

Workshop on Language Resources for Responsible AI at LREC-2020

May 12, 2020, Marseille, France

Website: <https://sites.google.com/view/LR4Responsible-AI-2020>

This one-day workshop will provide a forum to present and discuss research work on creation and use of language resources and tools specifically designed to ensure ethical behavior of Artificial Intelligence (AI) systems.

Traditionally, AI systems have been developed to maximize accuracy on benchmark tasks and datasets. However, when these systems are deployed in the real world, ethical considerations need to be taken into account in order to build trust in users and make sure that the systems do not cause any harm to individuals or society. With the emergence of societal awareness about the need for responsible AI, new regulations and standards are being released, such as GDPR enforced by the European Union (2018), China's CyberSecurity Law and the General Principles of the Civil Law (2017), and Canadian national standards for the ethical design and use of Artificial Intelligence (2019). However, the technology in its current state lacks the necessary tools for the AI developers to comply with these regulations. There is an urgent need for tools that can help:

- Researchers - to investigate how ethical considerations should be taken into account while designing AI systems;
- Companies - to ensure their products meet ethical requirements, to apply ethics-by-design frameworks, and to gain the trust of their clients;
- End users - to be able to understand and to challenge automatic decisions when necessary;
- Policy makers and governments - to be able to audit and scrutinize AI systems for compliance with policies and regulations.

Developing responsible AI needs a change of paradigm from accuracy-optimized models to models that prioritize the ethical use. This shift requires a change in our way of thinking when building resources for NLP applications, including datasets, annotation schemes, language models, and evaluation metrics. For example, each manually annotated dataset should be accompanied with detailed information on the data sources, the data sampling process, the annotation process, and all the other important decisions (Data Statements). Datasets should represent input from diverse and representative sample of users and be annotated by workers with diverse backgrounds. When building or using language representations, such as word and sentence embeddings, researchers should be aware that such representations often perpetuate and accentuate unfair biases and require mitigating techniques. Evaluation metrics should take into account the ethics and fairness related costs associated with different kinds of errors. The carbon footprint created by computationally

demanding models should also be considered as cost when measuring the effectiveness of AI systems.

There is also a need for tools and resources that are designed for translating the technical jargon into simple language and providing explanations of automatic outcomes to non-expert users. Explainable AI will empower the society to scrutinize the algorithms for ethical use (or the risk of misuse) in specific applications, and ensure that everyone has a voice in defining and validating ethical AI.

Topics of Interest

We invite papers describing original research on design, creation, and use of language resources (annotated and unlabeled corpora, lexicons, dictionaries, templates, language representations, evaluation metrics, etc.) and tools to address any of the following issues in responsible AI:

- Fairness and unintended biases
- Confidentiality and privacy
- Interpretability and explainability
- Safety and security
- Transparency
- Accountability
- Integrity.

The language resources and tools can be designed for any one or several NLP (or non-NLP) applications, including (but not limited to):

- Syntax parsing and tagging
- Lexical semantics
- Language representation
- Discourse analysis
- Information retrieval
- Information extraction
- Natural language generation
- Textual inference
- Speech processing
- Dialogue systems
- Argument mining
- Sentiment and emotion analysis
- Machine translation
- Question answering
- Summarization
- Social media analysis
- Computational social science

- Health and wellness applications
- Auditing in highly regulated fields, such as medical, financial, and legal.

Paper Submission

We solicit original papers that describe language resources, evaluation metrics, and tools designed to assist in developing and assessing ethical AI systems. We also welcome papers highlighting ethics related problems in existing, widely used language resources (e.g., labeled datasets, word embeddings). We invite regular papers describing completed projects, emerging research papers presenting ongoing work, and position papers arguing an opinion on one of the topics of interest.

The papers can be up to 8 pages long (plus unlimited pages for references) and should be formatted according to the LREC style guidelines. The review process will be double-blind, so please do not include the authors' names and affiliations in the submission. The submissions will be reviewed by at least two members of the Program Committee. Accepted papers will be invited for an oral (or poster) presentation during the workshop and will be published as workshop proceedings at the LREC website. At least one author for each accepted paper has to attend the workshop to present the paper.

Submissions to multiple venues are allowed, but papers must be withdrawn from other venues if accepted by the workshop.

For submission instructions, please go to the [workshop website](#).

Identify, Describe and Share your LRs

- Describing your LRs in the LRE Map is now a normal practice in the submission procedure of LREC (introduced in 2010 and adopted by other conferences). To continue the efforts initiated at LREC 2014 about “Sharing LRs” (data, tools, web-services, etc.), authors will have the possibility, when submitting a paper, to upload LRs in a special LREC repository. This effort of sharing LRs, linked to the LRE Map for their description, may become a new “regular” feature for conferences in our field, thus contributing to creating a common repository where everyone can deposit and share data.
- As scientific work requires accurate citations of referenced work so as to allow the community to understand the whole context and also replicate the experiments conducted by other researchers, LREC 2020 endorses the need to uniquely Identify LRs through the use of the International Standard Language Resource Number (ISLRN, www.islrn.org), a Persistent Unique Identifier to be assigned to each Language Resource. The assignment of ISLRNs to LRs cited in LREC papers will be offered at submission time.

Important Dates:

Paper submission deadline: **Feb. 19, 2020**

Notification of acceptance: **Mar. 11, 2020**

Camera-ready paper deadline: **Apr. 2, 2020**

Workshop: **May 12, 2020**

All deadlines are 11.59 pm UTC -12h ("anywhere on Earth").

Program Committee

Marta Ruiz Costa-jussà, Universitat Politècnica de Catalunya

Sepideh Ghanavati, University of Maine

Randy Goebel, University of Alberta

Christian Hardmeier, Uppsala University

Graeme Hirst, University of Toronto

Stan Matwin, Dalhousie University

Noman Mohammed, University of Manitoba

Malvina Nissim, University of Groningen

Viviana Patti, Università di Torino

Sameer Singh, University of California, Irvine

Patricia Thaine, University of Toronto

Lyle Ungar, University of Pennsylvania

Organizing Committee

Svetlana Kiritchenko, National Research Council Canada

Isar Nejadgholi, National Research Council Canada

Contact Information

Website: <https://sites.google.com/view/LR4Responsible-AI-2020>

Email: LR4Responsible-AI@googlegroups.com