



INTRODUCING A PRACTICE-BASED COMPLIANCE FRAMEWORK FOR ADDRESSING NEW REGULATORY CHALLENGES IN THE AI FIELD



BY
MONA SLOANE



&
EMANUEL MOSS

Mona Sloane, PhD. is a Senior Research Scientist at the NYU Center for Responsible AI, Faculty at the NYU Tandon School of Engineering, a Postdoctoral Researcher at the Tübingen AI Center, a Director at the NYU Tisch School of the Arts, and a Fellow with NYU Institute for Public Knowledge and The GovLab. Emanuel Moss, PhD. is a Joint Postdoctoral Fellow at Cornell Tech and the Data & Society Research Institute.

AI Ethics, Regulation & Firm Implications

By Benjamin Cedric Larsen & Yong Suk Lee



Regulation of Artificial Intelligence - Global Trends, Implications, and the Road Ahead

By Jayant Narayan



Toward a Non-Dispositive, Human-First Agenda for Public Sector AI

By Jerry Ma



Introducing a Practice-Based Compliance Framework (PCF) for Addressing New Regulatory Challenges in the AI Field

By Mona Sloane & Emanuel Moss



Algorithmic Pricing - A Black Box for Antitrust Analysis

By Max Huffman & Dr. Maria José Schmidt-Kessen



Reflections on the EU's AI Act and How we Could Make it Even Better

By Meeri Haataja & Joanna J. Bryson



Towards a Liability Framework for AI in Europe

By Miriam Buiten & Jennifer Pullen



The EU AI Act - Balancing Human Rights and Innovation Through Regulatory Sandboxes and Standardization

By Katerina Yordanova



Visit www.competitionpolicyinternational.com for access to these articles and more!

Introducing a Practice-Based Compliance Framework (PCF) for Addressing New Regulatory Challenges in the AI Field

By Mona Sloane & Emanuel Moss

Over the past years, regulatory pressure on tech companies to identify and mitigate the adverse impact of AI systems has been steadily growing. In 2022, we can expect this pressure to grow even further with transnational, national, federal, and local AI regulation kicking in. Many of these regulatory frameworks target both the design and the use of AI systems, often with a sector focus. AI practitioners and regulators alike are in need of new approaches that allow them to effectively respond to these regulations, and to enforce them competently. In this contribution, we will map out a **Practice-Based Compliance Framework ("PCF")** for identifying existing principles and practices that are already aligned with regulatory goals, that therefore can serve as anchor points for compliance and enforcement initiatives.

Scan to Stay Connected!

Scan here to subscribe to CPI's **FREE** daily newsletter.



01

INTRODUCTION

Over the past years, regulatory pressure on tech companies to identify and mitigate the harm AI systems can cause has been steadily growing. Facial recognition leading to wrongful arrest,² cover-ups of research³ into the psychological toll social media inflicts on teenagers, wildly disparate error rates⁴ from AI products for members of different racial groups, and a seemingly endless succession of privacy breaches⁵ have ensured this pressure is well-earned. In 2022, we can expect this pressure to grow even further with transnational, national, federal, and local AI regulation being proposed at an accelerating pace.⁶

These regulations will vary — some will ban specific uses of AI technology, some will establish guidelines for what companies are expected to do or not do when building AI products, still others will require companies to take specific steps to document the intended uses of their products or assess their likely impacts on society and the environment. Increasingly, regulations are more likely to enact sector-specific regulations that place different requirements on different kinds of companies, and different kinds of products, depending on what their intended uses are. While the exact details of any new regulations are hard to foresee, it is abundantly clear that regulations are coming.

AI practitioners and regulators alike need new approaches that allow them to effectively respond to — and even anticipate — these regulations. With past regulations, a wait-and-see approach has had significant opportunity costs; many firms found themselves flat-footed when the EU General Data Protection Regulation (“GDPR”) was rolled out and had to rapidly revise long-standing data management practices to quickly come into compliance. Data management was not new to such firms. It was key to their business practices, but was not necessarily part of their compliance strategy.

But while the intentions of GDPR were clearly telegraphed by policymakers years before its enactment, these firms missed an opportunity to shift their data management practices to better align with the likely goals of GDPR, and had to drastically reshape both their compliance and data management teams on a short time frame. Today, with a new regulatory landscape clearly on the horizon, as we discuss below, steps can be taken now to anticipate regulatory changes and adapt to their requirements competently. In this article, we will map out a Practice-Based Compliance Framework (“PCF”) for identifying existing principles and practices that already align with regulatory goals, that therefore can serve as anchor points for compliance and enforcement initiatives.

