

Law

[The first edition of the law section](#) for *Ethically Aligned Design* noted that the early stages in development of autonomous and intelligent systems (A/IS) have given rise to many complex ethical problems that translate directly and indirectly into discrete legal challenges. That is, of course, what the rule of law often intends to answer – how we should behave as a society when faced with difficult ethical decisions – and it should come as no surprise that the legal implications of A/IS continue to unfold as we witness the forms of its expression and use expand.

To consider the ongoing creep of A/IS ethical issues into the legal realm, one need look no further than the first section of this document: *Legal Status*. This section addresses what legal status should A/IS be granted and was not a topic in the original edition. That is to say, in just one revision of this paper, we felt the need to address the question of how A/IS should be labeled in the courts' eyes: a product that can be bought and sold? A domesticated animal with more rights than a simple piece of property, but less than a human? A person? Something new?

Our conclusion to that question is that A/IS are not yet deserving of any kind of “personhood” – yet the very fact that the question of whether A/IS could, or should, be granted such status demonstrates the rate at which the technology and the related legal and ethical questions are growing and provide two universal principles echoed throughout this document:

The development, design, and distribution of A/IS should fully comply with all applicable international and domestic law.

There is much work to be done: the legal and academic community must increase engagement in this rapidly developing field from its members.

Law

Concerns and recommendations fall into four main areas:

1. [Legal Status of A/IS](#)
2. [Governmental Use of A/IS: Transparency and Individual Rights](#)
3. [Legal Accountability for Harm Caused by A/IS](#)
4. [Transparency, Accountability, and Verifiability in A/IS](#)

While much debate continues to surround A/IS, its development, and use, these questions must be addressed *before* the proliferation of A/IS passes some kind of tipping point.

The authors hope this paper will inform the legislative process and inspire more members of the legal community to become involved *now*.

Disclaimer: While we have provided recommendations in this document, it should be understood these do not represent a position or the views of IEEE but the informed opinions of Committee members providing insights designed to provide expert directional guidance regarding A/IS. In no event shall IEEE or IEEE-SA Industry Connections Activity Members be liable for any errors or omissions, direct or otherwise, however caused, arising in any way out of the use of this work, regardless of whether such damage was foreseeable.

Section 1 – Legal Status of A/IS

There has been much discussion about how to legally regulate A/IS-related technologies, and the appropriate legal treatment of systems that deploy these technologies. Lawmakers today are wrestling with the issue of what status to apply to A/IS. Legal “[personhood](#)” (as is applied to humans and certain types of human organizations) is one possible option for framing such legal treatment, and the implications of granting that status to A/IS applications raises issues that have implications in multiple domains of human interaction beyond technical issues.

Issue:
What type of legal status (or other legal analytical framework) is appropriate for application to A/IS, given the legal issues raised by deployment of such technologies?

Background

The convergence of A/IS and robotics technologies has led to the development of systems and devices with attributes that resemble those of human beings in terms of their autonomy, ability to perform intellectual tasks and, in the case of some robots, their physical

appearance. As some types of A/IS begin to display characteristics that resemble those of human actors, some governmental entities and private commentators have concluded that it is time to examine how legal regimes should categorize and treat various types of A/IS. [These entities have posited questions](#) such as, “Should the law treat such systems as legal ‘persons,’ with all the rights and responsibilities that personhood entails?” Such status seems initially remarkable until consideration is given to the long-standing legal personhood status granted to corporations, governmental entities, and the like – none of which are human even though they are run by humans.

Alternatively, many entities have asked, should some A/IS be treated as mere products and tools of their human developers and users? Perhaps A/IS are something entirely without precedent, raising the question of whether one or more types of A/IS might be assigned an intermediate – and perhaps novel – type of legal status?

