



**The Journal of Robotics,
Artificial Intelligence & Law**

Editor's Note: Pandemic
Victoria Prussen Spears

Leading By Example Is Difficult: Europe's Approach to Regulating AI
Roch P. Glowacki and Elle Todd

Attorney General Charts Course for DOJ Counter-Drone Protection
James J. Quinlan and Elaine D. Solomon

What's in the FAA's Proposed Drone Remote Identification Rule
Brent Connor and Jason D. Tutrone

Insurance for Heightened Cyber Risk in the COVID-19 Era
Matthew G. Jeweler

Navigating Artificial Intelligence and Consumer Protection Laws in Wake of the COVID-19
Pandemic
Kwamina Thomas Williford, Anthony E. DiResta, and Esther D. Clovis

Does the FTC's Recent Influencer Guidance Address Robots?
Holly A. Melton

Second Circuit Takes Expansive Approach on the Definition of an ATDS
Jessica E. Salisbury-Copper, Scott A. King, and Doori Song

"Deepfakes" Pose Significant Market Risks for Public Companies: How Will You Respond?
Thaddeus D. Wilson, William T. Gordon, Aaron W. Lipson, and Brian M. Thavarajah

Artificial Intelligence at the Patent Trial and Appeal Board
Braden M. Katterheinrich, Ryan L. Duebner, and Sean Wei

Autonomous Vehicles, Ride Sharing, and the University
Louis Archambault and Kevin M. Levy

New Biometrics Lawsuits Signal Potential Legal Risks in AI
Debra R. Bernard, Susan Fahringer, and Nicola Menaldo

All Aboard! Major Shipping Lines Secure Antitrust Immunity for TradeLens Blockchain Agreement
Jeremy A. Herschaft and Matthew J. Thomas

Everything Is Not Terminator: An AI Hippocratic Oath
John Frank Weaver

- 293 Editor’s Note: Pandemic**
Victoria Prussen Spears
- 297 Leading By Example Is Difficult: Europe’s Approach to Regulating AI**
Roch P. Glowacki and Elle Todd
- 305 Attorney General Charts Course for DOJ Counter-Drone Protection**
James J. Quinlan and Elaine D. Solomon
- 311 What’s in the FAA’s Proposed Drone Remote Identification Rule**
Brent Connor and Jason D. Tutrone
- 317 Insurance for Heightened Cyber Risk in the COVID-19 Era**
Matthew G. Jeweler
- 323 Navigating Artificial Intelligence and Consumer Protection Laws in Wake of the COVID-19 Pandemic**
Kwamina Thomas Williford, Anthony E. DiResta, and Esther D. Clovis
- 329 Does the FTC’s Recent Influencer Guidance Address Robots?**
Holly A. Melton
- 333 Second Circuit Takes Expansive Approach on the Definition of an ATDS**
Jessica E. Salisbury-Copper, Scott A. King, and Doori Song
- 337 “Deepfakes” Pose Significant Market Risks for Public Companies: How Will You Respond?**
Thaddeus D. Wilson, William T. Gordon, Aaron W. Lipson, and Brian M. Thavarajah
- 341 Artificial Intelligence at the Patent Trial and Appeal Board**
Braden M. Katterheinrich, Ryan L. Duebner, and Sean Wei
- 347 Autonomous Vehicles, Ride Sharing, and the University**
Louis Archambault and Kevin M. Levy
- 353 New Biometrics Lawsuits Signal Potential Legal Risks in AI**
Debra R. Bernard, Susan Fahringer, and Nicola Menaldo
- 357 All Aboard! Major Shipping Lines Secure Antitrust Immunity for TradeLens Blockchain Agreement**
Jeremy A. Herschaft and Matthew J. Thomas
- 361 Everything Is Not *Terminator*: An AI Hippocratic Oath**
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

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 2020 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 Morrisette 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 © 2020 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 inquiries 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 contact:

Morgan Morrisette 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)

Everything Is Not *Terminator* An AI Hippocratic Oath

John Frank Weaver*

Periodically, I read opinion pieces advocating for a code of conduct governing the development of artificial intelligence (“AI”). The pieces typically recommend that developers, graduates of computer science programs, coders, etc. all sign an oath or agree to some sort of fiduciary duty before engaging in the creation of any AI application. The thought is that as AI grows more sophisticated, its uses expand, and the applications become part of more life-critical systems, the people developing those applications should adopt greater responsibility for how they develop them. In theory, it is similar to the Hippocratic Oath, which has historically obligated physicians to uphold certain ethical standards in their medical practice.

