JOURNAL OF
CROSS-DISCIPLINARY RESEARCH
IN COMPUTATIONAL LAW [CRCL]

Constitutive Practices of Administrative Law as Limits on Legal Automation

Frank Pasquale*

Abstract

At present, the appeal of automated legal systems rests on three pillars: speed, scale, and preference satisfaction. However, for many aspects of legal systems, there is a common sense that their translation into computation would be inappropriate. This concern about premature or unwise automation has many facets. The flexibility of natural language as opposed to computer languages is critical. A 'legal process' account of the rule of law hinges on the availability of human review, appeals, and dialogic interaction.

One possible rejoinder to these specific accounts of the limits of legal automation, for advocates of technology, is to characterize their critics' treatment of extant legal processes as ossification or naturalization. Ossification refers to a pathological hardening into permanence of practices that are merely contingent. Naturalization denotes the treatment of human-made processes as something like the laws of nature, premised on an errant assumption of their lasting endurance or value.

The simultaneous malleability of legal systems, and prevalence of constitutive practices within them, leads to a two-level consideration of a) what aspects of a liberal legal order are crucial, and b) for those that are crucial, what is lost when that aspect is either partially or fully automated. Within a sphere of human activity like a liberal legal order, some patterns of action are merely instrumental to achieving ends, while others are essential, or constitutive: the activity should no longer even be considered part of a liberal legal order when the practice ceases.

Administrative processes that are simply incidental and instrumental to the legitimate resolution of a case are well primed for automation. Other practices conducted by persons, for persons, are essential and intrinsically important, and properly resist being converted into machine-readable code. Distinctions between incidental and constitutive, or instrumentally and intrinsically important, aspects of law, should both bound and guide legal automation.

Keywords: administrative law, artificial intelligence, due process, constitutive practice, tradition, interpretive social science, legal order, rule of law

Replier: Samuel Judson • Nexus, sam@sjudson.com.

Journal of Cross-disciplinary Research in Computational Law

DOI: pending

www.journalcrcl.org

* Professor of Law, Cornell Tech and Cornell Law School, fp269@cornell.edu.

> 'The surest way to "kill" philosophy is to transmit it in the manner of a science…To deal with Plato without first sharing his problem is somewhat analogous to studying butterflies on the basis of a collection of pinned butterflies, without ever having seen one fly.'
>
> — Isabelle Stengers, *Thinking with Whitehead: A Free and Wild Creation of Concepts*[1]

Introduction

Legal automation is advancing in both public and private law. Coders and lawyers are working together to promote automatic contracting, monitoring of existing legal agreements, and dispute resolution. Firms market chatbots to solve pressing problems of access to justice, providing users with forms for an array of simple claims. Venture capital in the 'legal tech' space is betting on numerous 'turnkey solutions' to discovery, venue selection, and even legal search and argumentation.

Many of these initiatives are commendable, promoting both access to justice and commercial efficiencies. However, as Paracelsus advised, the dose makes the poison – and it is all too easy to overdose on AI in law. For example, one legaltech innovator proposed to pay an advocate $1 million to use a legal chatbot to make an oral argument before the Supreme Court. Levity ensued.[2] But there is a lingering worry that other, less dramatic forms of legal automation will exacerbate the very problems they are aiming to solve.

This concern about premature or unwise automation has been articulated in many ways.[3] The flexibility of 'natural languages', as opposed to computer languages, is critical.[4] Lawmakers insecure about their ability to forecast all possible states of affaris to which the law may apply, are likely to find the flexibility of language, rather than computer code, the optimal way of expressing the commitments embodied in law. The identity and responsibility of decisionmakers matters, too. For Kiel Brennan-Marquez and Stephen Henderson, human empathy is critical to the legitimacy of penalties imposed by the justice system.[5] A 'legal process' account of the rule of law also hinges on the availability of human review, appeals and dialogic interaction.[6] These specific accounts of the limits of legal automation have rightfully influenced both scholars and policymakers.

One possible rejoinder here, for advocates of technology, is to reject their critics' treatment of extant legal processes as ossification or naturalisation. Ossification refers to a pathological hardening into permanence of practices that are merely contingent, with questionable (or negative) utility. Naturalisation here denotes the treatment of human-made processes (like the legal practices discussed in this article) as something like laws of nature, and thus errantly assuming their lasting endurance or value.[7] The writing of contracts and adjudication of contract disputes have evolved significantly over centuries and even over recent decades. Public administration is a constant target of 'revitalisation' via neoliberal managerialism.[8] The critics of critics of AI say we must not become nostalgic, clinging to older modes of dispute resolution or planning simply be-

[1] Isabelle Stengers, *Thinking with Whitehead: A Free and Wild Creation of Concepts* (Harbard University Press 2014).

[2] Jacob Silverman, 'Angry Users Want DoNotPay to Pay Up: Troubles Accumulate for the Robot Lawyer That's Not a Robot or a Lawyer' (2023) ⟨https://www.jacobsilverman.com/p/angry-users-want-donotpay-to-pay⟩: 'Applying Silicon Valley business principles to something as important and potentially definitive as "the law" presents numerous practical challenges, ethical quandaries, and, of course, legal concerns.'

[3] Brian Sheppard, 'Incomplete innovation and the premature disruption of legal services' [2015] Michigan State Law Review 1797.

[4] Mireille Hildebrandt, 'The adaptive nature of text-driven law' (2021) 1(1) Journal of Cross-disciplinary Research in Computational Law: 'In suggesting that specific technologies generate specific "normativities" I am putting forward that people will form specific types of habits depending on the technologies they employ, thus reinforcing or transforming the mutual expectations people have of each other.'

[5] Kiel Brennan Marquez and Stephen E Henderson, 'Artificial Intelligence and Role-Reversible Judgment' (2019) 109(2) The journal of criminal law and criminology 1. Some commentators have countered that critics of AI may set too high a bar generally for natural language understanding (NLU). See Magnus Sahlgren and Fredrik Carlsson, 'The singleton fallacy: Why current critiques of language models miss the point' (2021) 4 Frontiers in Artificial Intelligence 682578. Even if one concedes this point, law presents a particular area where a demanding conception of understanding is necessary, including components of empathy, shared experience, and mutual intersubjective understanding.

[6] Frank Pasquale, 'A rule of persons, not machines: the limits of legal automation' (2019) 87 George Washington Law Review 1.

[7] Mark Bevir and Jason Blakely, *Interpretive social science: An anti-naturalist approach* (Oxford University Press 2018).

[8] Julie E Cohen, *Between truth and power* (Oxford University Press 2019).

cause of tradition.[9] This is a powerful objection, not least because commentators and even law reform commissions are now considering how to make the law more tractable to automation.

To respond to the naturalisation and ossification challenges, this article aspires to provide a more generalised account of why law is, or should be, more difficult to automate than, say, driving, manufacturing standardized consumer goods, or maintaining a comfortable range of temperatures in a house. None of these more mundane activities have constitutive practices separate and apart from the ends they are supposed to achieve. Technology could utterly transform how they are accomplished now and, and while there may be enormous disruption on a personal and even community level, few if any would want to stop the technology itself from advancing (even if they might be justifiably inclined to fight for strong social protections in the wake of its adoption).[10] Few mourn the disappearance of elevator or switchboard operators.

