

# Steps Towards Value-Aligned Systems

Osonde A. Osoba  
oosoba@rand.org  
RAND Corporation  
Santa Monica, CA

Benjamin Boudreaux  
bboudreaux@rand.org  
RAND Corporation  
Santa Monica, CA

Douglas Yeung  
dyeung@rand.org  
RAND Corporation  
Santa Monica, CA

## ABSTRACT

Algorithmic (including AI/ML) decision-making artifacts are an established and growing part of our decision-making ecosystem. They are now indispensable tools that help us manage the flood of information we use to try to make effective decisions in a complex world. The current literature is full of examples of how individual artifacts violate societal norms and expectations (e.g. violations of fairness, privacy, or safety norms). Against this backdrop, this discussion highlights an under-emphasized perspective in the body of research focused on assessing value misalignment in AI-equipped sociotechnical systems. The research on value misalignment so far has a strong focus on the behavior of individual tech artifacts. This discussion argues for a more structured systems-level approach for assessing value-alignment in sociotechnical systems. We rely primarily on the research on fairness to make our arguments more concrete. And we use the opportunity to highlight how adopting a system perspective improves our ability to explain and address value misalignments better. Our discussion ends with an exploration of priority questions that demand attention if we are to assure the value alignment of whole systems, not just individual artifacts.

## CCS CONCEPTS

• **Social and professional topics** → **Technology audits; Sociotechnical systems; Systems analysis and design**; • **Computing methodologies** → *Philosophical/theoretical foundations of artificial intelligence; Systems theory*.

### ACM Reference Format:

Osonde A. Osoba, Benjamin Boudreaux, and Douglas Yeung. 2020. Steps Towards Value-Aligned Systems. In *Proceedings of the 2020 AAAI/ACM Conference on AI, Ethics, and Society (AIES '20)*, February 7–8, 2020, New York, NY, USA. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3375627.3375872>

## 1 INTRODUCTION

Issues around bias in AI/ML and algorithmic artifacts are now common knowledge. There is a significant body of research demonstrating examples of algorithmic bias, its causes, and its features. Typical discussions of biases inherent in AI/ML artifacts include examinations of bias in facial recognition models [5], natural language processing models [7], and risk-estimation tools [1].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

AIES '20, February 7–8, 2020, New York, NY, USA

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.  
ACM ISBN 978-1-4503-7110-0/20/02...\$15.00  
<https://doi.org/10.1145/3375627.3375872>

Algorithmic bias fits into a larger body of concerns dealing with the integration and performance of technology artifacts in society. The problem of *value alignment* [2, 15, 22] captures this larger body of concerns: how can we design systems that achieve their specific objectives while remaining aligned with the broader values of society? The values implicated can include fairness, privacy, safety, trustworthiness, etc. The value alignment frame is useful because it primes us to think about the portfolio of values we hope to induce in our artifacts, as opposed to having an exclusionary focus on single values. This also helps us better appreciate the tensions between norms e.g. how do we balance the elevated data *privacy* risks in minority sub-demographics against the need for more minority data samples for training *fairer* ML models?

Much alignment-related research tends to focus narrowly on the nature of the technology artifacts and their specific value alignment pathologies. Efforts like the MIT Moral Machine [3] use crowd-sourcing to elicit moral intuitions in isolated artifacts in contrived scenarios, completely stripped of the systemic perspective. Other researchers have developed sophisticated learning mechanisms (e.g. inverse reinforcement learning models) [10, 13] to do more than just elicit moral intuitions and try to actually induce value alignment in individual intelligent artifacts.

