Volume 2, No. 5 | September–October 2019

The Journal of Robotics, Artificial Intelligence & Law

Editor's Note: Much Ado About AI
Victoria Prussen Spears

AI and Text-Based Analytics in Complex Environmental Litigation: What Environmental Practitioners Need to Know
Jennifer Hammitt

FDA's AI White Paper: To Be or Not to Be, That Is the Question
Nathan A. Brown, Christin Helms Carey, Howard R. Sklamberg, and Marlee P. Gallant

Criminal Liability and Artificial General Intelligence
David Atkinson

Will the EU be an Exporter of Ethical Artificial Intelligence?
Sophie Goossens and Roch P. Glowacki

Drones and Urban Air Mobility in the President's Budget
Jennifer L. Richter and Mark Aitken

How Smart Cities and Connected Cars May Benefit from Each Other
Brooke Kahn, Thomas Parisi, and Gregory Discher

Everything Is Not *Terminator:* What Does the Executive Order Calling for Artificial Intelligence Standards Mean for AI Regulation?
John Frank Weaver

fastcase FULL COURT PRESS

RAIL
The Journal of Robotics, Artificial Intelligence & Law
Volume 2, No. 5 | September–October 2019

305 **Editor's Note: Much Ado About AI**
Victoria Prussen Spears

309 **AI and Text-Based Analytics in Complex Environmental Litigation: What Environmental Practitioners Need to Know**
Jennifer Hammitt

323 **FDA's AI White Paper: To Be or Not to Be, That Is the Question**
Nathan A. Brown, Christin Helms Carey, Howard R. Sklamberg, and Marlee P. Gallant

333 **Criminal Liability and Artificial General Intelligence**
David Atkinson

351 **Will the EU be an Exporter of Ethical Artificial Intelligence?**
Sophie Goossens and Roch P. Glowacki

357 **Drones and Urban Air Mobility in the President's Budget**
Jennifer L. Richter and Mark Aitken

367 **How Smart Cities and Connected Cars May Benefit from Each Other**
Brooke Kahn, Thomas Parisi, and Gregory Discher

373 **Everything Is Not *Terminator:* What Does the Executive Order Calling for Artificial Intelligence Standards Mean for AI Regulation?**
John Frank Weaver

EDITOR-IN-CHIEF

Steven A. Meyerowitz
President, Meyerowitz Communications Inc.

EDITOR

Victoria Prussen Spears
Senior Vice President, Meyerowitz Communications Inc.

BOARD OF EDITORS

Miranda Cole
Partner, Covington & Burling LLP

Kathryn DeBord
Partner & Chief Innovation Officer, Bryan Cave LLP

Melody Drummond Hansen
Partner, O'Melveny & Myers LLP

Paul B. Keller
Partner, Norton Rose Fulbright US LLP

Garry G. Mathiason
Shareholder, Littler Mendelson P.C.

Elaine D. Solomon
Partner, Blank Rome LLP

Linda J. Thayer
Partner, Finnegan, Henderson, Farabow, Garrett & Dunner LLP

Mercedes K. Tunstall
Partner, Pillsbury Winthrop Shaw Pittman LLP

Edward J. Walters
Chief Executive Officer, Fastcase Inc.

John Frank Weaver
Attorney, McLane Middleton, Professional Association

THE JOURNAL OF ROBOTICS, ARTIFICIAL INTELLIGENCE & LAW (ISSN 2575-5633 (print) /ISSN 2575-5617 (online) at $495.00 annually is published six times per year by Full Court Press, a Fastcase, Inc., imprint. Copyright 2019 Fastcase, Inc. No part of this journal may be reproduced in any form—by microfilm, xerography, or otherwise—or incorporated into any information retrieval system without the written permission of the copyright owner. For customer support, please contact Fastcase, Inc., 711 D St. NW, Suite 200, Washington, D.C. 20004, 202.999.4777 (phone), 202.521.3462 (fax), or email customer service at support@fastcase.com.

Publishing Staff
Publisher: Morgan Morrissette Wright
Journal Designer: Sharon D. Ray
Cover Art Design: Juan Bustamante

Cite this publication as:

The Journal of Robotics, Artificial Intelligence & Law (Fastcase)

This publication is sold with the understanding that the publisher is not engaged in rendering legal, accounting, or other professional services. If legal advice or other expert assistance is required, the services of a competent professional should be sought.

Copyright © 2019 Full Court Press, an imprint of Fastcase, Inc.

