



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

Editor's Note: Big Data

Steven A. Meyerowitz

Big Data, Bigger Risk: Recognizing and Managing the Perils of Using Algorithms in Recruiting and Hiring

Mark J. Girouard

Artificial Intelligence and Associated Clinical Data Privacy Considerations

James A. Sherer, Taylor M. Hoffman, and Emily R. Fedeles

Risk and Opportunity with the Industrial Internet of Things

David T. Doot, Steven A. Cash, and James B. Blackburn, IV

The Robo-Criminal

Nina Scholten

AI and Image Recognition: The Next Generation Brand Protection?

Ronda Majure

Everything Is Not *Terminator*: Helping AI to Comply with the Federal Trade Commission Act

John Frank Weaver

- 231 Editor's Note: Big Data**
Steven A. Meyerowitz
- 235 Big Data, Bigger Risk: Recognizing and Managing the Perils of Using Algorithms in Recruiting and Hiring**
Mark J. Girouard
- 243 Artificial Intelligence and Associated Clinical Data Privacy Considerations**
James A. Sherer, Taylor M. Hoffman, and Emily R. Fedeles
- 257 Risk and Opportunity with the Industrial Internet of Things**
David T. Doot, Steven A. Cash, and James B. Blackburn, IV
- 263 The Robo-Criminal**
Nina Scholten
- 285 AI and Image Recognition: The Next Generation Brand Protection?**
Ronda Majure
- 291 Everything Is Not *Terminator*: Helping AI to Comply with the Federal Trade Commission Act**
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 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 © 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 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 call:

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*

Helping AI to Comply with the Federal Trade Commission Act

John Frank Weaver*

In my last column, I noted that there is general agreement that the authority of the Federal Trade Commission Act (the “Act”) is broad enough to govern algorithmic decision-making and other forms of artificial intelligence (“AI”).¹ Section 5(a) the Act prohibits “unfair or deceptive acts or practices in or affecting commerce” as unlawful.² The Federal Trade Commission (“FTC”) is authorized to challenge such acts or practices through administrative adjudication and to promulgate regulations to address unfair or deceptive practices that occur widely by multiple parties in the market.³

The FTC has a department that focuses on algorithmic transparency, the Office of Technology Research and Investigation, and has requested public comment on and scheduled hearings about algorithmic decision-making and AI.⁴ The FTC has already started adjudicating complaints related to unfair practices involving AI, such as the complaint the Electronic Privacy Information Center (“EPIC”) filed against Universal Tennis Rating (“UTR”), in which EPIC alleged that UTR relied on “a secret algorithm to score children” tennis players, which “created a substantial risk of harm because children’s development, educational, scholarship, and employment opportunities may be unfairly hindered by low and inaccurate scores, the calculation of which is secret and the validity of which parents are not permitted to dispute.”⁵

There is not clear guidance from the FTC at this point on what qualifies as an unfair or deceptive, and therefore unlawful, method of using AI. However, there are general practices that organizations can adopt that will help them minimize their potential for violating the Act:

1. Establish a governing structure;
2. Establish policies, internal and external, addressing the use and/or sale of AI and AI-reliant products;

3. Establish notice procedures;
4. Assess AI and algorithms for bias; and
5. Make sure third-party agreements properly allocate liability and responsibility.

Below, I briefly outline each recommended practice and provide suggestions for how organizations can adopt each one.

Governing Structure

The first step is to establish the group or the individual within the organization that will review AI implementation with an eye toward complying with the Act. Some organizations approach AI like any other technology or software update, but that is a mistake. AI is much more likely to introduce issues that are unique in terms of business operations, customer relations, and branding; organizations should implement a governing structure that creates a rubric for reviewing each AI proposal. That rubric can include the organization's philosophical concerns, legal interpretations, operations concerns, marketing, and branding concerns.

The governing structure does not have to be complicated. Rather, the size and composition of the governing structure should reflect the size and composition of the organization. Large companies that have sophisticated AI programs should have a group composed of key stakeholders. That might be the board of directors, a board committee, a committee formed from C-suite officers, etc. A smaller organization with more limited AI needs may designate only the president or the vice president of information technology to review each AI proposal in light of established principles.