In contrast to these efforts, this discussion aims to draw out the implications of the fact that no technology artifact exists in a vacuum; each artifact (including AI/ML models) fulfills at least one role within a complex sociotechnical system<sup>1</sup>. Individual artifacts can have complex nonlinear interactions with the rest of the sociotechnical system. They can have hard-to-predict influences on the system's overall performance. So even normatively perfect artifacts are not guaranteed to result in a normatively perfect system e.g. the rate of human overrides observed by Chouldechova et al. [8] after the introduction of an improved AFST tool. There are many other ways in which the use of seemingly innocuous algorithms will be unjust in practice. Our discussion is grounded primarily in work on fairness audits as this was the domain that informed our thinking [18]. Other relevant values for this kind of analysis include privacy and security.

Achieving value alignment in decision systems will require a broader systems-level focus in our research and implementation efforts. And the importance of system-level value alignment audits will likely only grow as complex algorithmic decision-making artifacts become a more established part of our decision-making ecosystem. Audits of individual artifacts will continue to be important insofar as they provide details for a more complete characterization of how decisions propagate in systems. But perfectly aligned

<sup>1</sup>We use the term *sociotechnical system* (STS) to refer aggregate systems that involve complex interactions among human agents, technology artifacts, and the environment around both [4]. For simplicity, we will focus most of our discussion on mission-oriented STSs within an institution e.g. a population of people working with diverse pieces of technologies within the criminal justice institution.

individual artifacts are not the goal; a fair or just sociotechnical system is the goal.

## 2 A SYSTEMS PERSPECTIVE IN PRACTICE?

An important starting point for auditing value alignment in STSs is identifying the structural map of the STS. Eliciting the exact structure of interactions in an STS can be difficult. The goal is to gain a *useful* understanding of the structure of decision-making in that system. So precision and fine detail may not be necessary. Understanding this structure can help explain or diagnose observed systemic problems.

In practice, we have framed this elicitation as the identification of simplified *decision pipelines* in an STS (e.g. in a criminal justice system, employee recruitment or promotion system, etc.). Decision pipelines are construed as collections of connected chains of decision-making points in an institution ('nodes'). Alternatively, decision pipeline can be construed as paths in an institution through which decision outcomes propagate ('edges'). The true complete structure of a typical STS is likely very complex and densely networked.

There are many advantages to this mapping exercise, not all of which need to be related to value alignment. For example, system decision maps can help:

- (1) **Identifying Opportunities for Automated Decisions:** Identify decision points at which algorithmic decision aids may be usefully applied or are already applied. The eligibility criteria generally includes: the existence of a well-defined decision goal, the existence of significant measurable data for training ML models, & the risks and costs associated with incorrect decisions.
- (2) **Performance tracking:** Validating whole-system performance will generally involve validating the system's sub-parts. Identifying the key sub-processes will facilitate this system validation effort.
- (3) **Accountability & Explanation:** Decision system maps help focus efforts to debug failures. *Credit assignment*, the ability to assign credit or blame to sub-parts of a system for observed overall errors or success, is a famously central function for effective system adaptation and optimization [16]. Credit assignment is only possible if the structure of the decision system is, at least, roughly identified.
- (4) **Intervention Planning:** Decision maps also enable better planning for corrective or restorative interventions when value misalignments are in progress.

The next subsection walks through a basic sample of one such decision map developed in the course of another exercise. It is also instrumental for framing the one benefit of the systems perspective we discuss later in this paper.

### 2.1 A Notional System Map: a US Criminal Justice System

There is significant prior literature exploring the use of algorithms at different phases of the criminal justice decision pipeline. Here we will focus on using these discussions to identify and describe a few

different sub-phases of the criminal justice decision pipeline, particularly those in which there is precedent for the use of algorithmic decision aids. For example, risk assessment tools are used to determine pretrial detention decisions, sentencing, treatment/rehabilitation for incarcerated offenders, and parole/probation decisions. These tools were first used for probation decisions, then pretrial decisions, and most recently for sentencing (including some sub-decisions about treatment/rehabilitation while incarcerated). Some key parts of this decision pipeline include (see Figure 1 for a stylized depiction of the pipeline):

**2.1.1 Pretrial detention.** Risk-assessment tools are often used for pre-trial detention and release decisions, and generally place more focus on static risk factors in this context. One pre-trial tool, the Public Safety Assessment (PSA), is used in three states (Arizona, Kentucky, and New Jersey) and dozens of local jurisdictions in other states. The PSA was developed by the Laura and John Arnold Foundation. It was built using data from 1.5 million crimes across 300 U.S. jurisdictions in an attempt to help make decisions about whether an individual should be detained or released before trial, but not any other considerations (e.g., rehabilitation).