All Rights Reserved.

A Full Court Press, Fastcase, Inc., Publication

Editorial Office

711 D St. NW, Suite 200, Washington, D.C. 20004
https://www.fastcase.com/

POSTMASTER: Send address changes to THE JOURNAL OF ROBOTICS, ARTIFICIAL INTELLIGENCE & LAW, 711 D St. NW, Suite 200, Washington, D.C. 20004.

Articles and Submissions

Direct editorial inquires and send material for publication to:

Steven A. Meyerowitz, Editor-in-Chief, Meyerowitz Communications Inc., 26910 Grand Central Parkway, #18R, Floral Park, NY 11005, smeyerowitz@meyerowitzcommunications.com, 646.539.8300.

Material for publication is welcomed—articles, decisions, or other items of interest to attorneys and law firms, in-house counsel, corporate compliance officers, government agencies and their counsel, senior business executives, scientists, engineers, and anyone interested in the law governing artificial intelligence and robotics. This publication is designed to be accurate and authoritative, but neither the publisher nor the authors are rendering legal, accounting, or other professional services in this publication. If legal or other expert advice is desired, retain the services of an appropriate professional. The articles and columns reflect only the present considerations and views of the authors and do not necessarily reflect those of the firms or organizations with which they are affiliated, any of the former or present clients of the authors or their firms or organizations, or the editors or publisher.

QUESTIONS ABOUT THIS PUBLICATION?

For questions about the Editorial Content appearing in these volumes or reprint permission, please call:

Morgan Morrissette Wright, Publisher, Full Court Press at mwright@fastcase.com or at 202.999.4878

For questions or Sales and Customer Service:

Customer Service
Available 8am–8pm Eastern Time
866.773.2782 (phone)
support@fastcase.com (email)

Sales
202.999.4777 (phone)
sales@fastcase.com (email)
ISSN 2575-5633 (print)
ISSN 2575-5617 (online)

Will the EU be an Exporter of Ethical Artificial Intelligence?

Sophie Goossens and Roch P. Glowacki*

The authors of this article explain the European Union's guidelines on how those involved in deploying AI-powered solutions should go about doing so in an ethical manner.

In April, the European Union ("EU") published a set of guidelines on how those involved in deploying artificial intelligence ("AI")-powered solutions should go about doing so in an ethical manner. This follows the publication of the guidelines' first draft in December 2018 and a consultation process during which the expert group working on the document received over 500 comments. The guidelines propose a set of seven key requirements that AI systems should comply with in order to be deemed trustworthy. The document is likely to heavily influence discussions surrounding the EU's future regulatory landscape for AI.

Why Are the AI Ethics Guidelines Important?

In its work program for 2019, the European Commission stated that it wants to be the "effective standard-setter and global reference point on issues such as data protection, big data, artificial intelligence and automation."[1] The values enshrined in the EU's General Data Protection Regulation 2016/679 are already shaping the global economy and laws of other countries.[2] It is clear that the European Commission has the same ambitions for its AI-related initiatives. In fact, it has been reported that Brussels' "ethics-first approach has already attracted attention from outside Europe, including Australia, Japan, Canada and Singapore."[3]

No other technology raises comparable ethical concerns (or even outright fear) and technical challenges quite like AI. For example, how should a music recommendation engine react to an individual who is depressed or even suicidal (assuming that the device in question can measure this) and who chooses to continue listening to melancholic music? Would it be acceptable for the

Robotics, Artificial Intelligence & Law / September–October 2019, Vol. 2, No. 5, pp. 351–355.
© 2019 Full Court Press. All rights reserved.
ISSN 2575-5633 (print) /ISSN 2575-5617 (online)

machine to refuse to play the desired music or should it nudge the user to listen to something more upbeat? As humans delegate more and more decisions to machines, there is a serious concern as to what this will mean for human autonomy and well-being.

Take another, more abstract, example: how should an AI solution handle the infamous "trolley problem"?[4] This involves a trolley heading down railway tracks toward five people (tied up and unable to move) but which can be diverted to a different track to kill just one person. If the machine does nothing, five people will die. If it acts, just one dies. There are a number of variations of this problem (for example, by replacing the trolley with an autonomous vehicle) but the fundamental questions are the same. These include:

1. How to program AI-powered solutions to uphold ethical values;
2. How and who should decide what these values should be;
3. Who should be liable for such AI agent's decision; and
4. Should a human be able to step in and exercise a degree of oversight?