02

NEW REGULATORY LANDSCAPES

The regulatory landscape for data-driven digital technologies is rapidly changing, following a lengthy period where it received little attention from lawmakers. From 1996, when the U.S. Congress updated the Communications Decency Act to protect common carriers from the content of their users' messages,⁷ to 2016, when the EU GDPR went into effect, little was done to address the many ways the technology industry has been reshaping society. As the first significant data regulation of the so-called age of “big data,” GDPR required sweeping changes to how “data controllers” — anyone who collects data — gain consent for collecting and using individuals' data, what they can do with that data once they have it, and what kinds of fines they face if they fail to comply. These changes rapidly upended how companies who collect and use data work; to demonstrate that they were in compliance with GDPR, they had to re-engineer database systems, redesign websites (including adding the now-familiar cookie consent popups we all know and love), and massively overhaul any machine learning services that used data covered by GDPR.

2 <https://www.nytimes.com/2020/12/29/technology/facial-recognition-misidentify-jail.html>, accessed on February 13, 2022.

3 <https://www.theguardian.com/technology/2021/sep/14/facebook-aware-instagram-harmful-effect-teenage-girls-leak-reveals>, accessed on February 13, 2022.

4 <https://www.newscientist.com/article/2166207-discriminating-algorithms-5-times-ai-showed-prejudice/>, accessed on February 13, 2022.

5 <https://www.reuters.com/technology/france-says-facial-recognition-company-clearview-breached-privacy-law-2021-12-16/>, accessed on February 13, 2022.

6 See especially the EU AI Act and the U.S. Algorithmic Accountability Act.

7 See <https://www.eff.org/issues/cda230>, accessed on February 13, 2022.

Since the enactment of GDPR, other more highly-specified regulations have been enacted (e.g. the California Consumer Privacy Act⁸ and Illinois' Biometric Information Privacy Act).⁹ But momentum is also building for a slate of subsequent regulations that have been drafted and that are sorely needed to protect the public and ensure data-driven technologies serve the public interest. In the United States, the Algorithmic Accountability Act, which stalled in 2019 but has just been reintroduced in Congress,¹⁰ would require developers to conduct impact assessments documenting how their products affect society and to involve community stakeholders in helping determine what potential impacts are assessed. In the European Union, the Artificial Intelligence Act¹¹ outlines what uses of AI ought to be considered risky in specific sectors, and would require that companies conduct "conformity assessments" to document the ways that the products they build are managing the appropriate degree of risk for its intended use case. What is common to these legislative proposals, and is likely to feature in any laws enacted in this current wave of AI regulation, is the need for companies to produce significant amounts of documentation about what they do and how it affects the public. What this means for companies, is that they will need to develop practices for complying with such requirements in ways that do not require starting from "square one" or reinventing their entire corporate management and compliance infrastructure, a need that the PCF described below addresses.

03

A PATHWAY FOR IMPLEMENTING NEW COMPLIANCE MANDATES

As we just discussed above, the regulatory landscape of AI within and across national borders is still in formation. As such, it is characterized by uncertainty. This uncertainty

affects regulators, technology companies, and civil society alike: the lines are blurry, and it is unclear how to best comply with and enforce new rules. PCF addresses this issue through a method that allows actors to comply with AI regulation rapidly and holistically by building on already existing organizational structures and baking AI compliance into these existing structures, practices and cultures — rather than deploying it top-down.

We propose that for developing such a method, a social science lens is extremely important. Specifically, we argue that social practice theory provides a particularly useful frame for considering strategies for encouraging behavioral change and social processes that do not depend on *linear* models of intervention implementation.¹² Social practice theory, the core of PCF, deploys a dynamic framework in which the central unit of inquiry — a social practice — comprises the three elements: meanings, competences, and materials. Meanings designate symbolic meanings, collective and emotional knowledge, shared aspirations, and social norms; competencies are skills, knowhow, techniques, and practical knowledges; and materials include tools, infrastructures, hardware, and other tangible entities, including the body itself.¹³

