Association for Information Systems
AIS Electronic Library (AISeL)

AMCIS 2021 TREOs | TREO Papers

8-9-2021

On Engineering The Ethical Requirements of Machine Learning

Abdelnasser Abdel-Aal
King Faisal University, dr.abdelaal@gmail.com

Follow this and additional works at: https://aisel.aisnet.org/treos_amcis2021

Recommended Citation
Abdel-Aal, Abdelnasser, "On Engineering The Ethical Requirements of Machine Learning" (2021). *AMCIS 2021 TREOs*. 48.
https://aisel.aisnet.org/treos_amcis2021/48

This material is brought to you by the TREO Papers at AIS Electronic Library (AISeL). It has been accepted for inclusion in AMCIS 2021 TREOs by an authorized administrator of AIS Electronic Library (AISeL). For more information, please contact elibrary@aisnet.org.

On Engineering the Ethical Requirements of Machine Learning

TREO Talk Paper

Abdelnasser Abdelaal
King Faisal University, Saudi Arabia
aabdelaal@kfu.edu.sa

Abstract

The current software engineering traditions classify system properties into functional requirements and non-functional requirements. The first include operations, processes, inputs, and outputs. The second captures performance constraints and quality specifications such as security, reliability, maintainability, usability, flexibility, interoperability. These requirements go through a process called 'Requirement Engineering" (RE), a process in which engineers elicit, analyze, negotiate& specify, and validate them.

Machine Learning (ML) has brought up five ethical issues that disrupt the RE traditions and suggest a new set of requirements that may be called 'Ethical Requirements.' First, ML has the capabilities to imitate human behavior, learn, and redesign itself, and these capabilities require a continuous update of both types of requirements. Second, machines have become part of the stakeholders because of the human capabilities they have and the training they need. Third, performance of giant machines has triggered severe errors, bias, and discrimination that need to be considered in system analysis and design. Fourth, ML has human capabilities that require certain ethical principles such as learning, speech recognition, computer vision, search, optimization, planning, and decision-making. Thus, stakeholders suggest charging them with an appropriate load of ethical principles. One may ask what are the key elements of this new class of requirements? Can we employ traditional models of RE to manage them or they need a special treatment? What are the classification, challenges and controls needed to manage these emerging requirements?

In general, ethics are reshaped by local cultures and social contexts and this may require special methodologies to elicit values of system users to be employed for developing moral systems. Therefore, in order for designers to consider the 'values wished for,' system analysts, or ethics analysts, need instruments to elicit and represent human values and translate them into technical requiremen1ts. In other words, to build a moral system, it is important to make sure that principles advocated by philosophers and ethicists are not only spoken of but also enacted in the designs and system functionalities. Then, the challenge is not how to identify ethics of AI but how to engineer them and integrate them in its algorithms.

Building on traditional RE models, we suggest an extra layer of RE that focuses on the elicitation, analysis & negotiation, documentation, validation, and maintenance of the Culture, Ethics, Spirituality, and Context (CESC) of autonomous systems. In the elicitation and discovery phase, the ethics engineer should consider that ethics and culture span over development teams, standardization bodies, clients, users, ethicists and religious figures. Scholars hold that machines have become members of the stakeholders and they should be engaged in the process through iterative and interactive training. Our suggested model is dynamic and recurring as it changes with the culture, ethics environment, and the context. In the validation process, engineers check for requirement realism, consistency, and resolve any requirement conflicts. The maintenance phase will capture any changes in the legal, social, cultural, and moral settings.

It is important to note that some of AI ethical decisions are sensitive to training data. This is to say that RE should consider data requirements and specifications as well. If we take driverless vehicle as an example, we find that traffic regulations, road rules, speed limits, and safety measures vary across countries. In addition, ethical requirements in Europe might be different from that of Africa. These example show the importance of considering system context. There is also a need to distinguish soft ethical requirements from hard requirements. For instance, fairness could be a hard requirement in an automated decision-making machine but a soft requirement for care robots which may consider conditions of patients. In addition, privacy of those who need special care may not be a hard requirement as well. The outputs of this process will feed the design activities to develop ethically aligned artifacts. Consequently, the development team may need an ethics engineer to assist in managing AI ethical requirements needed by stakeholders.