More normatively freighted social phenomena, such as legal dispute resolution by administrative agencies of the state, are distinct from rote and routinizable actions, or services where there is a clear destination at the outset (like logistics and driving).[11] One cannot locate justice on a map, or prescribe some algorithmic set of actions whereby it is accomplished. While agency's processes can evolve over time, they cannot lose a critical mass of the practices that constitute them, lest they become something else altogether.[12] This is partly because, in much of law, the ends or outcomes to be achieved are unknown at the outset. The chief way that we know that the outcome of a certain procedure is legitimate (if it is within a certain range of possibilities that do not violate principles of substantive justice), is by knowing it was generated via just procedures.

Structure, process and outcome evaluation illuminate this point, particularly when law is contrasted with other professions. In healthcare, for example, a surgery center may be evaluated on the basis of its facilities (structure), the actions of its staff (process), or the health of its patients (outcomes).[13] There is widespread agreement on what health is, so healthcare structures and processes may change substantially over time, so long as they improve health overall. Law is different when it deals with situations that have no clear 'correct' outcome before a ruling (such as a contested litigation) or agreement (such as contractual negotiations). The very fact of a dispute over the application of particular finance or environmental regulations indicates some degree of social disagreement over the dispute's proper resolution. Thus, processes of truth-seeking and norm-applying take centre stage, rather than any particular outcome the law is presumed to achieve.

Setting aside frivolous litigation, we do not know, we do not know which side of a case to be litigated should prevail, for example, because otherwise it would have been resolved at some earlier stage of litigation. Perhaps basic measures

[9] Caution is in order here, since denigrations of a practice as 'traditional' can be as mindless as the conception of unreflective traditionalism they assume. See Immanuel Maurice Wallerstein, *Historical capitalism with capitalist civilization* (Verso 1995), identifying the designation of certain practices as mere tradition (in the sense of lacking rational foundation), as more ideological than historical; James Alexander, 'Three rival views of tradition (Arendt, Oakeshott and MacIntyre)' (2012) 6(1) Journal of the Philosophy of History 20: 'Oakeshott sees tradition as something which conditions all action, and MacIntyre sees tradition as something which should condition right action'.

[10] To be sure, enormous personal disruption can be caused by such transitions, and communities may be dispersed via technocapital's rearrangement of economic activity. One need only look at the social decline and upheaval prevalent in the 'rust belt' of the United States to understand these effects. However, the process of, say, car-making itself is not regarded as degraded or corrupted by the introduction of robots or other machinery to complete tasks once done by humans.

[11] This article addresses utilisation of automation by agencies of the state. With respect to private dispute resolution, matters may be more flexible – i.e., a thinner set of constitutive practices bound the range of legitimate dispute resolution structures. Nevertheless, it is likely that even in these scenarios, constitutive practices will limit the degree to which automation may be recognised as legitimate.

[12] Of course, they may be altered over time, and still recognisable as such – no one claims that a court that uses a word processor is less legitimate than one which writes its judgments on parchment with a quill. But if, for example, no record of a proceeding were kept (or, worse, even possible), that would menace a constitutive practice of liberal legal orders (actually maintaining some record of the charge or claim involved, key evidence presented, and the rationale of judgment). For a jurisprudential account of the importance of such publicity, see Lon Luvois Fuller, *The morality of law* (Yale University Press 1969).

[13] Avedis Donabedian, 'Evaluating the quality of medical care' (1966) 44(3) The Milbank memorial fund quarterly 166. This structure/process/outcome model is valuable in policy evaluation.

like 'days to resolution of cases' can serve as rough indicators of the quality of a legal system, but note that they are very rough: a court could resolve all cases with a coin flip most speedily and have no claim to have done justice to litigants' concerns. To identify a process as legal, one will have to look beyond what it is aiming to achieve, to something internal to the process itself: its constitutive practices.

My discussion of these constitutive practices will focus on US law, and particularly a zone of contestation familiar to many writers in the field of legal automation: administrative hearings. Experts in administrative law have seized on the promise of automation using natural language processing (NLP), artificial intelligence (AI) and machine learning (ML). In a landmark report published in 2019, Engstrom, Ho, Sharkey and Cuellar identified numerous opportunities for automation of several dimensions of the administrative state.[14] Despite the many worthy initiatives they highlight and commend, there will be enduring challenges to AI in administrative hearings. Thus, legal automation must proceed carefully in order to avoid eliding constitutive practices that are in part constituted by direct human participation. And sometimes it should not proceed at all.

The discussion below titled 'Constitutive practices of due process in administrative hearings', examines how classic dimensions of due process in US administrative hearings might be conceived as constitutive practices of law and therefore particularly resistant to elision or reduction via automation, or even the types of full specification needed to prepare for automation. The article concludes with reflections on how recognition of the irreducibility of particularly important human practices to algorithms may enrich future discussions of legal technology.

Constitutive Practices of Due Process in Administrative Hearings

The simultaneous malleability of legal systems and persistence of constitutive practices within them leads to a two-level consideration of a) when human presence and skill is crucial to the legitimacy of a step of a legal process and b) what is lost when this presence and skill is either partially or fully replaced with computation. This part will consider each of these levels in a series of reflections on the 'Friendly factors': 11 aspects of the administrative process in the US that Judge Henry Friendly thought important to consider in due process determinations.[15] The factors are listed below:

1. unbiased tribunal;
2. notice of proposed action and grounds asserted for it;
3. opportunity to present reasons why the action should not occur;
4. right to call witnesses;
5. right to know opposing evidence;
6. right to have decision based exclusively on evidence presented;
7. right to counsel;
8. making of a record;
9. availability of statement of reasons;
10. public attendance;
11. judicial review.

[14] David Freeman Engstrom and others, 'Government by algorithm: Artificial intelligence in federal administrative agencies' [2020] (20-54) NYU School of Law, Public Law Research Paper.

[15] Henry Friendly, 'Some Kind of Hearing' (1975) 1268 University of Pennsylvania Law Review. Friendly's article was published as the US Supreme Court was attempting to clarify (and in many cases limit) the 'due process revolution' unleashed by its decision in *Goldberg v. Kelly* (1970). The *Goldberg* decision held that a welfare benefits recipient has a right to a hearing before benefits were terminated, but was soon limited by a series of decisions limiting the scope and force of due process rights.

For organisational purposes, these may be usefully categorised as a) Aspects of the Hearing (1, 10), b) Rights of the Person Judged (2, 3, 4, 5, and 7), c) Obligations of the Decisionmaker (6, 8, 9), and d) Right to Review (11). By considering the plausibility of automating dimensions of each of these aspects of administrative hearings, I hope to illuminate what ought to be considered *constitutive* of a liberal legal order and what is merely contingent to particular times and places. Within a sphere of human activity like a liberal legal order, some patterns of action are merely instrumental to achieving ends, while others are essential or constitutive: the activity should no longer even be considered part of a liberal legal order when the practice ceases.[16]

I should note the limitations of this approach at the outset. To structure an argument on the basis of a single, almost half-century old article by a single judge primarily concerned with a single level (federal) of law in a single (and in many ways unusual) common law legal system (the US), may strike the reader as parochial and even antiquarian. Hasn't legal scholarship moved beyond the 1970s? Shouldn't a project ostensibly based on the foundations of liberal legal orders, consult both civil and common law systems—and diverse common law systems at that? And isn't it the case that, in the wake of *Mathews v Eldridge* and its prodigious jurisprudential progeny, many US administrative hearings lack one or more of the Friendly factors mentioned above? In its most exteme form, this critique might demand: isn't the true path of law toward a singularitarian integration of all values and data into a computational-legal infrastructure, rather than attention to the singularity of a particular judge's views on the dimensions of due process?[17]

To be sure, there is much in legal scholarship elaborating on the type and complexity of procedures characteristic of (and desirable for) liberal legal orders. However, Friendly's typology remains a touchstone for many students and scholars of administrative law and due process in the US. It is deployed here less as an authoritative representation of the nature of liberal legal orders generally, than as a useful organising framework for practitioners of computational law who are interested in the types of protections for litigants that are often guaranteed in the fair administrative practices that are characteristic of liberal legal orders. Later work may explore the types of additions and subtractions from the Friendly list that may be characteristic outside the US or indeed in jurisprudential trends in the US itself.