Clarifying the legal status of A/IS in one or more jurisdictions is essential in removing the uncertainty associated with the obligations and expectations for organization and operation of these systems. Clarification along these lines will encourage more certain development and deployment of A/IS and will help clarify lines of legal responsibility and liability when A/IS cause harm. Recognizing A/IS as independent “legal persons” would, for example, limit or eliminate

Law

some human responsibility for subsequent “decisions” made by such A/IS (for example under a theory of “[intervening causation](#)” – akin to the “relief” from responsibility of a hammer manufacturer when a burglar uses a hammer to break the window of a house), thus potentially reducing the incentives for designers, developers, and users of A/IS to ensure their safety. In this example, legal issues that are applied in similar “[chain of causation](#)” settings (such as “[foreseeability](#),” “[complicity](#),” “[reasonable care](#),” “[strict liability](#)” for unreasonably dangerous goods, and other precedential notions) will factor into the design process. Different jurisdictions may reach different conclusions about the nature of such causation chains, inviting future creative legal planners to consider how and where to pursue design, development, and deployment of future A/IS in order to receive the most beneficial legal treatment.

The issue of the legal status of A/IS thus intertwines with broader legal questions regarding how to ensure accountability and assign and allocate liability when A/IS cause harm. The question of legal personhood for A/IS also interacts with broader ethical questions on the extent to which A/IS should be treated as moral agents independent from their human designers and operators, and whether recognition of A/IS personhood would enhance or detract from the purposes for which humans created the A/IS in the first place.

A/IS are at an early stage of development where it is premature to assert a single particular legal status or presumption for application in the many forms and settings in which those systems are

deployed. This uncertainty, coupled with the multiple legal jurisdictions in which A/IS are being deployed (each of which, as a sovereign, can regulate A/IS as it sees fit) suggests that there are multiple general frameworks through which to consider A/IS legal status. Below are some examples.

Candidate Recommendations

1. While conferring legal personhood on A/IS might bring some economic benefits, the technology has not yet developed to the point where it would be legally or morally appropriate to generally accord A/IS the rights and responsibilities inherent in the legal definition of personhood, as it is defined today. Therefore, even absent the consideration of any negative ramifications from personhood status, it would be unwise to accord such status to A/IS at this time. A/IS should therefore remain to be subject to the applicable regimes of property law.
2. Government and industry stakeholders alike should identify the types of decisions and operations that should never be delegated to A/IS, and adopt rules and standards that ensure effective human control over those decisions. Modern legal systems already address a number of other situations that could serve as appropriate analogues for the legal status of A/IS and how to allocate legal responsibility for harm caused by A/IS. These legal analogues may include the treatment of pets, livestock, wild animals, employees, and other “agents” of persons and corporations. Governments and courts should review

Law

these various potential legal models and assess whether they could serve as a proper basis for assigning and apportioning legal rights and responsibilities with respect to the deployment and use of A/IS.

3. In addition, governments should scrutinize existing laws — especially those governing business organizations — for mechanisms that could have the practical effect of allowing A/IS to have legal autonomy. If ambiguities or loopholes in the law could create a legal method for recognizing A/IS personhood, the government should review and, if appropriate, amend the pertinent laws.
4. Manufacturers and operators should gain an understanding of how each jurisdiction would categorize a given A/IS and how each jurisdiction would treat harm caused by the system. Manufacturers and operators should be required to comply with the applicable laws of all jurisdictions in which that system could operate. In addition, manufacturers and operators should be aware of standards of performance and measurement promulgated by standards development organization and agencies.
5. As A/IS become more sophisticated, governments should reassess the issue of legal status for these systems. In considering whether to accord legal protections, rights, and responsibilities to A/IS, governments should exercise utmost caution. Governments and decision-makers at every level must work closely with regulators, representatives of civil society, industry actors, and other stakeholders to

ensure that the interest of humanity — and not the interests of the autonomous systems themselves — remains the guiding principle.