**2.1.2 Sentencing.** Sentencing decisions are made for convicted defendants at this decision point. Virginia was the first state to implement a recidivism risk assessment instrument for sentencing decisions relatively recently in 1994 [20]. This instrument, created by the Virginia Criminal Sentencing Commission, was designed to divert 25% of the lowest-risk offenders to non-prison sanctions. In 1999, Virginia developed a second instrument; this time, the goal was to identify the highest-risk offenders. High-risk sex offenders' sentences could then be increased as much as threefold. In Utah, while the judge maintains discretion in the final decision, (s)he must consider both the sentencing guidelines for the offense and the recommendation of the Office of Adult Probation and Parole, which includes an assessment of risk and needs that account for factors including education level, substance abuse, and homelessness.

**2.1.3 Treatment/rehabilitation.** Some states use algorithmic scores similar to those mentioned in the sentencing section to assess which "life skills" classes best meet the inmate's needs and to determine conditions of probation. Newly imprisoned inmates in Pennsylvania are evaluated with the Risk Screening Tool to assess their risk of recidivism and are categorized into low, medium, and high-risk groups. Medium and high-risk inmates may complete further assessments and potentially complete a treatment plan addressing their "criminogenic needs" before being considered for parole [17].

**2.1.4 Parole/probation.** Rudimentary algorithms used for probability-based assessment long predate what we might consider the current algorithmic boom. In the first half of the 20th century, Ernest W. Burgess and others created expectancy tables for parole prediction, which included factors such as work history, marital status, and prison rule violations [6]. In the 1970s, the *Salient Factor Score* [14] was first developed for use in federal parole guidelines; in that era, it was considered unfair to include immutable characteristics (e.g., race, ethnicity, gender, and age) in parole and sentencing criteria. A 1982 report also examined *selective incapacitation*, a probabilistic approach to prioritizing incarceration targets by stratifying criminal risk [12].

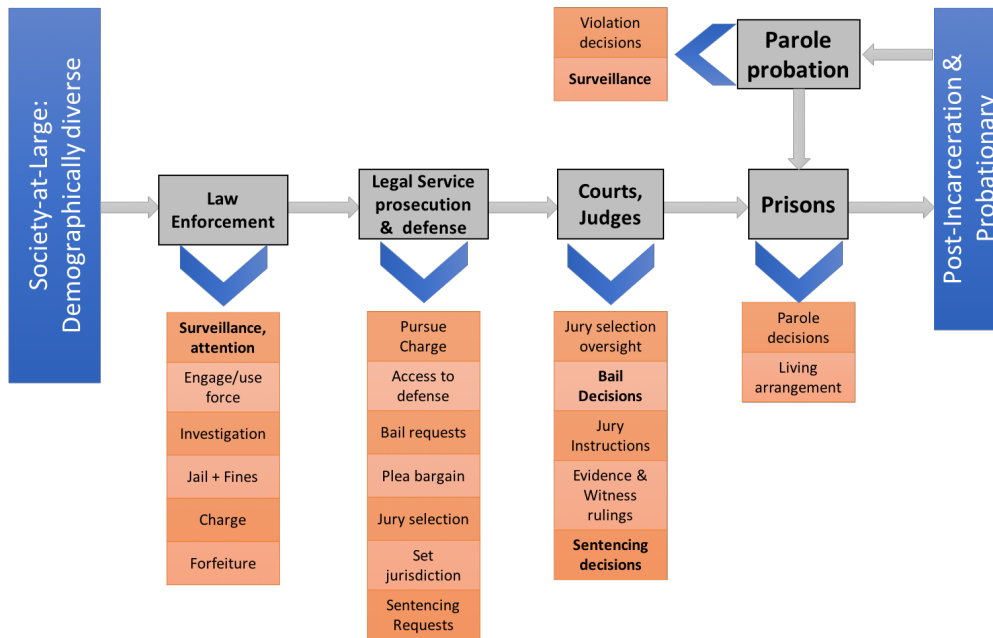


Figure 1: Notional example of a decision pipeline within a sociotechnical system (a US criminal justice system).