When these three elements combine in individual practices that continue to be reproduced, they stabilize the unit of a social practice. Broad examples of a social practice are cooking, driving, or exercising. More nuanced ones are shopping sustainably, keeping cool indoors, or doing AI design under full compliance with new AI regulation. The links between elements are made, broken, and re-made through individual reproduction. This process can transform elements. For example, the meaning of cooking can change under the observance of a new diet. Or the competence of AI design shifts based on new hardware that becomes available, or based on new (regulatory) requirements that are introduced into the practice. Elements can also disintegrate. For example, the meaning of computational work as secretarial and therefore feminized work disintegrated from the late 1960s. Computing jobs moved from being seen as so unskillful and unimportant that it was seen as inappropriate for men to take

8 <https://oag.ca.gov/privacy/ccpa>, accessed on February 13, 2022.

9 <https://www.ilga.gov/legislation/ilcs/ilcs3.asp?ActID=3004&ChapterID=57>, accessed on February 13, 2022.

10 <https://www.wyden.senate.gov/news/press-releases/wyden-booker-and-clarke-introduce-algorithmic-accountability-act-of-2022-to-require-new-transparency-and-accountability-for-automated-decision-systems>, accessed on February 13, 2022.

11 <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52021PC0206>, accessed on February 13, 2022.

12 Frost, J., Wingham, J., Britten, N. *et al.* The value of social practice theory for implementation science: learning from a theory-based mixed methods process evaluation of a randomised controlled trial. *BMC Med Res Methodol* 20, 181 (2020). <https://doi.org/10.1186/s12874-020-01060-5>.

13 Shove, E., Pantzar, M., & Watson, M. 2012. *The dynamics of social practice: Everyday life and how it changes*. Sage.

them, to becoming more synonymous with management, and thus masculinity, high status, and power — a meaning that forcefully stabilizes the social practice of computer work to this day.¹⁴

What follows from that is that the significance, purpose, and skill of a given practice is not contained to individual bodies or minds of people. Rather, people are “carriers of practice.” Relationships between practices and practitioners differ. Some are devoted practitioners (for example, stamp collectors) who keep practices alive, regardless of the status of a practices’ “career” (considering stamp collecting as a social practice that has been in a steady state of disintegration for a few decades). Others are reluctant practitioners, for example those who bought an expensive indoor exercise bike to motivate themselves to exercise more despite preferring to exercise by walking in the park.

Crucially important, however, is that policy can configure and reconfigure the elements of a social practice: subsidies can change availabilities of materials (for example computer chips), regulation can change the meaning of a practice (for example privacy in web surfing), and educational investments can change the competencies that are required for the participation in a practice (for example STEM degrees).

The point here is to underscore how a social practice theory approach can help to both identify systemic failure of interventions that sought to change behavior¹⁵ and serve as a basis for practitioners to identify the elements of practice (i.e. existing processes within and beyond their organization). Just as importantly, social practice theory can help identify high-potential “carriers of practice” and to specify how and where to implement concrete compliance processes - without them being based on linear, “top-down” implementation models. Below, we demonstrate how PCF accomplishes this.

04

PCF: HOW TO DO IT

PCF is a way of adapting existing social practices within a company to new regulatory goals without completely disrupting established ways of working. To do so requires analyzing new regulation and identifying the work practices that

are likely to be affected. Concurrently, work practices can be analyzed to identify the meanings, competences, and materials that can be maintained in shifting toward compliance with new regulations, and which ones ought to be altered.

PCF gives practitioners a three-step strategy to analyze the macro-level and micro-level of a new regulation and its impact on an organization, and to consider how a social practice theory approach can be leveraged to rapidly develop non-linear compliance processes. Practitioners should compose responses to the following catalog of questions:

Macro-Level: Regulation Analysis

- What is the regulation?
- Who is the authority, and what is the territory?
- What technology does it target and how is the technology defined?
- What are the interventions mandated by the regulation?
- What intervention should be in focus? [Out of the above list, pick one concrete intervention before you proceed with answering the rest of the questions in the catalog, then repeat for subsequent interventions]
- What behavioral change on an organizational level is required to comply with that intervention?

Micro-Level: Social Practice Analysis

- Within an organization, what are the existing social practices affected by the mandated intervention?
- What are the elements of that social practice (i.e. meanings, competencies, and materials)?
- Who are the carriers of that practice?