Below, after reviewing the Hippocratic Oath, I provide some examples of the oaths proposed for AI developers before providing a recommendation of my own.

Hippocratic Oath

According to the National Institutes of Health, new Greek physicians recited the Hippocratic Oath “to swear upon a number of healing gods” that they would uphold certain professional ethical standards.¹ It also bound new physicians to the community of physicians with responsibilities similar to that of a family member. The Oath has been updated frequently in order to reflect the values of different cultures using it. Although there is widespread belief in popular culture that all doctors take the Hippocratic Oath, most medical schools do not require it today.²

However, many schools still do, and the version they use was written in 1964 by Louis Lasagna, Academic Dean of Medicine at Tufts University:

I swear to fulfill, to the best of my ability and judgment,
this covenant:

I will respect the hard-won scientific gains of those physicians in whose steps I walk, and gladly share such knowledge as is mine with those who are to follow.

I will apply, for the benefit of the sick, all measures [that] are required, avoiding those twin traps of overtreatment and therapeutic nihilism.

I will remember that there is art to medicine as well as science, and that warmth, sympathy, and understanding may outweigh the surgeon's knife or the chemist's drug.

I will not be ashamed to say "I know not," nor will I fail to call in my colleagues when the skills of another are needed for a patient's recovery.

I will respect the privacy of my patients, for their problems are not disclosed to me that the world may know. Most especially must I tread with care in matters of life and death. If it is given me to save a life, all thanks. But it may also be within my power to take a life; this awesome responsibility must be faced with great humbleness and awareness of my own frailty. Above all, I must not play at God.

I will remember that I do not treat a fever chart, a cancerous growth, but a sick human being, whose illness may affect the person's family and economic stability. My responsibility includes these related problems, if I am to care adequately for the sick.

I will prevent disease whenever I can, for prevention is preferable to cure.

I will remember that I remain a member of society, with special obligations to all my fellow human beings, those sound of mind and body as well as the infirm.

If I do not violate this oath, may I enjoy life and art, respected while I live and remembered with affection thereafter. May I always act so as to preserve the finest traditions of my calling and may I long experience the joy of healing those who seek my help.³

Proposed Software Developers' Oath

Phillip Laplante, a professor of software and system engineering at Penn State Great Valley, offered an oath for software engineers in 2004 that is illustrative of what a similar oath for AI developers

could look like. He looked at the Hippocratic Oath and the nurse's version, the Nightingale Pledge, for inspiration. He noted that one of the most famous lines from the historic Hippocratic Oath—"First, do no harm"—served as a reminder to early physicians that their careers had the potential to injure if they are not careful.⁴ AI developers should keep the same advice in mind. The applications they create could be responsible for decisions that have legal impact on individuals or that analyze medical records for diagnoses and treatment. The designers should therefore remember that their applications should eliminate or minimize harm. Laplante warns, "In the course of fixing a problem we sometimes do more harm than good." That is good advice to keep in mind.

Quoting Thomas Carlyle, Laplante states that "Nothing is more terrible than activity without insight" and wonders if software engineers are doing exactly that, programming in languages they do not fully understand to perform tasks using functions they cannot fully explain.⁵ AI developers are frequently in the same position, creating algorithms and machine learning applications that operate in black boxes, delivering outputs the programmers cannot always explain. One of the potential benefits of the "right to an explanation"—popularized in a number of sources recently, including the European Union's General Data Protection Regulation—is that it forces developers to be able to explain why their AI applications produce certain results, to provide insight into its activity.⁶

Laplante relies on the Nightingale Pledge to propose an oath for software developers:

I solemnly pledge, first, to do no harm to the software entrusted to me; to not knowingly adopt any harmful practice, nor to adopt any practice or tool that I do not fully understand. With fervor, I promise to abstain from whatever is deleterious and mischievous. I will do all in my power to expand my skills and understanding, and will maintain and elevate the standard of my profession. With loyalty will I endeavor to aid the stakeholders, to hold in confidence all information that comes to my knowledge in the practice of my calling, and to devote myself to the welfare of the project committed to my care.⁷

Some of the provisions in this oath lend themselves more to a general professional creed ("devote myself to the welfare of the project"), but we can carve out a few provisions that are particularly

valuable for AI developers. As mentioned above, the “do no harm” ethos is appropriate, as is vowing to abstain from “whatever is deleterious and mischievous.” Similarly, promising not to adopt any practice or tool the engineer does not fully understand is important. I also like maintaining and elevating the standard of the profession, although having a clearer definition of the standard would be helpful. For further direction, though, we should turn to more recent and AI-specific examples.