Aspects of the hearing: unbiased tribunal and public attendance

Several types of bias may compromise a litigant's fair opportunity before a decisionmaker. There is a risk of *associational bias* when a decision-maker may have commercial or social ties with one side of a case. However, in practice, both agencies and the courts that are supposed to police them have been reluctant to closely examine such biases. Criticised for helping to decide a case involving then-Vice President Dick Cheney, with whom he had gone duck hunting, Justice Scalia once opined that 'a rule that required Members of this Court to remove themselves from cases in which the official actions of friends were at issue would be utterly disabling'.[18] Thus pragmatism reigns. Given the social station of Justices and indeed many judges of inferior courts, it is likely that the 'company they keep' would be far more well-heeled and conservative than the typical US citizen.[19] But this type of bias is largely seen as too diffuse to require action with respect to particular judges.

A second dimension of bias is *substantive*, concerning decision-makers who have or appear to have decided cases before they are even presented to them. For example, a judge who had written to condemn a firm's actions may not be permitted to judge those actions in court. Matters

[16] Perhaps there are counter-examples of state action in liberal legal orders that dispenses with one or more of these practices, but context is critical: they may just as easily be evidence of growing lawlessness by the state, as they are evidence that the practice is dispensable and not in fact constitutive.

[17] By singularity here I wish to evoke Lucien Karpik's emphasis on the value of unique, incommensurable insights of particular attorneys. See Lucien Karpik, *Valuing the unique: The economics of singularities* (Princeton University Press 2010). See also David M Dorsen, *Henry Friendly, greatest judge of his era* (Harvard University Press 2012).

[18] *Cheney v United States* 124 S. Ct. 1391 (District Court for the District of Columbia 2004) (Scalia J, mem.).

[19] On the importance of such social networks to judicial ideology, see Lawrence Baum and Neal Devins, *The company they keep: How partisan divisions came to the Supreme Court* (Oxford University Press 2019).

become considerably more complicated with leadership of agencies, though: an entity like the Federal Trade Commission, for example, has quasi-executive, quasi-legislative and quasi-judicial functions. It is impossible to lead such an agency without having some idea about what problems should be its highest priorities; this exercise of prosecutorial discretion cannot itself be a form of impermissible bias. The FTC's antitrust mandate covers industries where only one or a few firms may be dominant. To talk about the social media industry in the US, for example, is to talk at least somewhat about Meta and its acquisitions. So merely speaking about or analysing a particular firm's actions cannot be disqualifying either, without more evidence and context.

When exercising quasi-judicial authority, an agency decision-maker is to be disqualified if he or she 'has in some measure adjudged the facts as well as the law of a particular case in advance of hearing it'.[20] From the perspective of automated legal decision-making, this requirement raises fascinating concerns about the nature of prejudgment. An automated system may only be able to recognise a fact pattern in a binary or crude way. Is that an illicit 'prejudgment' when it comes to scenarios where the computational classifier does not fit well?[21] If so, might it require some human intervention to permit a more subtle or alternative classification? These are important questions which will need to be addressed more directly as 'robo-judging' of administrative claims becomes available.

The requirement of an unbiased tribunal is a rarely invoked, but still important part of US law. It cannot be dispensed with entirely, even if it has been applied in a limited manner. Could an automated process satisfy the standard? Here views divide sharply. On the one hand, there is a formalistic conception of law that would see the ultimate root of bias (both associational and substantive) in human beings' tendency to emotional or ideological thought. On this logic, a machine may be far less likely to consider extraneous factors when deciding a case. On the other hand, given the open-endedness of many legal requirements, the very definition of what is extraneous is itself an exercise of judgment that may be biased. So, while automation of 'closed' legal processes (where all relevant factors have been decided as a matter of law in advance and there are agreed upon ways of recognising them computationally) may not implicate the 'neutral tribunal' requirement, automation of far more common 'open-ended' legal processes will require some attention to this issue.

The types of scrutiny that may be necessary include algorithmic auditing to ensure that automated decision making (ADM) does not itself constitute or exacerbate a biased tribunal. Such auditing may try to insure both input and output legitimacy. On the output side, an ADM process that is continually ruling against certain groups of persons at a significantly higher rate than average should be flagged for further inquiry, and perhaps deserves abandonment on those grounds alone. On the input side, those developing an ADM system should be particularly careful about ensuring representative and unbiased data is used. Moreover, the major goal of the adoption of an ADM system should not be one-sided, such as 'cutting benefits'. As evidenced in controversies in Australia over 'robo-debt' and in the U.S. with respect to algorithmic Medicaid service level determinations, such goals tend to tilt the system toward wrong or strained interpretations of both fact and law in order to satisfy political goals.[22]

Just as some critical level of tribunal neutrality is in part constitutive of due process, the opportunity for public attendance or awareness should also be a norm for administrative hearings. This is one area where further technifica-

[20] *Cinderella Career & Finishing Schools, Inc. v FTC* 425 F.2d 583, 591 (D.C. Cir. 1970). At least one commentator has found it troubling for courts to concern themselves with decisionmakers' speeches that mentioned company names merely to 'illustrate the kind of policies the decisionmaker preferred'. See Richard J Pierce, 'Political control versus impermissible bias in agency decisionmaking: lessons from Chevron and Mistretta' (1990) 57(2) The University of Chicago Law Review 481, 494.

[21] For exercises of quasi-legislative authority, such as rulemaking, 'An agency member may be disqualified from such a proceeding only when there is a clear and convincing showing that he has an unalterably closed mind on matters critical to the disposition of the rulemaking'. See *Association of National Advertisers v FTC* 627 F.2d 1151, 1153 (D.C. Cir. 1979).

[22] Ryan Calo and Danielle Keats Citron, 'The automated administrative state: A crisis of legitimacy' (2020) 70 Emory Law Journal 797; Terry Carney, 'Robo-debt illegality: The seven veils of failed guarantees of the rule of law?' (2019) 44(1) Alternative Law Journal 4; Danielle Keats Citron, 'Technological due process' (2007) 85 Washington University Law Review 1249.

tion may be entirely consonant with transparency goals. What could be more conducive to public attendance than automatic broadcast of any given proceeding to anyone with an internet connection? However, for reasons of privacy, of both litigants and decision-makers, administrators may decide to keep certain sensitive proceedings out of the public eye. They may also restrict access to certain aspects of generally accessible proceedings, or even to shield some data about proceedings, putting conditions on its use. These practices suggest that *reasonable* 'public attendance and understanding of the adjudicative process' is the key constitutive practice to be preserved here.

France's nuanced approach to public access is instructive here. In 2019, it expanded public access to judicial decisions on the Legifrance website by adding lower court decisions to the extant collection of higher court decisions that were already being made available online.[23]. However, to promote privacy and public confidence in the judicial system, it also restricted certain uses of the data.[24] US law professors Livermore and Rockmore claimed that the law also made it illegal to engage in 'what is often called "judicial analytics", which is (roughly) the use of statistics and machine learning to understand or predict judicial behaviour.'[25] Numerous Anglophone commentators immediately condemned this alleged result of the law.