Fortunately, there are plenty of resources an organization can rely on when drafting those principles. For example, the Partnership on AI, an organization founded by several technology companies, is working to develop best practices for fair, transparent, and accountable AI; it has committed to making its research into the ethical, social, economic, and legal implications of AI open to the public.⁶ Similarly, the Software and Information Industry Association has published a brief on ethical principles for AI and data analytics,⁷ and the Institute of Electrical and Electronics Engineers ("IEEE") has published a treatise that attempts to provide recommendations on best practices, philosophies, and legal and ethical considerations

for AI.⁸ Any organization can review these to determine which principles and considerations are important to them and their compliance with the Act.

Policies

I have already written in this space about public-facing AI policies,⁹ and organizations should use them as their first public effort to demonstrate compliance with the Act to the FTC and consumers. However, internal policies are important as well. Properly written employee policies will establish how the governing structure incorporates elements of outside guidance and clearly states how the organization makes decisions about AI. These policies should not be huge documents, but they should be detailed enough that they are useful to the staff implementing AI. By following the policy, employees will comply with the Act and enact the organization's vision and beliefs for AI.

Notice

Notice is a key element to complying with the Act. In order to avoid using AI in an unfair and deceptive manner, informing consumers and concerned individuals is vital. The extent of the notice is important too, as notifying the individuals who are in an AI's training dataset can be just as important as informing customers who may interact with an organization's AI. A loan applicant should know if the lender relies on algorithmic decision-making to approve a mortgage, but IBM recently also ran into trouble because it did not inform the relevant people when the company used their Flickr photos to train its facial recognition AI.¹⁰ Even though the FTC has not issued an affirmative regulation about this, the trend in the industry is clearly running toward greater disclosure of AI usage. For example, IEEE has suggested a "government-approved labeling system like the skull and crossbones found on household cleaning supplies that contain poisonous compounds could be used for this purpose to improve the chances that users are aware when they are interacting with" AI.¹¹ While this is a ridiculously loaded comparison ("skull and crossbones," "poisonous compounds"), the point is clear.

Notice can take a variety of forms, depending on the AI in question. If the organization's website relies on AI to analyze website usage, a pop-up directing visitors to the organization's AI policy, like pop-ups that deliver privacy policies, is appropriate.¹² Similarly, if an organization uses autonomous chatbots as part of its customer service, the bot should clearly state that it is a bot to the consumer.¹³ Organizations that rely on AI in other forms or in other stages of their business operations should consider the most effective form of notice. In the example of the loan applicant above, the lender should include a clear statement in the application form, whether that is online or in hard copy, that the final decision will incorporate algorithmic decision-making.

Informing individuals whose data is being used to train an AI application can be much more difficult, as IBM's experience demonstrates. If an organization is collecting the data itself to train a specific application, it can provide notice directly to the data subjects. If an organization is obtaining datasets from a third-party vendor, they will likely need to rely on the third party to notify the participants; agreements with those vendors should address this, as discussed below. Alternatively, the organization can notify the data subjects the vendor used, but that may be impossible if the data is anonymized or the vendor does not have their contact information.

The FTC is likely to view notice in this context as a balancing act, weighing the interest of the data subjects to be notified that their data is being used to train the organization's AI against the cost and difficulty of informing them. The key issue is whether the use of the information in the dataset is unfair or deceptive. If there is no notice to data subjects in a dataset, the FTC will look at whether the data subjects were disadvantaged by the AI training and the extent to which each data subject would have behaved differently in providing his or her data if he or she knew it would be used to train the AI. At the time of this writing, the FTC is not considering any complaints against IBM regarding its use of Flickr photographs, but it is easy to see both how the issue could have been avoided with better notice and the difficulty in providing that notice.

Similar to providing notice, organizations should attempt to design and implement algorithms that allow key stakeholders—consumers, employees, vendors, leadership, etc.—to understand how and why AI applications make decisions, *e.g.*, the factors the AI weighs more heavily than others, data the AI does not consider,

etc. Admittedly, this is an easy to express concept that is difficult to execute, but it is important that organizations can show they are making good faith efforts to avoid AI that acts in an unfair or deceptive manner. Trying to make their AI more understandable to all parties involved is a good way.