Synthesis

- How does one or multiple elements of the social practice have to change in order to achieve the behavioral change?
- What are existing (organizational) processes that can be

¹⁴ Hicks, M., 2017. *Programmed inequality: How Britain discarded women technologists and lost its edge in computing*. MIT Press.

¹⁵ Frost et al. 2020.

leveraged to achieve the desired change on the level of the elements?

- Who are the high-potential carriers of practice who can spearhead this recalibration of the social practice?

We use a concrete example to illustrate the application of this process: the New York City Council bill on automated employment decision tools (Int 1894)¹⁶ which passed on November 10, 2021. This bill requires that “a bias audit be conducted on an automated employment decision tool prior to the use of said tool” and that “candidates or employees that reside in the city be notified about the use of such tools in the assessment or evaluation for hire or promotion, as well as, be notified about the job qualifications and characteristics that will be used by the automated employment decision tool,” with violations being subject to a civil penalty.

If we adopt the identity of an affected organization, such as a vendor of hiring AI, and use the above three-step strategy to effectively recalibrate and align social practices with regulatory goals, the following responses are possible:

Macro-Level: Regulation Analysis

- **What is the regulation?** The New York City Council bill on automated employment decision tools (Int 1894).¹⁷
- **Who is the authority, and what is the territory?** The authority is the New York City Council, and the territory is New York City.
- **What technology does it target and how is the technology defined?** The technology targeted is “automated decision tools.” In the bill, this technology is defined as “any system whose function is governed by statistical theory, or systems whose parameters are defined by such systems, including inferential methodologies, linear regression, neural networks, decision trees, random forests, and other learning algorithms, which automatically filters candidates or prospective candidates for hire or for any term, condition or privilege of employment in a way that establishes a preferred candidate or candidates.”
- **What are the interventions mandated by the regulation?** The interventions mandated by the regulation are bias audits, notification mechanisms for candidates and employees, and disclosure mechanisms

about the qualifications and characteristics used by the tool.

- **What intervention should be in focus?** The mandated intervention in focus here should be the mandated disclosure of the qualifications and characteristics used by the tool.
- **What behavioral change on an organizational level is required to comply with that intervention?** The behavioral change that is required on an organizational level is to make designing disclosure mechanisms a meaningful component of AI design practice.

Micro-Level: Social Practice Analysis

- **Within an organization, which existing social practice is most relevant to the intervention mandated by the regulation?** The existing social practice most relevant to the intervention mandated by the regulation is the AI design of a hiring tool, which here can be seen as a combination of machine learning engineering and user interface design applied to the hiring domain.
- **What are the elements of that social practice? (i.e. meanings, competencies, and materials)?** The materials of the social practice of AI design of a hiring tool are computer hardware, training data (e.g. qualifications and other characteristics of job candidates who have historically excelled in a job role, including characteristics that may not be directly or even indirectly relevant to evaluating a job candidate), a statistical model, input data / information (e.g. qualifications, characteristics, and other data solicited from individual job applicants), a hosting server, the web interface that connects the model to clients and users, and the devices used by clients and users to access that interface.

“The technology targeted is “automated decision tools.”

The *competencies* of AI design (stipulating for this case that it is a combination of machine learning and user interface design) is applied data science (i.e. being able to use applied statistical techniques to predict successful job appli-

¹⁶ <https://legistar.council.nyc.gov/LegislationDetail.aspx?ID=4344524&GUID=B051915D-A9AC-451E-81F8-6596032FA3F9>, accessed on February 13, 2022.

¹⁷ *Ibid.*

cants based on their qualifications and characteristics) and being able to design access and meaningful interaction with that model for multiple agents (clients, users).

The *meanings* of AI design are the informational content of data drawn from and supplied to clients and users (i.e. not the number of years of experience a candidate has and that might be entered into a data table, but rather what those years of experience mean for being able to succeed on-the-job), the classifications (or rankings or predictions) that are applied by the AI system to users and provided to clients (e.g. degree of suitability for a particular job), and what makes for a good (e.g. accurate, fair, robust) AI model.

Who are the carriers of that practice? The carriers of that practice are members of the engineering team at the organization, specifically those focused on model development and user interface design (rather than, for example, the marketing team).