Further Resources

- Bayern, S. "[The Implications of Modern Business-Entity Law for the Regulation of Autonomous Systems.](#)" *Stanford Technology Law Review* 19, no. 1 (2015): 93–112.
- Bayern, S. et al., "[Company Law and Autonomous Systems: A Blueprint for Lawyers, Entrepreneurs, and Regulators.](#)" *Hastings Science and Technology Law Journal* 9, no. 2 (2017): 135–162.
- Bhattacharyya, D. "[Being, River: The Law, the Person and the Unthinkable.](#)" *Humanities and Social Sciences Online*, April 26, 2017.
- Calverley, D. J. "[Android Science and Animal Rights, Does an Analogy Exist?](#)" *Connection Science* 18, no. 4 (2006): 403–417.
- Calverley, D. J. "[Imagining a Non-Biological Machine as a Legal Person.](#)" *AI & Society* 22 (2008): 403–417.
- European Parliament [Resolution of 16 February 2017](#) with recommendations to the Commission on Civil Law Rules on Robotics.
- Zyga, L. "[Incident of Drunk Man Kicking Humanoid Robot Raises Legal Questions.](#)" *Techxplore*, October 2, 2015.
- LoPucki, L. M. "[Algorithmic Entities.](#)" *Washington University Law Review* 95 (forthcoming 2017).

Law

- Scherer, M. "[Digital Analogues.](#)" *Imaginary Papers*, June 8, 2016.
- Scherer, M. "[Is Legal Personhood for AI Already Possible Under Current United States Laws?](#)" *Law and AI*, May 14, 2017.
- Solum, L. B. "[Legal Personhood for Artificial Intelligences.](#)" *North Carolina Law Review* 70, no. 4 (1992): 1231–1287.
- Weaver, J. F. [Robots Are People Too: How Siri, Google Car, and Artificial Intelligence Will Force Us to Change Our Laws.](#) Santa Barbara, CA: Praeger, 2013.
- European Parliament. European Parliament Resolution of 16 February 2017 with Recommendations to the Commission on Civil Law Rules on Robotics. February 16, 2017.

Law

Section 2 – Governmental Use of A/IS: Transparency and Individual Rights

Surveillance of populations by governments and the disruption of free elections will become ever easier as we deploy A/IS. How should we manage these systems to ensure that they act for the good of society?

Issue:

International, national, and local governments are using A/IS. How can we ensure the A/IS that governments employ do not infringe on citizens' rights?

Background

Government increasingly automates part or all of its decision-making. Law mandates transparency, participation, and accuracy in government decision-making. When government deprives individuals of fundamental rights, individuals are owed notice and a chance to be heard to contest those decisions. A key concern is how legal commitments of transparency, participation, and accuracy can be guaranteed when algorithmic-based A/IS make important decisions about individuals.

Candidate Recommendations

1. Government stakeholders should identify the types of decisions and operations that should never be delegated to A/IS, such as when to use lethal force, and adopt rules and standards that ensure effective human control over those decisions.
2. Governments should not employ A/IS that cannot provide an account of the law and facts essential to decisions or risk scores. The determination of, for example, fraud by a citizen should not be done by statistical analysis alone. Common sense in the A/IS and an ability to explain its logical reasoning must be required. Given the current abilities of A/IS, under no circumstances should court decisions be made by such systems alone. Parties, their lawyers, and courts must have reasonable access to all data and information generated and used by A/IS technologies employed by governments and other state authorities.
3. A/IS should be designed with transparency and accountability as primary objectives. The logic and rules embedded in the system must be available to overseers of systems, if possible. If, however, the system's logic or algorithm cannot be made available for

Law

inspection, then alternative ways must be available to uphold the values of transparency. Such systems should be subject to risk assessments and rigorous testing.