The Seven Requirements

The requirements put forward in the guidelines aim to provide a framework for analyzing and discussing the above-mentioned (and many other) issues. These requirements include:

1. *Human agency and oversight,* including evaluating AI systems in the context of fundamental rights, ensuring that users are able to make informed decisions regarding such systems and providing appropriate governance mechanisms;
2. *Technical robustness and safety,* including resilience to attack and security, fallback plans and general safety, accuracy, reliability and reproducibility of results;
3. *Privacy and data governance,* including respect for privacy, ensuring quality and integrity of data and appropriate controls on access to data;
4. *Transparency,* including traceability and explainability to enable identification of reasons why an AI decision was erroneous;

5. *Diversity, non-discrimination, and fairness,* including the avoidance of unfair bias, accessibility and universal design, and stakeholder participation;
6. *Societal and environmental well-being,* including sustainability and environmental friendliness, social impact, society, and democracy; and
7. *Accountability,* including auditability, minimization and reporting of negative impact, trade-offs, and redress mechanism.

The document is supplemented by a non-exhaustive Trustworthy AI assessment list (pilot version) that is intended to assist stakeholders operationalize Trustworthy AI.

Practical Considerations

Once an AI system makes a decision, it is important to understand why or how such a decision was reached. Without such information, a decision cannot be duly contested. The guidelines emphasize that explicability is crucial for building and maintaining users' trust in AI systems. However, the desire to achieve algorithmic transparency will need to be balanced against keeping information about business models and intellectual property related to the AI system confidential. The document acknowledges that current state-of-the-art technology is not, for example, capable of explaining AI systems that rely on neural networks. Some are also skeptical as to the value of such explanations. The often highly complex nature of the underlying data used in AI systems means that a meaningful explanation of their actions is likely to be incomprehensible to a layperson. However, it is likely that (in the future) safety-critical applications or those deployed in heavily regulated industries (such as health care or financial services) may need to be auditable and explainable.

The issues raised and the ethics-first approach enshrined in the guidelines is likely to influence the development of the European AI sector. In the future, AI-powered chatbots and other user-facing tools might need to contain disclaimers "AI-powered" or notices akin to the cookie consent banners. An entire sector focused on assessing and certifying the trustworthiness of AI solutions may emerge with new certifications ("made in Europe," etc.) or perhaps

even Trustpilot-type of websites for comparing AI-based solutions such as image-recognition software, recommendation engines, or instantaneous translators. Internal governance structures will also need to adapt to accommodate the complexity and pervasive nature of AI technology. Corporate governance policies already cover items such as privacy notices, data retention policies, and supply chain standards. There is every reason to believe that AI-related policies and procedures will become standard business practice.

What Comes Next?

The expert group clearly engaged with the 500+ responses received during the consultation process, many of which highlighted the need for a tailored approach to different use cases and the difficulties in achieving full explainability of AI systems. When the draft guidelines were first released in December 2018, the final version was intended to include a mechanism to allow stakeholders to voluntarily endorse them. Instead, a pilot phase is being set up to gather practical feedback on how the assessment list, that operationalizes the key requirements, could be improved. The pilot will commence in summer 2019 and interested stakeholders were invited to register their interest via the European AI Alliance. Feedback from the pilot will be used to re-evaluate the assessment list in early 2020. Based on this review, the Commission will propose any next steps.

Most recently, the same expert group that produced the guidelines has published the second deliverable consisting of AI Policy and Investment Recommendations. The document, released at the end of June, sets out 33 recommendations which are designed to ensure that Europe remains relevant in the global race for AI supremacy.

Notes

* Sophie Goossens is counsel at Reed Smith LLP where she leads the Entertainment and Media Industry practice in Paris and counsels media and technology companies. Roch P. Glowacki is an associate in the firm's Entertainment and Media Industry Group. The authors may be reached at sgoossens@reedsmith.com and rglowacki@reedsmith.com, respectively.

1. European Commission, "Commission Work Programme 2019, Delivering what we promised and preparing for the future," COM(2018) 800 final, October 23, 2018, page 4.

2. *See,* for example, California's Consumer Privacy Act.

3. Mehreen Khan and Madhumita Murgia, "EU publishes guidelines on ethical artificial intelligence," *Financial Times,* April 8, 2019.

4. Philippa Foot, "The Problem of Abortion and the Doctrine of the Double Effect," Oxford Review, No. 5, 1967.