please contact the programme inbox dsprogramme@turing.ac.uk

Please use the budget template provided in the Flexigrant application form. Please note, applicants will also need to upload on Flexigrant an approval of submission letter from your research/finance office to confirm costs are correct. The Principal Investigator must ensure the same is received for all collaborators / universities on multi party applications.

We must receive your application by **1600 on Friday 30 September 2022**.

What should be in the proposal?

The proposal should:

- Be in Word or PDF format
- Describe the scope and technical approach of the proposed work - This should be a narrative description of the principals/solutions the project would aim to achieve and how those solutions may relate to the problem.
- Describe how the approach would lead to the desired results.
- Include a description of how the task is decomposed, thematically i.e. by work package. For each work package, what activities will be undertaken to produce the results.
- Answer questions such as: What is innovative about the proposed research? Why are you uniquely placed to undertake this task? What is the expected scientific impact?
- Include an option of follow on work to further the development of Narrative Tracking Capability
- Identify any key risks and mitigations.
- Reference related work or/and relevant experience.
- Include FEC cost.

If you are employed by one of the Institute's 13 university partners, please contact your University Liaison Manager ([a list of University liaison managers is available on the Turing website](#)) to make them aware of your application. They can provide support, answer questions and involve you as part of the Turing community at your university from now on.

If you are employed at a university that received a Turing Network Development Award, please contact your Award lead ([a list of Turing Network Development Award Leads](#) - scroll to the bottom of the page) – to make them aware of your application.

Deliverables

We expect the project to achieve the following deliverables

Ref.	Deliverable	Due by Start Date = T0	Format		Description
<i>D0a</i>	Project Kick off Meeting	T0+1 M	Meeting		Meeting to initiate the project
<i>D0b</i>	Technical update meetings	Monthly	Meeting		technical update meetings
<i>D1a</i> <i>D1b</i>	Quarterly reports	T0+3 M, T0+6 M	Technical Reports		An update on technical progress that deliver the outputs identified above (including academic literature and system design choices as appropriate).
<i>D2</i>	Final report	T0+8 M	Technical report		Final technical report describing the methodology and technologies used, along with a discussion of recommendations for next steps.
<i>D3</i>	Narrative detection proof-of-concept tool	T0+8 M	Software		Source code for initial proof-of-concept narrative detection tool, along with documentation, and instructions for install and use.
<i>D4</i>	End-of-phase presentation and demonstration	T0+8 M	Presentation + demonstration		Technical presentation summarising progress up to T+8M, plus demonstration of the initial proof-of-concept tool.
<i>D0c</i>	Project Closure Meeting	T0+8 M	Meeting		Meeting to close the project or agree follow on steps

D5	OPTIONAL - Further Development of Narrative Tracking Capability	TBC	Software, technical reports, demonstration(s)		Exact details will depend on the outcome of phase 1 (and discussion with Dstl staff). Likely to include further development of the narrative detection concept, plus an increased focus on tracking and visualisation for Defence analysts.
----	---	-----	---	--	---

Assessment and review

The assessment and review will follow the following stages:

- 1) Stage 1.0: Eligibility and triage
- 2) Stage 2.0: Expert review

Following eligibility checks, proposals will be reviewed by an assessment panel who will rank the proposals based on score.

The assessment panel will consider the following criteria:

- Quality: This will consider the method and concepts for the proposed research. This will assess if the methods are suitable for delivering the desired outputs and pushing forward fundamental understanding in the field.
- Viability: This will assess how feasible it is to practically carry out the proposed research, and if it can be delivered in the time frame. This will account for the difficulty of the tasks, logistical factors surrounding delivery, and the track record of the proposed research team.
- Significance: This will consider the relevance to the call and the themes that are represented.
- Justification of resources: This will consider whether the proposal is appropriately resourced and suitable expenditure has been included in the budget.

Each of the criteria will be scored and while all criteria will have equal weighting in evaluation, there will be a minimum requirement on significance to be considered for approval.

Key Dates

Deadlines are as follows

Activity	Date
Proposals to be Submitted*	Friday 30 September 2022
Announcement of Results	Monday 10 October 2022
Research Project Starts by**	Tuesday 01 November 2022
Initial Research completed and deliverables submitted	Friday 30 June 2023
6 month extension completed and deliverables submitted	Friday 06 January 2024

*Proposals must be submitted via Flexigrant by 16:00 Friday 30 September 2022.

**Any project agreements not signed by Friday 28 October 2022 may result in funding offer being withdrawn and going to an application on the reserve list

Post-award information

Project meetings

Successful applicants will be expected to attend a kick-off meeting and a project close meeting, with a Technical Partner from the D&S programme Partner/s. These may take place online, at the Turing, at Dstl, or at the project lead's university.

Screening of researchers

This research is not at a classified level so formal security clearance (see <https://www.gov.uk/guidance/security-vetting-and-clearance>) is not required. Successful applicants will however be required to complete a 'Personal Particulars - Research Workers & Dstl Sponsored PhD Students form' for all staff working on their project. A sample form can be requested by emailing dsprogramme@turing.ac.uk or found [here](#) (this is an example and the actual form may differ).

Outputs required

- We require that the project will produce the following:
- If applicable, the application to and approval from the relevant research ethics committee.
- Progress summaries (up to one page) / meetings.
- Literature review on the evaluation of known data fusion models and application to the Integrated Chemical Sense concept after the first month
- Delivery of theoretical model for data fusion in the Integrated at the end of the project.
- If applicable, any source code, compilation, use documentation and material associated with the outputs delivered.

Outputs acceptance criteria

Delivered research outputs must be commensurate with those standards expected in peer-review, open literature publications. Must include methodological details and examples of the use of the model sufficient to be reproduced. The theoretical model developed must be accompanied by clear explanations of its use and implementation suitable to allow non-data scientists/mathematicians to implement (e.g. Sensing SMEs). If outputs do not meet the acceptance criteria, re-work will be requested before final acceptance.

Publications

Please note, approval from the D&S programme is sometimes required prior to publication; in such cases, approval will not be unreasonably withheld.

The funders are committed to full and open publication of the research outputs in line with academic practices.

We encourage researchers to submit their findings to a high-quality peer-reviewed journal or conference, on an open-access basis (funding for open-access fees will be available on a case-by-case basis).

We expect a 'green' open access version of any papers to be published (if allowed by journal/conference - please check <http://www.sherpa.ac.uk/romeo/index.php>) either as a pre-print on (e.g.) the ArXiv (<https://arxiv.org/>) or in an institutional repository.

We also encourage datasets and research code to be openly shared too where possible - for example on the Turing's Github repository. All publications, reports and code should reference the support of the Turing Defence & Security programme.

Reporting and dissemination

Extracts from reports may be collated into update papers for the D&S Programme Board, Strategic Partners Board, Turing Innovations Ltd Board, and the Turing's Trustee Board.

Awardees may also be required to present their work to members of the D&S programme, the D&S Programme Board and/or other invited audience during the award period.

Reporting allows further identification and signposting of potential additional opportunities for the benefit of the awardees and the Turing; for example, opportunities from across the Turing's network such as new collaborations, external/public engagement, media/press, other funding availability, speaking slots at or invitations to events/conferences/seminars.

Queries

Please contact Alaric Williams, The Alan Turing Institute, Programme Manager
dsprogramme@turing.ac.uk.