By contrast, Professor G'sell of Sciences Po offered a more positive assessment. For G'sell, the new French law signified a major advance in transparency, since it opened up so much new text to be analysed, even if some forms of analysis were forbidden. But transparency should not be absolute. She observed a deep cultural divide between more personality-focused common law systems of law and the French system:

> In the French legal system, judicial decisions are written in an impersonal manner and rendered in the name of the French people ('*au nom du peuple français*'). They take in principle the form of a syllogism. Decisions are usually made by a panel of three judges whose deliberations are confidential: no one knows whether a particular judge decided with the majority or dissented. In some sense, judges, as individuals, disappear behind the institution of justice.[26]

To preserve this institution-centred (rather than person-centred) judicial culture, the law prohibits analytics which attempt to draw conclusions about particular judges or clerks. As G'Sell puts it, 'it is perfectly legal to make predictions regarding the possible decisions of the Paris Court of Appeal or of the 3rd civil division of the Paris Court of Appeal. The only action that is prohibited is data analytics applied to Judge Dupont [that is, any particular judge], who has been sitting in the Paris Court of Appeal for the past 10 years, in order to provide information about Judge Dupont's opinions or professional practice.'[27]

G'Sell's nuanced approach is well worth exploration in the US, given what David Pozen has called 'transparency's ideological drift': its tendency to frustrate the very ends that it was originally proposed to advance.[28] As Pozen argues, transparency is 'not as an end in itself, but rather as a means to achieve particular social goods', since it is not 'a transcendent normative ideal, but [...] an administrative technique like any other – with contestable moral, political, and distributional implications.'[29] This is a welcome reminder that, in an age of big data, not all that can be done, should be done. Consider, for instance, this vision of an administrative panopticon:

> In an age in which government can watch our every move, why can't we watch its moves? [...] Government actors could be required to videotape their *ex parte* meetings with private entities and place them online, along with automatically produced,

[23] Loi n°2019-222 du 23 mars 2019 de programmation 2018-2022 et de réforme pour la justice, publiée au Journal Officiel du 24 mars 2019.
[24] ibid.
[25] M Livermore and D Rockmore, 'France kicks data scientists out of its courts' (*Slate*, 2019) ⟨https://slate.com/technology/2019/06/france-has-banned-judicial-analytics-to-analyze-the-courts.html⟩.
[26] Florence G'sell, 'Predicting courts' decisions is lawful in France and will remain so' (2019) ⟨https://gsell.tech/predicting-courts-decisions-is-lawful-in-france-and-will-remain-so⟩.
[27] ibid.
[28] David E Pozen, 'Transparency's ideological drift' (2018) 128 Yale Law Journal 100.
[29] ibid 103.

> easily searchable transcripts. Officials could even be placed in remote offices, such as in Montana or Alaska. They could, of course, still communicate with regulated parties, but only by email, which would be open to the public. Similarly, if the PRISM project [of the U.S. National Security Administration, revealed in the Edward Snowden leaks] records the metadata from our every phone call and Internet request, government phone calls should be equally transparent, particularly because the public foots the bill for all the chatter. (...) Using a technology like Google street view, as well as face recognition software, firms could easily track who eats lunch with whom at Washington, D.C. lobbying lunch hotspots...[30]

In this way, 'public attendance' of all parts of the administrative process could be achieved via automatic surveillance. AI could complete these tasks far more comprehensively than any person could, via an expedient as simple as a constant 'recording' app placed on the administrator's cell phone (which they would be required to carry at all times). Nevertheless, there must be some limits on the degree of distrust in government personnel, lest the reporting requirements of administration become so invasive that few, if any, are willing to shoulder its burdens.

Thus, the 'omniveillant' capacities of drones and robotics present a paradoxical perfection of transparency requirements for administrative action. They can make exposure so absolute and exacting that it becomes less a mainstay of democratic openness than neo-Maoist sousveillance. The balance between openness and confidentiality, transparency and privacy is a matter requiring *phronesis,* individualised application of laws to facts reflected in the thicket of litigation spawned by the US Freedom of Information Act's exceptions. Openness may be maximised computationally and secrecy can be as well; but the proper balance between them is what requires human judgment.[31]

Rights of the person judged: notice; opportunity to present reasons, call witnesses, and know opposing evidence

Every person should have the right to know about adverse legal action proposed against them and the asserted basis for it. Sometimes this notice requirement is satisfied via actual notice – for example, a person signing a document indicates that they have received notice or they may be captured on video being told of the proposed action and given relevant documentation. In other cases, constructive notice may be enough – for example, the affected person may be deemed to be aware of an advertisement in a publication or (in the case of rule-making) a publication in the Federal Register.

Like an unbiased tribunal, notice, too, is a constitutive practice of a liberal legal order. Authoritarians may favour star chambers as far more efficient, since notice gives the affected party not only the opportunity to plead their case, but also to flee. Nevertheless, liberal legal orders demand that the accused have some chance to state why a certain action should not occur – yet another constitutive practice of law.

There are, of course, *de minimis* or triviality exceptions. For example, in his classic article, Judge Henry Friendly describes the following:

> Good sense would suggest that there must be some floor below which no hearing of any sort is required. One wonders whether even the most outspoken of the Justices [who applied due process protections to withdrawals of social welfare benefits] would require one on the complaint of an AFDC [Aid to Families with Dependent Children] recipient, recounted by Professor Bernard Schwartz, that 'I didn't receive one housedress, underwears [...] They gave me two underwears for $14.10 [...] it should have been $17.60 instead of $14.10'. Although the value of even small benefits should not be deprecated, given the

[30] Adam Candeub, 'Transparency in the administrative state' (2013) 51 Houston Law Review 385.

[31] On this crucial balance between judgment and maximising calculation, see Brian Cantwell Smith, *The promise of artificial intelligence: reckoning and judgment* (MIT Press 2019). For a rich history of the efforts of US courts to balance demands of privacy and openness, see Amy Gajda, *Seek and hide: The tangled history of the right to privacy* (Penguin Random House 2022). This question of secrecy and openness is an ongoing negotiation: for example, the publicity of judges' addresses became a highly contested matter in the US after a judge's son was murdered at the judge's home by a self-described 'anti-feminist lawyer'. The US Congress tightly restricted publication of judges' personal information after this tragedy.

> precarious financial condition of the recipients of AFDC, the cost of providing an evidentiary hearing in such a case must so far outweigh the likelihood or the value of more accurate determinations that final reliance should be placed on the informed good faith of program administrators.[32]

Judge Henry Friendly's implied quantification of the value of additional procedure anticipated the US Supreme Court's important ruling in *Mathews v Eldridge,* which required litigants to give some estimate of the likely increase in accuracy concomitant with a judicial grant of the additional procedure(s) they requested before adverse action could be taken against them. Traffic administration may be thought of, in the abstract, as another area where the stakes are small enough and the benefit of lengthier proceedings so minimal, that no notice (and therefore no opportunity to give reasons) is given before a penalty is imposed. However, even in the case of red light cameras, which penalize drivers for driving through a red light, some jurisdictions have established elaborate procedures, including the obligation of a human reviewer to confirm the machine vision program's recognition of a given situation as 'driving through a red light'.

Transparency, in terms of knowing the opposing evidence, is also critical to a liberal legal order. To be denied access to such evidence compromises the foundational assurance of telling one's own side of the story. However, the right to call witnesses may be circumscribed in instances where their utility is reasonably questioned. In other words, the ways in which a litigant may present a case may be varied but the litigant must know the basis of the case of the party they oppose.