Assessing for Bias

One of the greatest concerns even the strongest supporters of AI have is its tendency to incorporate bias into decisions. I recently discussed biased AI in hiring,¹⁴ criminal sentencing,¹⁵ and lending.¹⁶ Bias is, in and of itself, a problem, but bias in AI is particularly troublesome because consumers typically have no access to how AI makes its decisions. This is commonly referred to as the “black box” problem: data enters the AI’s black box, the algorithm in the black box analyzes the data, and the black box produces a decision based on the data. Except for a small number of key people in the organization, no one knows how the AI makes the decision. Although there are no regulations under the Act governing this directly, the FTC is actively exploring rules, and organizations need to be careful.¹⁷

Absent specific regulations, the best strategy to avoid FTC action due to impermissible bias is to conduct regular tests. If there is an investigation, an organization wants to show a history of checking its AI for bias. I also recommend involving outside counsel in that process, both to comply with the Act (as well as other federal and state laws, as state attorneys general are also investigating bias under their states’ consumer protection acts) and to protect with attorney/client privilege the test results from regulators and discovery during litigation.

In testing for bias, identify the types of bias that might be a problem: race, gender, age, etc. After that, the organization should create test datasets that will demonstrate whether or not the AI can properly incorporate data regarding the areas of concern without evidencing bias. If the AI is used to make hiring decisions, but the organization is worried it will evidence a preference for hiring men, the test dataset should be designed to show how the AI application incorporates gender into its final hiring decisions.

If the test dataset has returned results that indicate impermissible bias is baked into the AI’s algorithm, the organization needs

to show efforts to reduce and eliminate that bias in order to comply with the Act. This involves training the AI using other datasets, which are designed to teach the AI how to incorporate data without evidencing impermissible bias, *i.e.*, machine learning. In the hiring example, the datasets should train the AI to ignore applicants' gender or to favor women in order to counteract the existing training that led the AI to favor men.

An organization that can show a history of testing its AI and attempting to remediate any impermissible bias it discovers will have a strong defense against any FTC action.

Third-Party Agreements

Part of helping to make sure that an organization's AI complies with the Act is to make sure that its vendors and contractual partners comply with the Act. It is not enough to assume. Organizations need to include language in their contracts in which the appropriate parties (a) represent that the relevant individuals have received notice or given consent, (b) provide proof of that notice or consent, and (c) indemnify the other party for losses and costs caused by the relevant AI at issue.

Aggressive and/or sophisticated organizations may also seek to assign most or all liability to the other party, even when that is not appropriate given the responsibilities of the parties. For example, in a contract where an organization agrees to provide AI analysis of website usage for a third party, the party who maintains the website should represent that it provides notice of the AI while the organization performing the analysis should indemnify the third party for all losses and damages caused by the AI, generally speaking. However, if the analyzing organization is aggressive, it might attempt to assign all liability for the AI to the website operator under the theory that the AI is only being used on behalf of the website.

At this point, that type of assignment is permitted. It is possible that some assignments of liability associated with AI will be prohibited in the future. Similar statutes and regulations govern liability in other contexts. Some states require landlords to accept liability for their negligence and willful misconduct, making void any lease clause that would force the tenant to release the landlord of such liability.¹⁸ Under the European Union's General Data Protection Regulation, the processor of an individual's personal data

is liable to individuals for a subprocessor's violations; it cannot contract that liability away.¹⁹ Until similar prohibitions exist for liability associated with AI, organizations may try to aggressively limit their own risk exposure. For this reason alone, organizations should review their contracts with third parties to make sure AI representations and liability are properly addressed.

But reviewing those contracts is also part of complying with the Act. For example, an organization that obtains datasets from a third party should review the contract to make sure that the dataset provider represents it has given notice to or obtained consent from the relevant individuals, that the organizations can review documentation to confirm such notice or consent, and that the provider indemnifies the organization for losses and costs caused by the provider failing to give notice or obtain consent. If there is a complaint against the organization, its failure to take such precautions could lead the FTC to determine that it had engaged in unfair or deceptive trade practices because it aided and abetted the third-party dataset provider.

Conclusion

By following these practices, organizations will have a strong defense in the event a consumer files a complaint with the FTC. Even in the absence of the FTC and the Act, I recommend the above as best practices that organizations should adopt. They help the organization make thoughtful decisions about AI, allow the organization to develop a desirable brand in AI management with consumers, and give consumers appropriate notice and protection regarding potentially harmful AI.