Examples of AI Developer Oaths

When Oren Etzioni, the chief executive officer of the Allen Institute for Artificial Intelligence, proposed a Hippocratic Oath for AI “practitioners,” he noted that historically “much power and responsibility over life and death was concentrated in the hands of doctors. Now, this ethical burden is increasingly shared by the builders of AI software.”⁸ He comments that other technologies—cloud computing, smartphones, social media platforms, etc.—also collect and use a lot of sensitive and personal information, with corresponding potential to cause problems for individuals regarding privacy, profiling, manipulation, and personal safety. But he goes on to accurately differentiate AI from those technologies by writing that “[i]t is these issues that AI, if not developed responsibly, will further amplify.”⁹

With that in mind, he proposed to revise the Hippocratic Oath to apply to AI developers specifically:

I swear to fulfill, to the best of my ability and judgment, this covenant:

I will respect the hard-won scientific gains of those scientists and engineers in whose steps I walk, and gladly share such knowledge as is mine with those who are to follow.

I will apply, for the benefit of the humanity, all measures required, avoiding those twin traps of over-optimism and uniformed pessimism.

I will remember that there is an art to AI as well as science, and that human concerns outweigh technological ones.

Most especially must I tread with care in matters of life and death. If it is given me to save a life using AI, all thanks. But it may also be within AI’s power to take a life; this awesome

responsibility must be faced with great humbleness and awareness of my own frailty and the limitations of AI. Above all, I must not play at God nor let my technology do so.

I will respect the privacy of humans for their personal data are not disclosed to AI systems so that the world may know.

I will consider the impact of my work on fairness both in perpetuating historical biases, which is caused by the blind extrapolation from past data to future predictions, and in creating new conditions that increase economic or other inequality.

My AI will prevent harm whenever it can, for prevention is preferable to cure.

My AI will seek to collaborate with people for the greater good, rather than usurp the human role and supplant them.

I will remember that I am not encountering dry data, mere zeros and ones, but human beings, whose interactions with my AI software may affect the person's freedom, family, or economic stability. My responsibility includes these related problems.

I will remember that I remain a member of society, with special obligations to all my fellow human beings.¹⁰

Similarly, The Academy.ai in Barcelona requires its AI students to sign an ethical contract:

I will apply AI towards the benefit of humanity at all costs.

I will respect every human's privacy as if it was my own.

I will do everything in my power to acquire knowledge and share it with others.

I will set positive models for others to emulate.

I will consider the impact of my models and disobey unjust requests.

I will train my models again and again until I succeed.

I will consider the impact of historical and new bias in my work.

I will preserve human concerns over technological ones.

I will work to create a new set of conditions that reduce inequalities.

My AI models will be designed to prevent harm at all costs.

I will keep my word.¹¹

Jan Carbonell, the chief executive officer, says that the academy tries to give its students “an understanding of the other implications of AI and how they can make a positive impact in society.”¹²

There are some key points to parse out of these proposals.

First, the idea that human concerns outweigh technological ones must be fundamental to AI development.

Second, AI applications will have access to tremendous personal information, and the developers of those applications need to remember that that information is not abstract data, but important to other people.

Finally, recognizing that AI applications need to be trained to run properly and remove bias is vital to developing AI that is widely beneficial.

Proposed AI Oath

The examples above are fine, but in each case there are significant issues. Laplante’s reads like it was not written to apply to AI, which is, of course, the case. Etzioni, by his own admission, uses the Hippocratic Oath as a model, which shows in places where the language or format more relevant to medical practitioners. The oath from The Academy.ai is very altruistic, perhaps more so than is realistic in the profession, or even desirable.