It is difficult to imagine how that case could be 'known' by the litigant where black box, unexplainable AI is utilized, a deficiency which once again may severely limit the scope of automation of law. At the very least, an AI used to 'accuse' a suspect would need to disclose the data it used and perhaps the algorithms as well. Law professor Chad Squitieri has noted the constitutional dimensions of the problem:

> When government investigators request data from companies such as Google, they obtain data on targeted individuals with a guarantee that the data has been collected, stored, and analyzed properly. These guarantees constitute a testimonial statement under the Confrontation Clause [of the US Constitution]. Similar to lab analysts who submit test results of cocaine samples or blood alcohol levels [...] analysts involved with the collection, storage, and analysis of big data must be available for confrontation under the Sixth Amendment.[33]

Squitieri's approach is particularly illuminating, because it emphasises the inevitable human role behind the AI (or big data, as was the common parlance for many statistical processes now rebranded as AI). The AI becomes more legitimate to the extent its human authors are known and accountable for the choices they made at each step of the process of generating the relevant tool.

Obligations of the decision-maker: to make a decision based exclusively on the evidence presented, and to state reasons for the decision

The obligation to make a decision based exclusively on the evidence presented is one of the most problematic of the Friendly factors. Every decision-maker brings certain experiences and understandings to a case. These will, at the very least, colour the decision-maker's reception of the evidence and perhaps will become determinative in some instances. So perhaps the obligation may be better reframed as one of not searching for or receiving evidence about a particular case that is not presented or known by the litigants – and to ensure that all factors that can be made explicit to the litigant, are actually made so. In other words, the requirements to make a record and to state reasons for a decision, subsume the requirement to rule 'exclusively on the evidence provided'.

[32] Friendly (n 15) 1275-76.

[33] Chad Squitieri, 'Confronting Big Data: Applying the confrontation clause to government data collection' (2015) 101 Virginia Law Review 2011. Squitieri is approaching the problem in a criminal context, but the same normative considerations arise in any 'star chamber' situation where the government is refusing to give the rationales or bases for its action.

Focusing specifically on the creation of a record and stating reasons for a decision: this is certainly critical for final decisions in formal administrative hearings. The principle has its limits, though: not all evidentiary rulings need to be justified, nor are all procedural decisions reviewed. Nevertheless, in order to keep open the possibility for such records and explanations, the person affected by the hearing must be able to question the decisionmaker about the bases of decision, seek review of whatever record has been left (the subject of the next section), and seek an overturning of the action based on the inadequacy or bias of the record.

The prospect of automating these critical obligations of the decision-maker (to make a record and give reasons) is an enticing one for many governments. The US case law on due process is littered with grave (if misleading) declarations that any dollars allocated to procedural protections are likely to be denied to recipients.[34] There is widespread impatience and frustration with bureaucracy. But here, as in so many other proposed governmental reforms, the devil is in the details.

The need for explanation of legal decisions is a critical impediment to the spread of such legal automation. While OpenAI has disclaimed offering legal advice on publicly available versions of ChatGPT, promoters of large language models (LLMs) may soon claim that such technology can write legal opinions justifying decisions. Indeed, use of ChatGPT by jurists has already been documented.[35] Concerned about possible misuse of the technology, Professor Juan David Gutiérrez has recommended that '(i) the user must understand how the technology works, acknowledges its limitations and risks, and makes sure that the tool is adequate for the required task (informed use); (ii) the user is transparent about the use of the technology in proceedings (transparent use); (iii) the user distinguishes clearly which sections of the judicial decision or legal document are AI-generated text (ethical use); and, (iv) the user rigorously checks information retrieved from the AI system against reliable sources and explicitly informs about such examination (responsible use).'[36] These are all vital conditions for adoption of automatic text generation in judicial contexts, especially given the shortcomings of current LLM evaluation methodology identified by Research Scientist Sasha Luccioni and Professor Anna Rogers.[37] Given the opacity of even basic aspects of many LLMs, such as the text corpora and methods of reinforcement learning by human feedback (RHLF) used to train them, even the first condition prescribed by Gutierrez is unlikely to be met any time soon.

Even if all four of Gutierrez's conditions are met, there are still serious concerns about enabling justification via the next token prediction characteristic of LLMs. While the 'unreasonable effectiveness of data' (to evoke Alon Halevy, Peter Norvig, and Fernando Pereira's resonant phrase) may well enable LLMs to pass what Eugene Volokh has called a 'John Henry' test (that is, generating work product and oral questions and responses that even seasoned and experienced legal professionals cannot distinguish from that of other seasoned and experienced legal professionals), the relevant AI cannot actually understand or feel what it is doing.[38] For commentators like Kiel Brennan-Marquez and Stephen Henderson, such empathy is a *sine qua non* for the legitimate application of state power.[39] It is also one reason why Joseph Weizenbaum, a professor of computer science at MIT at the dawn of AI research, found the automation of judicial proceedings so troubling:

[34] *Mathews v Eldridge* 424 US 319 (1976). This is misleading on two levels: taxes can be raised, and at least in the US (a sovereign currency issuer), the real constraint on government spending is the productivity of the economy as a whole. More exacting scrutiny may result in better classification of claimants and matching them to proper sources of work and/or aid. Of course, conditions may be different in Europe, where member states have ceded authority for sovereign currency issuance to a supranational entity.

[35] Juan David Gutiérrez, 'ChatGPT in Colombian Courts: Why we need to have a conversation about the digital literacy of the judiciary' (2023).

[36] ibid.

[37] Alexandra Sasha Luccioni and Anna Rogers, 'Mind your Language (Model): Fact-Checking LLMs and their Role in NLP Research and Practice' (2023) ⟨https://arxiv.org/abs/2308.07120⟩.

[38] On Volokh's test, see Volokh Eugene, 'Chief Justice Robots' (2019) 68(58) Duke Law Journal 1137.

[39] Marquez and Henderson (n 5). For philosophical work on the importance of the fusion of knowledge and feeling, see the discussion of 'epistemopathy' in Steven Connor, *The madness of knowledge: On wisdom, ignorance and fantasies of knowing* (Reaktion Books 2021); Martha C Nussbaum, *Upheavals of thought: The intelligence of emotions* (Cambridge University Press 2003); Agnes Heller, *A theory of feelings* (Lexington Books 2009).

> What could be more obvious than the fact that, whatever intelligence the computer can muster, however it may be acquired, it must always necessarily be absolutely alien to any and all authentic human concerns? The very asking of the question, 'What does a judge... know that we cannot tell a computer', is a monstrous obscenity. That it has to be put into print at all, even for the purpose of exposing its morbidity, is a sign of the madness of our times… What emerges as the most elementary insight is that, since we do not now have any ways of making computers wise, we ought not now to give computers tasks that demand wisdom.[40]

This problem of meaning is exacerbated whenever 'black box AI' is utilised, whether in the form of LLMs or other computation. 'Black box AI' refers to any natural language processing, machine learning, textual analysis or similar software which uses data not accessible to the data subject and/or which deploys algorithms which are either similarly inaccessible or so complex that they cannot be reduced to a series of rules and rule applications comprehensible to the data subject. The less transparent AI is, the less it is capable of meeting the demands of a public record.

Post-hearing right to review alleged errors of fact and law

There is a strong commitment in liberal legal orders to permitting a review of a decision, to ensure that errors of fact or law were not determinative of an outcome. In principle, an automated legal system could provide such a review. Indeed, there is some movement toward limited automated review in the US Social Security Administration for very specific errors (for example, to spot a wrongful denial of benefits).[41] However, anyone familiar with the semantic complexity of judgment involved in the process of deciding whether and how to resolve appeals should be sceptical of a robot judge's ability to complete this task. There is not enough data to allow for competent automation here, particularly given the 'case of first impression' status of many appeals.