Synthesis

- **How does one or multiple elements of the social practice have to change in order to achieve the behavioral change?** The material element of the social practice of AI design of a hiring tool must change to include a piece of text disclosing the qualifications and characteristics used by the tool. To make that material change, however, requires resolving a challenging question for AI design. Namely, *which* qualifications and characteristics, of the many characteristics an AI designer might have access to, are relevant to predicting a successful job applicant? The competencies of AI design do not necessarily already include that degree of precision, as effective tools for predicting and classifying job applicants can be built without knowing which specific characteristics contributed to the overall accuracy of an AI model. To comply with the disclosure requirements, however, requires changing this competency of AI design of hiring tools.
- **What are existing (organizational) processes that can be leveraged to achieve the desired change on the level of the elements?** AI design of hiring tools already has processes in place to test and evaluate models. This process can be adapted to include benchmarks that include metrics for evaluating the relevance of each qualification or characteristic to a model, as part of the overall evaluation of model performance. The old maxim "you can't manage what you can't measure" is apt here; metrics are already a key competency of AI design that can be modified here to shift the over-

all social practice toward being able to comply with this regulatory intervention.

- **Who are the high-potential carriers of practice who can spearhead this recalibration of the social practice?** The machine learning engineers who practice AI design for hiring tools are well-positioned to recalibrate the social practice toward being able to offer the disclosures mandated by the New York City bill. They hold the competencies in applied statistics, and can tackle the challenges involved with creating relevance measures for qualifications and characteristics. Framing this challenge as an exciting research problem (which it is) aligns with the incentives that give meaning to the work of these carriers of practice. These incentives are strengthened by the fact that addressing this problem could improve the entire field of AI and machine learning, and also would burnish the credentials and skills of those who work on it. But engineers cannot accomplish this alone; they must be supported by project managers (e.g. by allocating work hours to their engineering team for addressing this task) and by the user-interface designers who must hold visual space in the finished product's interface in which to place the disclosure.

05

CONCLUSION

In this article, we have mapped out the emerging regulatory landscape around AI and suggested a new Practice-based Compliance Framework ("PCF") that can help practitioners rapidly recalibrate their existing professional practice to comply with new regulatory mandates. PCF is based on a social practice theory approach that focuses on identifying the elements of a practice (i.e. existing processes within and beyond their organization) as well as high-potential "carriers of practice" to specify how and where to implement concrete compliance processes. We have argued that this approach can help avoid systematic failures that can be caused by top-down intervention models that are blind to how the relevant actors make sense of what they do.

We have argued that practitioners should use the three-step PCF to analyze the macro-level and micro-level of a new regulation and its impact on an organization in order to derive effective strategies for realizing desired behavioral change. To illustrate PCF, we have proposed a set of

questions pertaining to the regulation, the relevant social practice, and the synthesis of both. We have demonstrated the applicability of our approach by taking on the perspective of an AI vendor and walking the reader through the example of the New York City bill on automated hiring tools.

There are, of course, limitations to this approach. It could, for example, be argued that a social practice theory approach leads to a narrow engagement with a new regulation that is overly compliance-focused, distracting from more sweeping shifts in the culture of AI design and deployment that regulation might seek to encourage (such as user empowerment through mandates requiring that users have more power over what data is collected on them). It could also be argued that our interpretation of social practice theory is too focused on pushing behavioral change, rather than assessing failures of past attempts. In the same vein as both of these points, it could also be said that the proposed approach deliberately leaves larger issues around power, oppression, and capital, and how AI regulation can address such issues (for example in the realm of taxation), untouched.

However, we argue that PCF can help mitigate one of the most pressing issues the field of responsible AI is currently facing: a polarization between technologists and social scientists, and between regulators and industry. A focus on how the professional practice of AI stabilizes can direct attention onto how and where issues show up, and where what kinds of knowledges and tactics, as well as interdisciplinary collaborations, can be deployed to slowly, but steadily, shift a whole industry towards more accountability and equity. That, for sure, is a topic that is relevant beyond tech regulation. ■



*Processes in place to test and evaluate models.
This process can be adapted*

CPI SUBSCRIPTIONS

CPI reaches more than **35,000 readers** in over **150 countries** every day. Our online library houses over **23,000 papers**, articles and interviews.

Visit [competitionpolicyinternational.com](https://www.competitionpolicyinternational.com) today to see our available plans and join CPI's global community of antitrust experts.

