

Introducing the AI Governance and Regulatory Archive (AGORA): An Analytic Infrastructure for Navigating the Emerging AI Governance Landscape

Zachary Arnold¹, Daniel S. Schiff², Kaylyn Jackson Schiff², Brian Love¹, Jennifer Melot¹, Neha Singh¹, Lindsay Jenkins¹, Ashley Lin¹, Konstantin Pilz¹, Ogadinma Enwereazu², Tyler Girard²

¹Georgetown University, Center for Security and Emerging Technology, Emerging Technology Observatory

²Purdue University, Governance and Responsible AI Lab

zachary.arnold@georgetown.edu, dschiff@purdue.edu, schiffk@purdue.edu, bl711@georgetown.edu, jennifer.melot@georgetown.edu, neha.singh@georgetown.edu, oenwerea@purdue.edu, tgirard@purdue.edu

Abstract

AI-related laws, standards, and norms are emerging rapidly. However, a lack of shared descriptive concepts and monitoring infrastructure undermines efforts to track, understand, and improve AI governance. We introduce AGORA (the AI Governance and Regulatory Archive), a rigorously compiled and enriched dataset of AI-focused laws and policies encompassing diverse jurisdictions, institutions, and contexts related to AI. AGORA is oriented around an original taxonomy describing risks, potential harms, governance strategies, incentives for compliance, and application domains addressed in AI regulatory documents. As of its launch in July 2024, AGORA included data on several hundred instruments, with new entries being added continuously. We describe the manual and automated processes through which these data are systematically compiled, screened, annotated, and validated, enabling deep, efficient, and reliable analysis of the emerging AI governance landscape. The dataset, supporting information, and analyses are available through a public web interface (<https://agora.eto.tech>) and bulk dataset.

Overview and Motivation

As machine learning techniques and generative AI models proliferate, the question of how to govern AI has attracted increasing attention from scholars, policymakers, and the general public. Open issues facing AI policy stakeholders across social and economic sectors include the ideal level of centralization (Cihon, Maas, and Kemp 2020); the role of self-regulation as contrasted against independent oversight (Raji 2022); the creation or adaptation of institutions (Maas 2023); the translation of social and ethical values into policy (Schiff 2023); and the adequacy of provisions for effective monitoring and enforcement (?), among many other issues.

Yet contemporary discussions of AI governance often lack a shared understanding of existing and proposed laws, rules, and norms that govern AI technologies—seriously impeding efforts to evaluate their effects. This missing baseline complicates discussions about what additional measures (or changes to existing measures) are needed to govern increasingly powerful AI systems, makes it difficult for AI developers and users to understand their current obligations,

and prevents those crafting new AI-related regulations from drawing on precedents and prior lessons learned elsewhere.

In short, as AI governance moves from agenda-setting and policy formulation to adoption and implementation, a conceptual and analytic infrastructure is needed to support policy learning. At present, efforts to understand the evolving AI governance landscape face two basic challenges. First, standardized descriptive concepts are necessary for meaningful regulatory comparison, aggregation, and trend analysis (Murdick, Dunham, and Melot 2020; OECD 2022). However, at present there is no widely accepted set of concepts used to describe AI regulations; despite growing efforts, analysts have struggled to taxonomize concepts like AI risks, harms, and incidents (OECD 2023).

Second, and more fundamentally, there is currently no sufficiently systematic, robust, and publicly-available compilation of AI regulations to describe in the first place. AI-related laws, regulations, standards, ethical and normative frameworks, and similar documents (hereinafter “instruments”) are produced nearly daily by a dizzying array of actors with different definitions, jurisdictions, and purposes. Properly understanding this geographically, institutionally, and thematically diverse enterprise requires systematic data collection and processing capabilities that have been lacking to date.

To address these gaps, we present AGORA (the AI Governance and Regulatory Archive), a rigorously compiled dataset of AI governance instruments. AGORA includes instrument text, summaries, and extensive metadata describing the AI-related risks, ethical concerns, governance strategies, incentives for compliance, and application domains addressed in each instrument. Annotators systematically compile, screen, annotate, and validate data using standardized processes and an original taxonomy of regulatory concepts designed to support an effective analytic infrastructure for governance stakeholders.

To our knowledge, AGORA is the most systematic effort to date to characterize the emerging AI regulatory landscape. Earlier and ongoing efforts to map laws and regulations in AI-related domains provide valuable data for some analyses, but each has shortcomings that complicate analysis of trends in AI governance.¹ Among other challenges, these sources

¹Existing efforts include the International Association of

tend to omit essential methodological details, such as the criteria by which they include or exclude documents, the sources and processes used to locate documents nominally in scope, and the criteria and conceptual definitions used to produce original metadata at the document level. Most lack sufficiently detailed metadata and are not captured via interoperable or download-friendly data formats that can support detailed analysis and research of AI governance trends. Many appear to be maintained inconsistently, if at all.

We seek to build on these important prior efforts by defining a well-elaborated taxonomy of concepts relevant specifically to AI governance, and by creating data infrastructure to efficiently and systematically compile information on AI governance instruments as they emerge.