Using the examples above as models, drawing out the key points I identified, and relying on important concepts that I frequently return to when advocating for the regulation of AI, I propose the following as an oath for AI developers:

As I develop software, applications, systems, and programs that rely on or incorporate machine learning, algorithmic analysis, and other forms of artificial intelligence that can make qualitative decisions without input or action by human beings, I promise to:

1. Consider the impact of my work on people;
2. Prioritize the people who my work will impact over the technological achievements of my work;
3. Respect and remember the other people in my field whose work has contributed to or is contributing to my work;
4. Be inventive in my work;

5. Understand how my work functions, in all ways;
6. Be able to explain how my work functions and how it produces any particular output;
7. Work to minimize and reduce any harmful effects my work may produce;
8. Remember that my work may interact with the personal information of other people and that their information should only be used in ways and for purposes they have consented to;
9. Do my best to remove any bias that may impact my work and that my work may produce in order to reduce inequality; and
10. Remember that I am a member of society and that I and my work have obligations of consideration, hard work, and kindness to other people in society

Many of the ideas expressed in this oath will be familiar from the above. Each line item has the reciter's reference "my work" in order to emphasize that the AI is separate from them, but also that they own it. It does not refer to maintaining or establishing industry standards for AI, but by its terms it creates aspirational standards for each reciter: consideration of people; inventiveness; respect for the work of others; reduction in inequality; and the importance of the societal good.

How should this oath be administered, and is there any way to make it legally binding? Unlike physicians, there is no school that you have to graduate from to become an AI developer, so there is no easy point of entry where all or nearly all developers have to take this oath. I encourage companies, academic institutions, and government agencies that work with AI applications and software to require this oath from new hires that work with AI. If this code were included in employee handbooks, the terms of which typically become conditions of employment per employment contracts, upholding the oath would be legally binding, at least between the employee and the employer. In the event that any forms of AI development require licensing in the future, all licensees should be required to take this oath.

Even though this oath and its administration provide limited options to create fiduciary or legally binding obligations on individual AI developers, the oath is able to highlight important considerations for AI developers as they create new applications.

Those considerations might exist only in their minds, and not in law, but frequently that is a more effective place for them to be.

Notes

* John Frank Weaver, a member of McLane Middleton's privacy and data security practice group, is a member of the Board of Editors of *The Journal of Robotics, Artificial Intelligence & Law* and writes its "Everything Is Not Terminator" column. Mr. Weaver, who may be contacted at john.weaver@mclane.com, has a diverse technology practice that focuses on information security, data privacy, and emerging technologies, including artificial intelligence, self-driving vehicles, and drones.

1. Greek Medicine, National Institutes of Health, *available at* https://www.nlm.nih.gov/hmd/greek/greek_oath.html ("Greek Medicine Page").

2. *Id.*

3. The Hippocratic Oath: Modern Version, PBS, *available at* https://www.pbs.org/wgbh/nova/doctors/oath_modern.html.

4. Kevin A. Laplante, "First Do No Harm: A Hippocratic Oath for Software Developers?," *Development* (Vol. 2, No. 4; August 31, 2004), *available at* <https://queue.acm.org/detail.cfm?id=1016991>. Despite that phrase's fame, it does not traditionally appear in the Hippocratic Oath. Greek Medicine Page, *supra* note 1.

5. Laplante, *supra* note 4.

6. *See* Council Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC, 2016 O.J. (L119) 1, Art. 13(2)(f); Bryce Goodman & Seth Flaxman, "European Union regulations on algorithmic decision-making and a 'right to explanation,'" 6-7, August 31, 2016, *arXiv.org*, <https://arxiv.org/pdf/1606.08813.pdf>; John Frank Weaver, "Artificial Intelligence Owes You an Explanation," *Slate*, May 8, 2017, http://www.slate.com/articles/technology/future_tense/2017/05/why_artificial_intelligences_should_have_to_explain_their_actions.html.

7. Laplante, *supra* note 4.

8. Oren Etzioni, "A Hippocratic Oath for artificial intelligence," *TechCrunch*, March 14, 2018, <https://techcrunch.com/2018/03/14/a-hippocratic-oath-for-artificial-intelligence-practitioners/>.

9. *Id.*

10. *Id.*

11. Navanwita Sachdev, "AI Academy asks students to sign ethical contract akin to Hippocratic Oath," *The Sociable*, April 15, 2019, <https://sociable.co/technology/ai-academy-asks-students-sign-ethical-contract-hippocratic-oath/>.

12. *Id.*