4. Individuals should be provided a forum to make a case for extenuating circumstances that the A/IS may not appreciate – in other words, a recourse to a human appeal. Policy should not be automated if it has not undergone formal or informal rulemaking procedures, such as interpretative rules and policy statements.
5. Automated systems should generate audit trails recording the facts and law supporting decisions and such systems should be amenable to third-party verification to show that the trails reflect what the system in fact did. Audit trails should include a comprehensive history of decisions made in a case, including the identity of individuals who recorded the facts and their assessment of those facts. Audit trails should detail the rules applied in every mini-decision made by the system. Providers of A/IS, or providers of solutions or services that substantially incorporate such systems, should make available statistically sound evaluation protocols through which they measure, quality assure, and substantiate their claims of performance, for example, relying where available on protocols and standards developed by the National Institute of Standards and Technology (NIST) or other standard-setting bodies.
6. Investor list(s), developers, and promoters of any given A/IS being developed should be required by law to be made public when the A/IS are used for governmental purposes. There should also be transparency of the specific ethical values promoted by the designer, and *how* they were embedded in the system. Transparency should also apply to the input data selection process.

Further Resources

- Schwartz, P. "[Data Processing and Government Administration: The Failure of the American Legal Response to the Computer.](#)" *Hastings Law Journal* 43 (1991): 1321–1389.
- Citron, D. K. "[Technological Due Process.](#)" *Washington University Law Review* 85 (2007): 1249–1313.
- Citron, D. K. "[Open Code Governance.](#)" *University of Chicago Legal Forum* 2008, no. 1 (2008): 355–387.
- Crawford, K., and J. Schultz. "[Big Data and Due Process: Toward a Framework to Address Predictive Privacy Harms.](#)" *Boston College Law Review* 55, no. 1 (2014): 93–128.
- Pasquale, F. [Black Box Society](#). Cambridge, MA: Harvard University Press, 2014.

Law

- Bamberger, K. "[Technologies of Compliance: Risk and Regulation in the Digital Age.](#)" *Texas Law Review* 88, no. 4 (2010): 669–739.
- Kroll, J. [Accountable Algorithms](#). Princeton, NJ: Princeton University Press, 2015.
- Desai, D., and J. A. Kroll. "[Trust But Verify: A Guide to Algorithms and the Law.](#)" *Harvard Journal of Law and Technology*, forthcoming.
- ICRC. "[Views of International Committee of Red Cross \(ICRC\) on Autonomous Weapon System.](#)" April 11, 2016.
- Rainie, L., J. Anderson, and J. Albright. "[The Future of Free Speech, Trolls, Anonymity and Fake News Online.](#)" Pew Research Center, March 29, 2017.
- Marwick, A., "[Are There Limits to Online Free Speech?](#)" *Data & Society: Points*, January 5, 2017.
- Neier, A., "[Talking Trash: What's More Important, Human Dignity or Freedom of Speech?](#)" *Columbia Journalism Review*, September/October 2012.

Law

Section 3 – Legal Accountability for Harm Caused by A/IS

As A/IS becomes more prevalent while also potentially becoming more removed from the human developer/manufacturer, what is the correct approach to ensure legal accountability for harms caused by A/IS?

Issue:

How can A/IS be designed to guarantee legal accountability for harms caused by these systems?

Background

One of the fundamental assumptions most laws and regulations rely on is that human beings are the ultimate decision-makers. As autonomous devices and A/IS become more sophisticated and ubiquitous, that will increasingly be less true. The A/IS industry legal counsel should work with legal experts to identify the regulations and laws that will not function properly when the “decision-maker” is a machine and not a person.

Candidate Recommendations

Any or all of the following can be chosen. The intent here is to provide as many options as possible for a way forward for this principle.

1. Designers should consider adopting an identity tag standard – that is, no A/IS agent should be released without an identity tag to maintain a clear line of legal accountability.
2. Lawmakers and enforcers need to ensure that the implementation of A/IS is not abused by businesses and entities employing the A/IS to avoid liability or payment of damages. Governments should consider adopting regulations requiring insurance or other guarantees of financial responsibility so that victims can recover damages for harm that A/IS cause.
3. Companies that use and manufacture A/IS should be required to establish written policies governing how the A/IS should be used, including the real-world applications for such AI, any preconditions for its effective use, who is qualified to use it, what training is required for operators, how to measure the performance of the A/IS, and what operators and other people can expect from the A/IS. This will help to give the human operators and beneficiaries an accurate idea of what to expect from the A/IS, while also protecting the companies that make the A/IS from future litigation.