Moreover, even if one could posit an exceptionally well-tuned large language model that generated appellate decisions that were affirmed by judges for some period of time (as posited in Volokh's 'John Henry' test mentioned above), how would its operator demonstrate the continuing validity of its opinions? What types of data should be added to 'update' the algorithm over time? The difficulty of answering such questions without knowing the full range of cases to be decided in advance (which is itself impossible to predict) cautions against automation of appellate review functions, and applies *a fortiori* to most visions of a 'legal singularity' or pervasively personalized law.

Conclusion

While computational Intelligence Augmentation (IA) promises many improvements to the legal system, there are some normative limits on automation, particularly the types of AI now hyped as human-like text generators. The most important lesson of this article is to avoid elimination of constitutive practices of liberal legal orders by automating them prematurely. There are many constitutive practices of law that could be substituted for the word 'action' below, in Agre's classic characterisation of formalisation as 'a highly organized form of social forgetting':

> A word like 'action' might present real challenges to a philosophical project that aims to respect ordinary usage, but the assimilation of action to formal language theory reduces the word to a much simpler form: a repertoire of possible 'actions' assembled from a discrete, finite vocabulary of 'expressive elements' or 'primitives'. Having thus taken its place in the technical vocabulary of AI, the word's original semantic ramifications are lost as potential resources for AI work. The ideology surrounding formalization accords no intrinsic value to these left-over materials. As a result, formalization becomes a highly organized form of social forgetting -- and not only of the semantics of words but of their historicity as well. This is why the historical

[40] Joseph Weizenbaum, *Computer power and human reason: From judgment to calculation* (WH Freeman and Company 1976) 226-27.

[41] Kenneth Abbott, Yen-Yi Ho, and Jennifer Erickson, 'Automatic health record review to help prioritize gravely ill social security disability applicants' (2017) 24(4) Journal of the American Medical Informatics Association 709.

> provenance and intellectual development of AI's underlying ideas claim so little interest among the field's practitioners.[42]

Recall that toward the beginning of this article, I mentioned the 'naturalisation' or 'ossification' critique of constitutive practices of person to person interaction, which dismissed them as mere tradition. From this futurist perspective, the role of legal engineers (of the present and future) is expansive, even totalising: all practices must be scrutinised for their outcomes and if there are more efficient means to those outcomes, those means are to be developed and used. Agre's wisdom here is to highlight all the data (including qualitative impressions of persons in the moment of decision) left out when an automated process is said to be 'data-driven'. Especially in the complex legal contexts discussed above, the term 'data-driven' should be replaced with 'driven by some selection of data, representative of a particular place in time, and always in need of updating.'[43] As Thea Snow argues, '[w]e should constantly be asking ourselves, what data is missing? Whose voices are missing? What untested assumptions are we making, and how do these obscure other truths? [...] how [do] our position and power in the world shape how we [...] make decisions using the information collected?'[44]

Agre spent years attempting to build systems that could accomplish the types of searches and navigation necessary for a robot to make its way through the physical world, making plans and completing tasks. He became deeply knowledgeable about the history of diverse research programs in AI, from symbolic analysis to connectionism to early prototypes of the neural networks that have become so pervasive and successful today. But he ultimately became one of the great critics of artificial intelligence, only after years of trying to accomplish some of its central aims. As a recent celebration of his work noted:

> Agre began to rebel, in a sense, from his profession, seeking out critics of artificial intelligence, studying philosophy and other academic disciplines. At first, he found the texts 'impenetrable', he wrote, because he had trained his mind to dissect everything he read as he would a technical paper on math or computer science. 'It finally occurred to me to stop translating these strange disciplinary languages into technical schemata, and instead simply to learn them on their own terms', he wrote.[45]

The constitutive practices of law depend on 'strange disciplinary languages' and resist standardisation. Agre's dissatisfaction with AI arose with his frustration at translating aspects of human experience into formalised terms. The same frustration has recurred among critics of the automation of both public law (via AI in public administration) and private law (via 'smart contracts' and similar technology). There are aspects of adjudication or even many deals that cannot be reduced to machine-readable code. When those aspects are also constitutive practices of liberal legal orders, they should not be eliminated merely in order to make legal practices more tractable to automation.

References

Abbott K, Ho Y.-Y, and Erickson J, 'Automatic health record review to help prioritize gravely ill social security disability applicants' (2017) 24(4) Journal of the American Medical Informatics Association 709.

Agre PE, 'The soul gained and lost: Artificial intelligence as a philosophical project' (1995) 4(2) Stanford Humanities Review 1.

Albergotti R, 'He predicted the dark side of the Internet 30 years ago. Why did no one listen?' (*Washington Post*,

[42] Philip E Agre, 'The soul gained and lost: Artificial intelligence as a philosophical project' (1995) 4(2) Stanford Humanities Review 1. As literary critic Leo Robson has argued, '[The] most pervasive way in which we currently make sense of the world [is an] ingratitude or amnesia that parades as a cleansing iconoclasm.' See Leo Robson, 'Response to Query: What Forms of Art, Activism, and Literature Can Speak Authentically Today' (Bookforum 2021) ⟨https://www.bookforum.com/print/2802/what-forms-of-art-activism-and-literature-can-speak-authentically-today-24492⟩.

[43] The words in quotes are my formulation, but inspired by works in Lisa Gitelman, *Raw data is an oxymoron* (MIT Press 2013) and Charles Taylor's critique of 'brute data' in Charles Taylor, 'Human agency and language' (1985) 2 Philosophical Papers.

[44] Thea Snow, 'The (il)logic of legibility – Why governments should stop simplifying complex systems' (*LSE Impact of Social Sciences*, 2021) ⟨https://blogs.lse.ac.uk/impactofsocialsciences/2021/02/12/the-illogic-of-legibility-why-governments-should-stop-simplifying-complex-systems⟩.

[45] Reed Albergotti, 'He predicted the dark side of the Internet 30 years ago. Why did no one listen?' (*Washington Post*, 2021) ⟨https://www.washingtonpost.com/technology/2021/08/12/philip-agre-ai-disappeared⟩.

2021) ⟨https://www.washingtonpost.com/technology/2021/08/12/philip-agre-ai-disappeared⟩.

Alexander J, 'Three rival views of tradition (Arendt, Oakeshott and MacIntyre)' (2012) 6(1) Journal of the Philosophy of History 20.

Baum L and Devins N, *The company they keep: How partisan divisions came to the Supreme Court* (Oxford University Press 2019).

Bevir M and Blakely J, *Interpretive social science: An anti-naturalist approach* (Oxford University Press 2018).

Calo R and Citron DK, 'The automated administrative state: A crisis of legitimacy' (2020) 70 Emory Law Journal 797.

Candeub A, 'Transparency in the administrative state' (2013) 51 Houston Law Review 385.

Carney T, 'Robo-debt illegality: The seven veils of failed guarantees of the rule of law?' (2019) 44(1) Alternative Law Journal 4.

Citron DK, 'Technological due process' (2007) 85 Washington University Law Review 1249.

Cohen JE, *Between truth and power* (Oxford University Press 2019).

Connor S, *The madness of knowledge: On wisdom, ignorance and fantasies of knowing* (Reaktion Books 2021).

