Methodologies for Automated Landscape Scan and Gap Analysis

Date: July 19, 2019

To: NIST RFI on "Developing a Federal AI Standards Engagement Plan"

Submitted by: Marcello Balduccini (Saint Joseph's University)

We would like to thank NIST for giving the community the opportunity to comment on the proposed plan for federal engagement in AI standards.

In response to step 3 of the list "PRACTICAL STEPS FOR AGENCY ENGAGEMENT IN AI STANDARDS" ["PLAN FOR FEDERAL ENGAGEMENT IN AI STANDARDS -DRAFT FOR PUBLIC REVIEW 2-JUL-2019", pages 14-15], we would like to point out that the body of potentially relevant items is extremely large. Those preforming the scan will likely need to take into account general IT standards as well as other types of standards and best practices that are adjacent to the AI-focused standardization.

For a significant body of documents, manual processing may be time-consuming and open to subjective judgment, biased by an expert's familiarity with items and/or the expert's interpretation of the expectations. The result may not be inclusive, and will make it difficult to detect true gaps and also to evaluate building blocks in standards and related documents as opposed to their full text.

A possible way of limiting these issues could rely on the adoption of an approach for landscape scan and gap analysis focused on automated analysis. Probably, only the initial stages of the analysis will need to be automated. Several technologies may help in this endeavor. For instance, methodologies have recently been developed to scan document repositories and assess the relevance of the documents they contain to a topic based on statistical parameters. In one such case, a concept of interest might be defined by a collection of keywords, e.g. trustworthiness might be defined by privacy, security safety, reliability, resilience, and additional elements the definition of Trustworthiness may contain. The keywords might then be used to measure the degree of relatedness of each part of a structured document (e.g., a clause) to the concept of interest, e.g. by calculating the number of occurrences of the keywords in the heading of a document's clause. A more refined metric could also take into account synonyms, word senses and context, or reduce the degree of relevance in proportion to the length of the text.

An approach of this kind would allow one to obtain a comprehensive, largely unbiased view of relevant information already available. Additional stages of automation or review by humans, depending on the size of the body of relevant documents, would then be able to complete the task much more quickly.

In addition, the availability of such a sub-repository of likely relevant standards, specifications, and best practices would allow the reviewers to create an ontology that organizes the concepts from the documents in a hierarchical structure. Note that a distinctive feature of ontologies (and similar knowledge representation approaches) is the ability to specify properties of the ontology's components, including relationships between elements such as similarity and conflict. Ontologies additionally come with a set of well-understood automated reasoning mechanisms. As a result, it will be possible to leverage the ontology for a deeper scan of the documents, aiming to detect gaps, conflicts and affinities. The process can be guided by the user, who could pose queries such as “What are the typical components of the existing standards documents related to AI trustworthiness?” or “What components should be included in a new standard for Federated Learning?” Finally, when some level of automation is available, the results of the analysis could be updated as the new standards and best practices are created.