Law

4. Because the person who activates the A/IS will not always be the person who manages or oversees the A/IS while it operates, states should avoid adopting universal rules that assign legal responsibility and liability to the person who “turns on” the A/IS. For example, liability may attach to the manufacturers or to the person who directs, monitors, and controls the A/IS’s operations, or has the responsibility to do so.
6. For the avoidance of repeated or future harm, companies that use and manufacture A/IS should consider the importance of continued algorithm maintenance. Maintenance is an essential aspect of design. Design does not end with deployment. Thus, there should be a clear legal requirement of (1) due diligence, and (2) sufficient investment in algorithm maintenance on the part of companies that use and manufacture A/IS that includes monitoring of outcomes, complaint mechanism, inspection, correction, and replacement of harm-inducing algorithm, if warranted. Companies should be prohibited from contractually delegating this responsibility to unsophisticated end-users.
7. Promote international harmonization of national legislations related to liability in the context of A/IS design and operation (through bi- or multilateral agreements) to enhance interoperability, and facilitate transnational dispute resolution.
8. Courts weighing A/IS litigation cases based on some form of injury should adopt a similar scheme to that of [product liability litigation](#), wherein companies are not penalized or held

responsible for installing post-harm fixes on their products designed to make the product safer. In other words, because courts have recognized that it is good public policy to encourage companies to fix dangerous design flaws, retroactively fixing a design flaw that has caused injury is not considered an admission or a sign of culpability. The same approach should be used in A/IS litigation.

Further Resources

- Allan, T., and R. Widdison. [“Can Computers Make Contracts?”](#) *Harvard Journal of Law and Technology* 9 (1996): 25–52.
- Asaro, P. M. [“The Liability Problem for Autonomous Artificial Agents.”](#) Palo Alto, CA: Association for the Advancement of Artificial Intelligence, 2015.
- Chopra, S., and L. F. White. [A Legal Theory for Autonomous Artificial Agents](#). Ann Arbor, MI: University of Michigan Press, 2011.
- Colonna, K. [“Autonomous Cars and Tort Liability.”](#) *Case Western Journal of Law, Technology & The Internet* 4 no. 4 (2012): 81–130.
- Field, C. [“South Korean Robot Ethics Charter 2012.”](#) PhD thesis (part), Sydney, Aus.: University of Technology, 2010.
- Grossman, M. R., and G. V. Cormack. [“Technology-Assisted Review in E-Discovery Can Be More Effective and More Efficient Than Exhaustive Manual Review.”](#) *Richmond Journal of Law and Technology* 17, no. 3 (2011): 1–48.

Law

- Kalra, N., J. Anderson, and M. Wachs. "[Liability and Regulation of Autonomous Vehicle Technologies.](#)" State of California Department of Transportation Technical Report. Berkeley, CA: Institute of Transportation Studies, University of California, 2009.
- Krakow, C. E. A. "[Liability for Distributed Artificial Intelligences.](#)" *Berkeley Technology Law Journal* 11, no. 1 (1996): 147–204.
- Rivard, M. D. "Toward a General Theory of Constitutional Personhood: A Theory of Constitutional Personhood for Transgenic Humanoid Species." *UCLA Law Review* 39, no. 5 (1992): 1425–1510.
- Scherer, M., "[Regulating Artificial Intelligence Systems: Risks, Challenges, Competencies, and Strategies.](#)" *Harvard Journal of Law and Technology* 29, no. 2 (2016): 353–400.
- Tobin, R., and E. Schoeman. "[The New Zealand Accident Compensation Scheme: The Statutory Bar and the Conflict of Laws.](#)" *The American Journal of Comparative Law* 53, no. 2 (2005): 493–514.
- Wachter, S., B. Mittelstadt, and L. Floridi. "[Transparent, Explainable, and Accountable AI for Robotics.](#)" *Science Robotics* 2, no. 6 (May 31, 2017). DOI: 10.1126/scirobotics.aan6080.
- Weaver, J. F. [Robots Are People Too: How Siri, Google Car, and Artificial Intelligence Will Force Us to Change Our Laws.](#) Santa Barbara, CA: Praeger, 2013.
- Weiss, A. "Validation of an Evaluation Framework for Human-Robot Interaction. The Impact of Usability, Social Acceptance, User Experience, and Societal Impact on Collaboration with Humanoid Robots." PhD thesis, University of Salzburg, 2010.
- Wooldridge, M., and N. R. Jennings. "[Intelligent Agents: Theory and Practice.](#)" *The Knowledge Engineering Review* no. 2 (1995): 115–152.