Besides these uses, other modern decision points in the criminal justice pipeline include predictive policing, facial recognition, and probabilistic DNA analysis [9].

### 3 THE FRUITS OF A SYSTEMS-LEVEL FRAMING

The value of a systems perspective on value alignment is at least two-fold: it has more explanatory power for the observed alignment pathologies; & it gives us more tools and flexibility for correcting alignment problems. We highlight the advantages with two examples: 1) the concern around feedback cycles, and 2) an unexpected adaptation to the introduction of algorithms.

- (1) **Algorithm overrides:** Chouldechova et al. tells of an STS pathology in a child welfare decision pipeline [8]. The team developed and deployed an algorithmic decision-making aid to flag high-risk cases in the child welfare system for caseworkers. The model was certified to have accuracy improvements over the status quo. However, case-workers were overriding the model's risk estimation decisions at a higher than justified rate. The reasons for the overrides were still under investigation at the time they published their results. The example illustrates that focusing on the efficiency of sub-parts of a decision-pipeline may be insufficient to improve the value-alignment of the whole system. There are *impedance mismatches* or transmission barriers as information and decisions flow through an STS. This is especially true when the agents in the STS (human or algorithmic)

have significant discretion in adapting to changes<sup>2</sup>. Whole-system value alignment requires careful accounting for such transition costs.

- (2) **Runaway feedback loops:** A myopic examination of a criminal justice decision pipeline may predispose a researcher to expect that as long as only perfectly fair algorithms are introduced the overall system will operate fairly. But simulation studies of criminal justice decision pipelines [11, 19] show that seemingly innocuous and fair predictive policing algorithms based on discovered crime rates (the form of data most easily available for training ML models) can lead to runaway feedback loops in inequitable patterns of criminalization (where equity is defined in terms of probability of discovery given crime commission). This particular alignment pathology can be traced to dynamics earlier in the decision pipeline, at the surveillance decision point (also further back in time). A myopic focus on just current predictive policing algorithms would be insufficient for under and addressing this misalignment.

These documented issues make a strong case that researchers and decision-makers need to adopt a systems perspective if they are to be effective at adapting (or designing) value-aligned systems. And this is true regardless of what value, theory of justice, or philosophy the practitioner brings to the table.

#### 3.1 Priority Research Questions


A systems-level perspective is just a first step. Research programs based on this perspective (ours included) often still revert to just using system decision maps to identify individual algorithms worth auditing. Making full effective use of this framing will require a

<sup>2</sup>The effects of such discretion or latitude are not necessarily normatively good or bad.

few new tools in our methodological arsenal, a minimal selection of which could include work on:

- (1) **Characterizing Decision Transition Costs:** there can sometimes be barriers to the transmission of decisions at interfaces between AI/ML and human decision agents in a system (what we called Impedance Mismatches earlier). The reasons for such barriers can include: lack of trust in the technology, cultural aversion, the practice of human discretion, explanatory gaps (e.g. when the model is opaque), etc. There are already bodies of research on some of these concerns. Translating these types of work to better characterize and understand pathologies in sociotechnical decision systems could be enlightening.
- (2) **Complex Adaptive Models of STS:** STSs are comprised of adaptive human agents/organizations and relatively static artifacts. Such complex adaptive systems (CASs) can exhibit hard-to-predict behaviors. We need principled tools and methods for modeling and simulating the interaction of human and artificial agents in these STSs. These approaches will enable better diagnostics and situational awareness for systems-level value-alignment issues.
- (3) **Designing Advanced Mechanisms for Redress or Compensation and:** The focus on single artifact performance limits the range of possibilities for redress. There is a heavier focus on just designing more advanced technologies to overcome normative misalignment. This can be a myopic response. A systems perspective allows for a broader range of interventions e.g. robust oversight and compensation procedures built around imperfect artifacts.
- (4) **Statistical Learning Models of Regulation:** The increasing complexity of technology artifacts (including AI/ML models) has placed a significant strain on the ability of regulating agencies to constrain the adverse impacts of these technologies. Current events routinely highlight this shortfall in regulatory capacity e.g. Boeing safety crisis of 2019, the financial crash of 2008, etc. There is a case for exploring and introducing regulatory models based on statistical learning theory and control theory, especially when the regulatory goal is focused on *measurable* outcomes and consequences (as opposed to procedural norms) e.g. the regulator only cares about safety or outcome fairness. Optimal regulatory schemes in such outcome-focused or consequentialist “regulatory games” can be modeled as reinforcement learning problems or optimal control theory (Figure 2) i.e. what optimal policies can a regulator learn to achieve a desired observable outcome in a population of agents?
- (5) **Value Alignment in Multi-stakeholder STSs:** An implicit assumption in much of the work on value alignment is that *the* controlling value exists and is discoverable. But recent work is beginning to highlight problems with this framing [18]. First, the assumption of a universal prescriptive controlling value for an application is very much in direct opposition with most practical experience with ethical decision-making. Second, even if such a value exists, there is no guarantee that all relevant stakeholders agree on the same

controlling value(s) for an application. Resolving this conflict in a multi-stakeholder application requires resorting to preference aggregation mechanisms. There is already some sophisticated work on preference aggregation (e.g. [10, 21]). There is less work targeting the elicitation of truthful preferences. And far less work on the computational mechanisms for robustly managing the process of compromise in preferences as is commonly observed in participatory modeling of stakeholder norms.



**Figure 2: Illustration of the interaction between free actor and governing regulator agents in an outcome-focused regulatory game. The regulator agent observes state information relevant to their regulatory mandate (in this case fairness). While the actor agent pursues its goals subject to its internal efficient mandate. Regulators may take actions to impose costs or rewards on the actors to incentivize value-aligned behavior in the population of actors.**

## 4 CONCLUSION

In summary, this paper calls for adopting a *systems engineering* approach for doing research on AI value alignment problems. The systems perspective has advantages over the focus to-date on individual misbehaving AI/ML artifacts: it gives us more levers for redressing imperfect individual AI/ML decision artifacts; it can help us anticipate and think more clearly about mismatches that can occur as we embed more AI/ML in human decision systems; and it helps us anticipate and diagnose mechanisms that can lead to adverse feedback cycles more clearly.

Moral values and norms have been central to this discussion and its related research program. But the discussion is still agnostic as to which specific set of values are desirable to induce on STSs and their constituent artifacts. This highlights an opportunity and a need for more philosophy and anthropology research on what broadly-acceptable binding norms to induce in technology artifacts in pluralist societies.

## REFERENCES

- [1] Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. 2016. Machine bias. *ProPublica*, May 23 (2016), 2016.
- [2] Thomas Arnold, Daniel Kasenberg, and Matthias Scheutz. 2017. Value Alignment or Misalignment—What Will Keep Systems Accountable?. In *Workshops at the Thirty-First AAAI Conference on Artificial Intelligence*.
- [3] Edmond Awad, Sohan Dsouza, Richard Kim, Jonathan Schulz, Joseph Henrich, Azim Shariff, Jean-François Bonnefon, and Iyad Rahwan. 2018. The moral machine experiment. *Nature* 563, 7729 (2018), 59.
- [4] Gordon Baxter and Ian Sommerville. 2010. Socio-technical systems: From design methods to systems engineering. *Interacting with Computers* 23, 1 (08 2010), 4–17. <https://doi.org/10.1016/j.intcom.2010.07.003> arXiv:<http://oup.prod.sis.lan/iwc/article-pdf/23/1/4/2038336/iwc23-0004.pdf>
- [5] Joy Buolamwini and Timnit Gebru. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*. 77–91.