Donabedian A, 'Evaluating the quality of medical care' (1966) 44(3) The Milbank memorial fund quarterly 166.

Dorsen DM, *Henry Friendly, greatest judge of his era* (Harvard University Press 2012).

Engstrom DF and others, 'Government by algorithm: Artificial intelligence in federal administrative agencies' [2020] (20-54) NYU School of Law, Public Law Research Paper.

Eugene V, 'Chief Justice Robots' (2019) 68(58) Duke Law Journal 1137.

Friendly H, 'Some Kind of Hearing' (1975) 1268 University of Pennsylvania Law Review.

Fuller LL, *The morality of law* (Yale University Press 1969).

G'sell F, 'Predicting courts' decisions is lawful in France and will remain so' (2019) ⟨https://gsell.tech/predicting-courts-decisions-is-lawful-in-france-and-will-remain-so⟩.

Gajda A, *Seek and hide: The tangled history of the right to privacy* (Penguin Random House 2022).

Gitelman L, *Raw data is an oxymoron* (MIT Press 2013).

Gutiérrez JD, 'ChatGPT in Colombian Courts: Why we need to have a conversation about the digital literacy of the judiciary' (2023).

Heller A, *A theory of feelings* (Lexington Books 2009).

Hildebrandt M, 'The adaptive nature of text-driven law' (2021) 1(1) Journal of Cross-disciplinary Research in Computational Law.

Karpik L, *Valuing the unique: The economics of singularities* (Princeton University Press 2010).

Livermore M and Rockmore D, 'France kicks data scientists out of its courts' (*Slate*, 2019) ⟨https://slate.com/technology/2019/06/france-has-banned-judicial-analytics-to-analyze-the-courts.html⟩.

Luccioni AS and Rogers A, 'Mind your Language (Model): Fact-Checking LLMs and their Role in NLP Research and Practice' (2023) ⟨https://arxiv.org/abs/2308.07120⟩.

Marquez KB and Henderson SE, 'Artificial Intelligence and Role-Reversible Judgment' (2019) 109(2) The journal of criminal law and criminology 1.

Nussbaum MC, *Upheavals of thought: The intelligence of emotions* (Cambridge University Press 2003).

Pasquale F, 'A rule of persons, not machines: the limits of legal automation' (2019) 87 George Washington Law Review 1.

Pierce RJ, 'Political control versus impermissible bias in agency decisionmaking: lessons from Chevron and Mistretta' (1990) 57(2) The University of Chicago Law Review 481.

Pozen DE, 'Transparency's ideological drift' (2018) 128 Yale Law Journal 100.

Robson L, 'Response to Query: What Forms of Art, Activism, and Literature Can Speak Authentically Today' (Bookforum 2021) ⟨https://www.bookforum.com/print/2802/what-forms-of-art-activism-and-literature-can-speak-authentically-today-24492⟩.

Sahlgren M and Carlsson F, 'The singleton fallacy: Why current critiques of language models miss the point' (2021) 4 Frontiers in Artificial Intelligence 682578.

Sheppard B, 'Incomplete innovation and the premature disruption of legal services' [2015] Michigan State Law Review 1797.

Silverman J, 'Angry Users Want DoNotPay to Pay Up: Troubles Accumulate for the Robot Lawyer That's Not a Robot or a Lawyer' (2023) ⟨https://www.

jacobsilverman.com/p/angry-users-want-donotpay-to-pay⟩.

Smith BC, *The promise of artificial intelligence: reckoning and judgment* (MIT Press 2019).

Snow T, 'The (il)logic of legibility – Why governments should stop simplifying complex systems' (*LSE Impact of Social Sciences*, 2021) ⟨https://blogs.lse.ac.uk/impactofsocialsciences/2021/02/12/the-illogic-of-legibility-why-governments-should-stop-simplifying-complex-systems⟩.

Squitieri C, 'Confronting Big Data: Applying the confrontation clause to government data collection' (2015) 101 Virginia Law Review 2011.

Stengers I, *Thinking with Whitehead: A Free and Wild Creation of Concepts* (Harbard University Press 2014).

Taylor C, 'Human agency and language' (1985) 2 Philosophical Papers.

Wallerstein IM, *Historical capitalism with capitalist civilization* (Verso 1995).

Weizenbaum J, *Computer power and human reason: From judgment to calculation* (WH Freeman and Company 1976).

A reply: Striving Towards Non-Interference with Constitutive Practices

Samuel Judson • Nexus, sam@sjudson.com

In Constitutive Practices of Administrative Law as Limits on Legal Automation, Frank Pasquale convincingly argues that justice is not simply an output of a legal process that can be replicated through automation, but rather is situated within the correct and consistent application of just procedures. To Pasquale, skeptics of legal automation are not trapped in a stubborn embrace of tradition due to ossification or naturalisation. Instead, they are being justifiably protective of constitutive practices that are an essential part of the law, not just the minutia of applying it. By rushing to automate constitutive practices, technologists risk delegitimizing the law's claim to substantive justice – and that is before even considering all of the ways that (machine learning-based systems in particular) can encode biases, hallucinate falsehoods, or otherwise fail.

Pasquale therefore argues against yielding control of information flows and decision making to legal automation. But not all such limits would be created equal. Restricting automation to decision support for a human-in-the-loop cannot, *e.g.*, on its own guarantee the integrity of a constitutive practice. Consider the '[o]bligations of the decision-maker'. Perhaps the most infamous deployments of automated legal decision support are for risk scoring of recidivism[46], while Pasquale himself raises the generative use of LLMs in judicial opinion drafting. Used appropriately – if one believes that is possible – neither directly infringe on the obligation of the decision maker to decide based on the evidence presented and to articulate a reason for their decision. The constitutive practice is only undermined as a matter of degree, when the decider improperly substitutes the tool's (mimicry of) reasoning and judgement for their own, rendering the decision 'something else altogether'. As we cannot know from the outside where the spark of a judgement rests inside the mind of a judge it is hard to set practical and effective limits on the use of such decision support tools. Determining when their use goes too far can only rest on empirical assessment of outcomes and earnest but ultimately subjective self-criticism on behalf of the decision makers.

More broadly, as Pasquale notes, 'no one claims that a court that uses a word processor is less legitimate than one which writes its judgments on parchment with a quill', yet nonetheless it is well-established in the human-computer interaction literature that automating even 'rote and routinizable actions' can significantly impact the processes that consume their output[47]. Greater societal acceptance of and trust in automation – whether deserved or not – may also lead to a commensurate broadening of the scope of legal responsibilities accepted as routine enough to be safe and appropriate for (partial) automation. Rather than wholesale replacement of constitutive practices – which can be recognized and confronted – a more insidious threat may very well be the cumulative weight of 'acceptable' automation gradually hollowing them out. For the constitutive practices of law to maintain their integrity, human responsibility cannot become little more than a well-worn patina coating an otherwise algorithmic engine.

In this I do not believe I disagree with Pasquale. Rather, I interpret his argument as carrying an implicit call for something stronger than its explicit call for limiting the au-

[46] Julia Angwin and others, 'Machine Bias' [May 23rd, 2016] ProPublica (https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing); Vivian Lai and others, 'Towards a Science of Human-AI Decision Making: An Overview of Design Space in Empirical Human-Subject Studies' (2023).