Section 4 – Transparency, Accountability, and Verifiability in A/IS

Transparency around A/IS is a difficult issue because it impinges on the differing needs of developers for trade secrets and users to be able to understand the technology to guard against problems occurring with it, and to hold accountable the correct entity in the event of a system failure.

Issue:
How can we improve the accountability and verifiability in autonomous and intelligent systems?

Background

Decision-making algorithms can be designed for various purposes, and the applications are wide-ranging for both the public and the private sectors. We must assume that virtually every decision that we make as humans can be mediated or replaced by an algorithm. Therefore, we cannot overestimate both the current and future role of A/IS across different sectors. Algorithms and automated decision-making (e.g., resume/cv screening during job applications)

have the potential to be more fair, and less biased than humans, provided that the systems are designed well. This requires, in particular, that effective preventative measures are put in place to avoid an algorithm-based information and/or value bias.

At the same time, most users of A/IS will not be aware of the sources, scale, varying levels of accuracy, intended purposes, and significance of uncertainty in the operations of A/IS, or that they are interacting with A/IS in the first place. The sources of data used to perform these tasks are also often unclear. Furthermore, users might not foresee the inferences that can be made about them or the consequences when A/IS are used. The proliferation of A/IS will result in an increase in the number of systems that rely on machine learning and other developmental systems whose actions are not pre-programmed, and that may not produce logs or a record of how the system reached its current state.

These systems are often opaque (frequently referred to as “black boxes”) and create difficulties for everyone, from the engineer, to the lawyer in court, to the online shopper, to the social media user. The result is an abundance of ethical issues of ultimate accountability.

Law

Candidate Recommendations

1. Given that many of the desired design specifications regarding accountability and verifiability are not technologically possible at this time, for now, this is an ethical issue that is best addressed by disclosure. If users are aware that they are interacting with an A/IS in the first place, and know exactly what information is being transferred to it, they will be better suited to tailor their inputs. 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 A/IS.
2. Designers and manufacturers must remain accountable for the risks or externalities their systems cause. This is a balancing act since the level of risk that is acceptably mitigated through disclosure is not always clear. Recommending specific levels (whether a manufacturer of A/IS acts responsibly, or whether there is enough disclosure, or whether total disclosure would even be enough to mitigate the risk to users) is often a fact-specific discussion that doesn't suit itself well to broad rules.
3. There is a demand for algorithmic operation transparency. Although it is acknowledged this cannot be done currently, A/IS should be designed so that they always are able, when asked, to show the registered process which led to their actions to their human user, identify to the extent possible sources of uncertainty, and state any assumptions relied upon.
4. A/IS should be programmed so that, under certain high risk situations where human decision-making is involved, they proactively inform users of uncertainty even when not asked.
5. With any significant potential risk of economic or physical harm, designers should conspicuously and adequately warn users of the risk and provide a greater scope of proactive disclosure to the user. Designers should remain mindful that some risks cannot be adequately warned against and should be avoided entirely.
6. To reduce the risk of A/IS that are unreasonably dangerous or that violate the law from being marketed and produced, we recommend lawmakers provide whistleblower incentives and protections. As in many industries, insiders may often be the first to know that the A/IS are acting illegally or dangerously. A well-crafted law to protect whistleblowers and allow a public interest cause of action would improve accountability and aid in prevention of intentional, reckless, or negligent misuses of A/IS.
7. Government and industry groups should consider establishing standards that require A/IS to create logs (or other means of verification of their decision-making process) regarding key aspects of their operations and store those logs for a specified period of time. Designers should leverage current computer science regarding accountability and verifiability for code. New verification techniques may need to be developed

Law

to overcome the technical challenges in verifiability and auditability of A/IS operations; A/IS oversight systems (“A/IS guardians”) or methods such as [Quantitative Input Influence](#) (“QII”) measures could facilitate this process. Making sure, *ex ante*, that such information is, or can be made, available will also provide a higher degree of trust in verifiability and a sense of transparency in A/IS operations.