- [6] Ernest W Burgess. 1936. Protecting the public by parole and by parole prediction. *Am. Inst. Crim. L. & Criminology* 27 (1936), 491.
- [7] Aylin Caliskan, Joanna J Bryson, and Arvind Narayanan. 2017. Semantics derived automatically from language corpora contain human-like biases. *Science* 356, 6334 (2017), 183–186.
- [8] Alexandra Chouldechova, Diana Benavides-Prado, Oleksandr Fialko, and Rhema Vaithianathan. 2018. A case study of algorithm-assisted decision making in child maltreatment hotline screening decisions. In *Conference on Fairness, Accountability and Transparency*. 134–148.
- [9] Jessica Gabel Cino. 2017. Deploying the Secret Police: The Use of Algorithms in the Criminal Justice System. *Ga. St. UL Rev.* 34 (2017), 1073.
- [10] Vincent Conitzer, Walter Sinnott-Armstrong, Jana Schaich Borg, Yuan Deng, and Max Kramer. 2017. Moral decision making frameworks for artificial intelligence. In *Thirty-first aaai conference on artificial intelligence*.
- [11] Danielle Ensign, Sorelle A Friedler, Scott Neville, Carlos Scheidegger, and Suresh Venkatasubramanian. 2018. Runaway Feedback Loops in Predictive Policing. In *Conference on Fairness, Accountability and Transparency*. 160–171.
- [12] Peter W Greenwood, Allan F Abrahamse, et al. 1982. *Selective incapacitation*. Rand Santa Monica, CA.
- [13] Dylan Hadfield-Menell, Stuart J Russell, Pieter Abbeel, and Anca Dragan. 2016. Cooperative inverse reinforcement learning. In *Advances in neural information processing systems*. 3909–3917.
- [14] Peter B Hoffman. 1983. Screening for risk: A revised salient factor score (SFS 81). *Journal of Criminal Justice* 11, 6 (1983), 539–547.
- [15] Tae Wan Kim, Thomas Donaldson, and John Hooker. 2018. Mimetic vs Anchored Value Alignment in Artificial Intelligence. *arXiv preprint arXiv:1810.11116* (2018).
- [16] Marvin Minsky. 1961. Steps toward artificial intelligence. *Proceedings of the IRE* 49, 1 (1961), 8–30.
- [17] John Monahan and Jennifer L Skeem. 2013. Risk redux: The resurgence of risk assessment in criminal sanctioning. *Fed. Sent'g Rep.* 26 (2013), 158.
- [18] Osonde A Osoba, Benjamin Boudreaux, Jessica Saunders, J Luke Irwin, Pam A Mueller, and Samantha Cherney. 2019. Algorithmic Equity. (2019).
- [19] Osonde A Osoba and William Welser IV. 2017. *An intelligence in our image: The risks of bias and errors in artificial intelligence*. Rand Corporation.
- [20] B Ostrom, Matthew Kleiman, Fred Cheesman, R Hansen, and Neal B Kauder. 2002. *Offender risk assessment in Virginia: a three-stage evaluation*. The National Center for State Courts and the Virginia Criminal Sentencing Commission.
- [21] Dorsa Sadigh, Anca D Dragan, Shankar Sastry, and Sanjit A Seshia. 2017. Active Preference-Based Learning of Reward Functions. In *Robotics: Science and Systems*.
- [22] Nate Soares and Benja Fallenstein. 2014. Aligning superintelligence with human interests: A technical research agenda. *Machine Intelligence Research Institute (MIRI) technical report* 8 (2014).