Notes

* John Frank Weaver, an associate at McLane Middleton and a member of the firm's privacy and data security practice group, is the "Everything Is Not *Terminator*" columnist for *The Journal of Robotics, Artificial Intelligence & Law*. Mr. Weaver, who may be contacted at john.weaver@mcclane.com, has a diverse practice that focuses on land use, real estate, telecommunications, and emerging technologies, including artificial intelligence, self-driving vehicles, and drones.

1. John Frank Weaver, “Everything Is Not *Terminator*: Value-Based Regulation of Artificial Intelligence,” *The Journal of Robotics, Artificial Intelligence & Law* (Vol. 2, No. 3; May-June 2019).

2. 15 U.S.C. § 45(a)(1).

3. 15 U.S.C. § 45(b); 15 U.S.C. § 57a.

4. Press Release, “FTC Announces Hearings On Competition and Consumer Protection in the 21st Century,” Federal Trade Commission (June 20, 2018), <https://www.ftc.gov/news-events/press-releases/2018/06/ftc-announces-hearings-competition-consumer-protection-21st>.

5. Complaint, Request for Investigation, Injunction, and Other Relief Submitted by Electronic Privacy Information Center at 12, *In re: Universal Tennis* (Federal Trade Commission, May 17, 2017).

6. Partnership on AI, About Us, <https://www.partnershiponai.org/about/> (accessed on 3-14-2019); Partnership on AI, Tenets, <https://www.partnershiponai.org/tenets/> (accessed on 3-14-2019).

7. *SIIA Issue Brief: Ethical Principles for Artificial Intelligence and Data Analytics*, Software and Information Industry Association, available at <http://www.siiia.net/Portals/0/pdf/Policy/Ethical%20Principles%20for%20Artificial%20Intelligence%20and%20Data%20Analytics%20SIIA%20Issue%20Brief.pdf?ver=2017-11-06-160346-990>.

8. *Ethically Aligned Design*, v.2, Institute of Electrical and Electronics Engineers, available at http://standards.ieee.org/develop/indconn/ec/ead_v2.pdf (“Ethically Aligned Design”). Full disclosure—I am a member of the lawyers committee that contributed to *Ethically Aligned Design*.

9. John Frank Weaver, “Everything Is Not *Terminator*: Public-Facing Artificial Intelligence Policies—Part I,” *The Journal of Artificial Intelligence & Law* (Vol. 2, No. 1; January-February 2019); John Frank Weaver, “Everything Is Not *Terminator*: Public-Facing Artificial Intelligence Policies—Part II,” *The Journal of Artificial Intelligence & Law* (Vol. 2, No. 2; March-April 2019).

10. Shannon Liao, “IBM didn’t inform people when it used their Flickr photos for facial recognition training,” *The Verge* (March 12, 2019), <https://www.theverge.com/2019/3/12/18262646/ibm-didnt-inform-people-when-it-used-their-flickr-photos-for-facial-recognition-training>.

11. *Ethically Aligned Design*, *supra* note 8, at 159.

12. In the European Union, this requirement is explicit, as the General Data Protection Regulation grants data subjects the right not to be subject to decisions based solely on automated processing, including profiling (*i.e.*, decisions made by AI), which produces legal effects concerning him or her or similarly significantly affects him or her. 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, Article 22(1) (the “GDPR”).

13. In California, state law requires this disclosure to lawfully use bots to incentivize a purchase or sale of goods or services in a commercial transaction or to influence a vote in an election. CAL. BUS. & PROF. CODE § 17941.

14. Jeffrey Dastin, “Amazon scraps secret AI recruiting tool that showed bias against women,” *Reuters* (October 9, 2018), <https://uk.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUKKCN1MK08G>.

15. Julia Angwin, Jeff Arson, et al., “Machine Bias,” *ProPublica* (May 23, 2016), <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>.

16. Bryce Goodman & Seth Flaxman, “European Union regulations on algorithmic decision-making and a ‘right to explanation,’” 4, 6, August 31, 2016, *arXiv.org*, <https://arxiv.org/pdf/1606.08813.pdf>.

17. Katie McNinnis, “The consumer welfare implications associated with the use of algorithmic decision tools, artificial intelligence, and predictive analytics,” *Consumers Union*, August 8, 2018, https://www.ftc.gov/system/files/documents/public_comments/2018/08/ftc-2018-0056-d-0031-155157.pdf.

18. See M.G.L. c. 186, §15.

19. GDPR, *supra* note 12, at Art. 82.