[47] Raja Parasuraman, Thomas B Sheridan, and Christopher D Wickens, 'A Model for Types and Levels of Human Interaction with Automation' (2000) 30(3) IEEE Transactions on Systems, Man, and Cybernetics – Part A: Systems and Humans 286.

tomation of constitutive practices. To me, Pasquale functionally argues that legal automation should strive for a standard of *non-interference*, to adopt – and to some extent abuse – a term of art from the computer security literature. A central definition in the study of secure software systems due to Goguen and Meseguer:

> one group of users, using a certain set of commands, is noninterfering with another group of users if what the first group does with those commands has no effect on what the second group of users can see.[48]

When formalized (often in probabilistic terms), non-interference is a common desiderata of cryptographic engineering and program analysis methods – and even whole programming languages – supporting the design of secure software.[49] Mapping to Pasquale's thesis, the 'second group of users' here would be the people who carry out the constitutive practices, and in so doing maintain their integrity. The 'first group of users' would be the automated technologies they are mandated to or choose to rely upon. Beyond just limiting the technologies of the second group from taking over the roles and responsibilities of the first group, the standard of non-interference would further require that the technologies neither improperly influence decision making nor unduly constrain access to information.

Rather than being a seemingly arbitrary negative limitation on technology that invites claims of ossification and naturalisation, non-interference recasts the maintenance of constitutive practices as a positive target – and interesting challenge – for technologists to both theoretically model and empirically measure. It also would apply to the analysis of legal technologies not usually understood as 'automation', from the humble word processor to more advanced (and even speculative) tools like zero-knowledge proofs and verifiable computation that have been proposed to help manage tradeoffs between privacy and accountability.[50] There is also an important distinction to be made regarding whether the decision making under consideration is human or automated. The use of automated tooling may be more easily justified for legal scrutiny of algorithmic decision making, as 'fighting fire with fire' through program analysis, causal reasoning, interpretability, and related methods[51] are often our only means for understanding the root causes of algorithmic harms. While an arbitrary limit on the investigatory use of automation would need a special case for algorithmic decision making to be actively carved out, non-interference is sufficient on its own – such tools cannot be harmfully interfering with processes that are incapable of assessing responsibility without them.

Finally, both Pasquale and myself write in general terms, but in order for interdisciplinary efforts to substantively protect constitutive practices concreteness must be a virtue. Projects such as the COHUBICOL Typology[52] that catalogue and classify legal technologies – and crucially, how they interface with legal practice – present a promising path forwards. As a technologist, I know that in order for me to fruitfully work while respecting constitutive practices would require the guidance of lawyers able and willing to enumerate such practices, articulate their dimensions, and frame their integrity in a way approachable from a computational mindset.

References

Angwin J and others, 'Machine Bias' [May 23rd, 2016] ProPublica (https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing).

Diver L and others, 'Typology of Legal Technologies' (2022).

[48] Joseph A Goguen and José Meseguer, 'Security Policies and Security Models' (1982).

[49] Andrei Sabelfeld and Andrew C Myers, 'Language-based Information-Flow Security' (2003) 21(1) IEEE Journal on Selected Areas in Communications 5.

[50] Jonathan Frankle and others, 'Practical Accountability of Secret Processes' (2018); Joshua A Kroll and others, 'Accountable Algorithms' (2017) 165(3) University of Pennsylvania Law Review 633.

[51] Samuel Judson and others, ''Put the Car on the Stand': SMT-based Oracles for Investigating Decisions' (2024); Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin, '"Why Should I Trust You?" Explaining the Predictions of Any Classifier' (2016); Sandra Wachter, Brent Mittelstadt, and Chris Russell, 'Counterfactual Explanations Without Opening the Black Box: Automated Decisions and the GDPR' (2017) 31 Harv. JL & Tech. 841.

[52] Laurence Diver and others, 'Typology of Legal Technologies' (2022).

Frankle J and others, ‘Practical Accountability of Secret Processes’ (2018).

Goguen JA and Meseguer J, ‘Security Policies and Security Models’ (1982).

Judson S and others, ‘‘Put the Car on the Stand’: SMT-based Oracles for Investigating Decisions’ (2024).

Kroll JA and others, ‘Accountable Algorithms’ (2017) 165(3) University of Pennsylvania Law Review 633.

Lai V and others, ‘Towards a Science of Human-AI Decision Making: An Overview of Design Space in Empirical Human-Subject Studies’ (2023).

Parasuraman R, Sheridan TB, and Wickens CD, ‘A Model for Types and Levels of Human Interaction with Automation’ (2000) 30(3) IEEE Transactions on Systems, Man, and Cybernetics – Part A: Systems and Humans 286.

Ribeiro MT, Singh S, and Guestrin C, ‘“Why Should I Trust You?” Explaining the Predictions of Any Classifier’ (2016).

Sabelfeld A and Myers AC, ‘Language-based Information-Flow Security’ (2003) 21(1) IEEE Journal on Selected Areas in Communications 5.

Wachter S, Mittelstadt B, and Russell C, ‘Counterfactual Explanations Without Opening the Black Box: Automated Decisions and the GDPR’ (2017) 31 Harv. JL & Tech. 841.

Author's reponse

Frank Pasquale

In his thoughtful reply, 'Striving Towards Non-Interference with Constitutive Practices', Samuel Judson translates some of the key concerns of my article into the language of computer science. In this response, I want to reflect on the non-interference principle as a useful guiding principle for future encounters between technologists and the legal system, as well as Judson's welcome call for concreteness in further work in the area.

Judson describes non-interference as a commitment to prevent one set of users of a secure software system from affecting what a second group of users 'can see', given a certain set of commands accessible to the first group of users. A simple example might be drawn from online banking. A secure online banking system should not allow one set of, say, depositors, to use commands in order to prevent other depositors from seeing their own account balances (illicit invisibilization), or to make others' bank balance visible to everyone else (illicit hypervisibilization). Nor would a secure system of this type permit any bank employee from doing the same, without some legitimate process taking place beforehand. For example, law enforcement officials might present specialized bank employees with sufficient evidence to hide the balance of an alleged thief from the alleged thief, based on a narrow exception to the general principle of non-interference.

I can foresee several potential applications of a principle of non-interference (as well as limited exceptions) in the ongoing automation of legal processes. For example, as administrative adjudicators require claimants for benefits to upload forms and to check boxes to indicate aspects of their case, they should encourage developers to try to maintain as many positive affordances of once-more-common face-to-face interactions with bureaucrats. In this way, as Judson observes, '[r]ather than being a seemingly arbitrary negative limitation on technology... non-interference recasts the maintenance of constitutive practices as a positive target – and interesting challenge – for technologists to both theoretically model and empirically measure.' There are already software services trying to imitate some aspects of in-office experience online, now that so many US professionals are predominantly working from home. Perhaps a similar effort can reduce the alienation accompanying many forms of legal digitization.

I also agree with Judson that, 'in order for interdisciplinary efforts to substantively protect constitutive practices, concreteness must be a virtue.' This concreteness was recently exemplified in Jennifer Pahlka's excellent book, *Recoding America*. Though an advocate of automation, Pahlka describes in great detail how the downsizing of civil services left the US paying 'vendors hundreds of thousands of dollars for basic web forms that don't work, and [made] applying for government services feel like the Inquisition.'[53] Pahlka describes an 'an application for food stamps that requires answering 212 separate questions.'[54] While in the abstract, each question probably has some merit, this hypertrophy of process reflects an excessive task specification and differentiation that has almost certainly gone beyond what is statutorily required for the provision of aid, however well it reflects the algorithmatization of law. Attention to what is truly constitutive of such determinations – whether requiring human judgment or computational power – can help refocus policymakers on their true mission, rather than abstract representations of it.

[53] Jennifer Pahlka, *Recoding America* (New York: Metropolitan Books 2023) 62.
[54] ibid.