Deploy and operate ML applications in compliance with FATES requirements

Tristan Gouaichault
Institut de Recherche en Informatique de Toulouse (IRIT), France. Directed by O. Teste & J.-M. Bruel.
tristan.gouaichault@irit.fr

Objectives

FATES is a set of extra-functional properties of Machine Learning based systems that are becoming crucial to address. The five properties are not covered to the same extent by research. For example, there are far more resources in fairness than in the others. In addition to the disparity in the research volume, none of the properties has a standard method of measurement or correction.
To our knowledge, there is currently no systematic study or support to guide ML scientists and/or engineers on indicators for monitoring FATES properties. It impacts the entire software lifecycle in variable ways, depending on the issues and advances in the field.
By reusing and refining software engineering concepts and tools, our PhD aims to study FATES properties and propose a systematic tool-based approach for considering these fundamental properties in the lifecycle of software developed using MLOps frameworks. The key concerns are: (i) to characterize the FATES properties and their underlying concepts; (ii) link those definitions with their corresponding process (metrics, methods, potential correction methods, ...); (iii) integrate those definitions in an MLOps pipeline.

FATES Properties

Fairness
Accountability
Transparency
Ethics
Safety & Security

First introduced in 2014 with FAT/ML (for Fairness, Accountability, and Transparency in Machine Learning), it was later completed with the addition of Ethics and, more recently, safety.

MLOps

MLOps emerges from DevOps to facilitate the management of machine learning projects at all stages of the application construction and lifecycle [Testi2022].

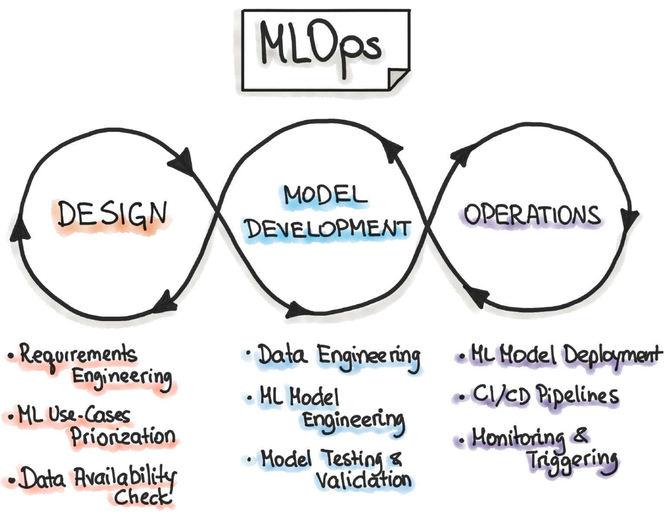

https://ml-ops.org/

References

[Testi2022]
Testi, M., Ballabio, M., Frontoni, E., Iannello, G., Moccia, S., Soda, P., & Vessio, G. (2022). MLOps: a taxonomy and a methodology. IEEE Access, 10, 63606-63618.

[Polacsek2018]
Polacsek, T., et al. (2018). The need of diagrams based on toulmin schema application: an aeronautical case study. EURO Journal on Decision Processes, 6(3-4), 257-282.

Key concerns

We plan to carry out the study and characterization of FATES properties by defining their requirements. As these definitions evolve over time and are increasingly systematically considered in the future, the definition framework will have to be highly adaptable. The form of these definitions (metamodel, ontologies, ...) is not yet determined, but it must support versioning, relationship definitions, and query capabilities.

Most of the properties can be interpreted in various ways, so none can be evaluated with a single and absolute metric. Each FATES property will be characterized by more precise concepts (such as *bias* for Fairness, for example), and each concept can be associated with several metrics.
Behind a metric, there is always a goal that can be translated into a concept. Knowledge of the many concepts covered by one of the FATES properties is essential to offer a set of metrics and corrections that are not redundant.
We will have to establish links between these high-level requirements defined for FATES properties, the algorithmic components, and the metrics or artifacts. These links will sometimes require the development of dedicated components responsible for verifying or even implementing certain properties' characteristics (e.g., anonymization).

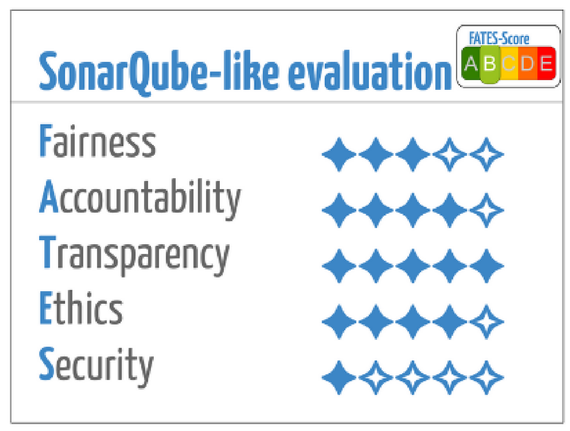

FATES properties are impacted by the entire MLOps pipeline. By defining the right part of the pipeline in which a process should be considered (a metric, an algorithm, a test, ...), we make it easier to produce ready-to-use tools, but also to inspect whether one of the FATES properties was taken into account, measured, in a given project.

Future work

In critical contexts, documentation plays an essential role in product accreditation, establishing confidence in the development and design process. The aim is to provide evidence that the development process has been properly managed and that standards have been respected. Justifying that a software system complies with FATES properties is a good way to achieve the same objective.

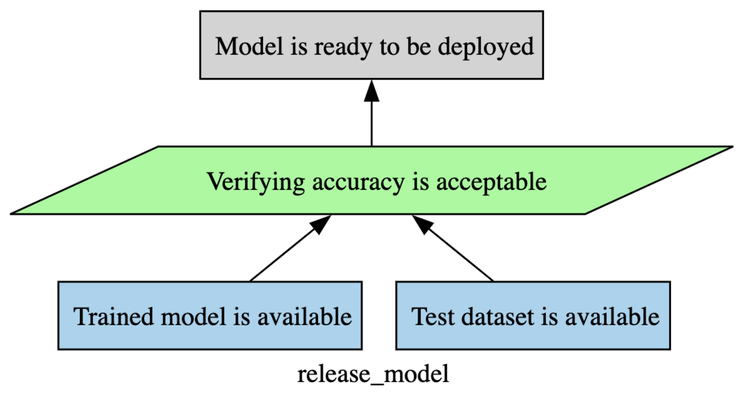

In [Polacsek2018], Polacsek et al. introduced justification diagrams (JD) to organize the elements that contribute to the justification of a result following IEC 62304.
We propose using justification diagrams to formally describe the reasoning behind a property's assessment.

Acknowledgments

This work is supported by the FATES-MLOps project funded by the French ANR TSIA 2024 – Specific Topics in Artificial Intelligence Under Grant Agreement No. ANR-24-IAS2-0002.