8. In Europe, the discussion on the so called [“right to explanation”](#) when automated decision-making occurs is important to address. Even though it is not yet guaranteed in Europe, future jurisprudence or Member State laws could grant individuals the right to ask for an explanation when a solely automated decision (e.g., refusal of an online credit application or e-recruiting practices) is being made about them that has legal or other significant effects. Such a right could provide a mechanism to increase the transparency and accountability of A/IS, and should therefore be seriously considered. In addition, other accountability enhancing tools such as ethical audits or certification schemes for algorithms should be explored. In addition, users should have the right to be informed, possibly through an interactive training program, on the areas of uncertainty, risks, and circumstances where safety or harm issues could arise, without this increasing user’s accountability for A/IS decision-making consequences.
9. Lawmakers on national and international levels should be encouraged to consider and carefully review a potential need to introduce new regulation where appropriate, including rules subjecting the market launch of new A/IS driven technology to prior testing and approval by appropriate national and/or international agencies. Companies should establish an A/IS ethics statement that includes statements about discrimination, addressing in that matter data-driven profiling and commitment to take measures to avoid user discrimination. In addition, companies should have internal systems that allow employees to identify and escalate issues related to discrimination in data and A/IS. Laws should create whistleblower protection for those who can and wish to reveal explicit violation of discrimination law. In particular, a well-crafted law to protect whistleblowers and to allow a public interest cause of action would improve accountability and aid in prevention of intentional misuse of A/IS.
10. The general public should be informed when articles/press releases related to political figures or issues are posted by an A/IS, such as a bot.

Further Resources

- Barocas, S., and A. D. Selbst. [“Big Data’s Disparate Impact.”](#) *California Law Review* 104 (2016): 671–732.
- Datta, A., S. Sen, and Y. Zick. [“Algorithmic Transparency via Quantitative Input Influence: Theory and Experiments with Learning Systems.”](#) *2016 IEEE Symposium on Security and Privacy*, May 22–26, 2016. DOI: 10.1109/SP.2016.42

Law

- Etzioni, A., O. Etzioni. "Designing AI Systems That Obey Our Laws and Values." *Communications of the ACM* 59, no. 9 (2016): 29–31.
- Kroll, J. A., and J. Huey, S. Barocas, E. W. Felten, J. R. Reidenberg, D. G. Robinson, and H. Yu. "[Accountable Algorithms](#)." (March 2, 2016). *University of Pennsylvania Law Review* 165 (2017 Forthcoming); *Fordham Law Legal Studies Research Paper No. 2765268*.
- Mittelstadt, B., P. Allo, M. Taddeo, S. Wachter, and L. Floridi. "[The Ethics of Algorithms: Mapping the Debate](#)." *Big Data & Society* (July–December, 2016): 1–21.
- Regulation (EU) 2016/679 of the European Parliament and of the Council, General Data Protection Regulation (). "[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](#)." April 27, 2016.
- Wachter, S., B. Mittelstadt, and L. Floridi. "[Why a Right to Explanation of Automated Decision-Making Does Not Exist in the General Data Protection Regulation](#)." *International Data Privacy Law* 7, no. 2 (2017): 76–99.
- Zarsky, T. "[The Trouble with Algorithmic Decisions: an Analytic Roadmap to Examine Efficiency and Fairness in Automated and Opaque Decision Making](#)." *Science, Technology & Human Values* 41, no. 1 (2016): 118–132.