As of its beta launch in July 2024, AGORA included several hundred AI regulatory instruments,² with especially thorough coverage of relevant United States federal laws and regulations enacted between 2020 and 2023. Full text for each instrument is continuously compiled, and additional documents and thematic metadata added as new AI-related laws and regulations emerge. AGORA is available through a public web-based interface (<https://agora.eto.tech>) designed in consultation with subject-matter experts and potential users in government, academia, and the private sector, as well as a bulk dataset suitable for academic analysis.

Below, we discuss AGORA’s taxonomy and scope, methodological procedures, and trajectory. We conclude by considering limitations and proposing substantive policy questions that may be amenable to rigorous analysis and academic research using AGORA.

Defining AGORA: Scope and Key Concepts

Scope of Coverage

AGORA includes *laws, regulations, standards, and similar instruments that directly and substantively address the development, deployment, or use of artificial intelligence technology*. The intent of this scoping definition is to encompass the large majority of instruments created by lawmakers, regulators, and standard-setters in direct response to advances in modern machine learning and related technologies.

Applying subjective elements of this definition, such as “directly and substantively,” inevitably involves judgment. When screening documents for inclusion in AGORA, we try to constrain this judgment by defining heuristics; for example, screeners are instructed that an instrument does not “substantively” address artificial intelligence if it only mentions AI in contextual or non-operative language, e.g., “findings of Congress” provisions in bills, passing mentions, or

Privacy Professionals (IAPP) Global AI Law and Policy Tracker (<https://iapp.org/resources/article/global-ai-legislation-tracker>), OECD AI Policy Observatory (<https://oecd.ai/en>), AI Standards Hub (<https://aistandardshub.org>), and OCEANIS collection of standards (<https://ethicsstandards.org/repository>), as well as more specialized compilations by organizations such as the American Action Forum, Electronic Privacy Information Center, National Conference of State Legislatures, Future of Life Institute, and various legal and consultancy groups.

²For current data, visit <https://agora.eto.tech>.

Federal Register explanatory text accompanying a new regulation. Screeners are also instructed that instruments addressing AI-related concepts such as machine learning, machine autonomy, or algorithmic decision making should generally be considered in scope; that is, the presence or absence of the specific term “artificial intelligence” is not determinative. In turn, instruments addressing technologies such as autonomous vehicles and synthetic media are all potentially within AGORA’s scope (and indeed, the current dataset includes instruments related to these topics). The full set of current screening heuristics and accompanying guidance is publicly available (see online appendices).

Critically, the requirement that instruments “directly” address artificial intelligence generally excludes laws predating the rise of modern machine learning, even if they are broad enough in scope to bear on AI. We draw this line to ensure that AGORA’s scope is manageable in practice and to reinforce the dataset’s emphasis on policies created in response to 21st century developments in AI, rather than the entire set of policy instruments that may affect individual sectors and AI governance writ large. Note, however, that more recent instruments that tailor these broad laws to the specific context of AI would qualify for inclusion in AGORA. For example, while the Civil Rights Act of 1964 would not be included in AGORA, a related federal regulation or guidance document applying the Act to racially discriminatory AI is within AGORA’s scope.

Taxonomizing AI Governance

AGORA is built around a conceptual taxonomy that is inspired by scholarly and policy literature, but intended to be useful to a wide range of potential users and reasonably intuitive to both those users and AGORA annotators. This necessarily involves balancing comprehensiveness with parsimony and interpretability. The taxonomy was drafted by an interdisciplinary team with training in law, data engineering, public policy, political science, AI governance, and quantitative and qualitative social science methods, with input from potential AGORA users in government, academia, and the private sector, and has been refined iteratively based on annotator and user feedback. The approach is thus both theory-driven and in part inductive, drawing on existing literature and adapting where gaps were identified.

The taxonomy consists of discrete concepts (“codes”) organized into five overarching domains, defined below:

- **Risk factors governed:** The characteristics addressed by the instrument that affect AI systems’ propensity to cause harm, closely related to ethical concerns. Surveying these characteristics is a focus of many current AI policy efforts (Barrett et al. 2023; Bengio et al. 2023; European Commission 2021), while others focus on specific risk characteristics such as those related to bias, safety, or security. AGORA’s risk codes are adapted from the risk categories outlined in the widely-used AI Risk Management Framework issued by the U.S. National Institute of Standards and Technology (2023).
- **Harms addressed:** The potential harmful consequences of the development or use of AI that this instrument

Code	Definition	Example of text that would justify applying this code	Keywords that may signal code should be applied
1. Risk factors governed			
Bias	The instrument governs undesirable biases in the outputs of AI systems, including biases according to commonly protected classes such as race or gender.	<i>“It shall be unlawful for an employer or an employment agency to use an automated employment decision tool unless such tool has been the subject of a bias audit conducted no more than one year prior to the use of such tool.”</i>	<i>Bias, fairness, discrimination, unequal</i>
2. Harms addressed			
Harm to physical health/safety	This includes death, injury, or a reduction in lifespan.	<i>“Clinical performance testing must demonstrate that the device performs as intended under anticipated conditions of use, including detection of gastrointestinal lesions and evaluation of all adverse events.”</i>	<i>Injury, death, disability, public health, disease, illness, health outcome</i>
3. Governance strategies			
Evaluation	Requiring, encouraging, etc. the systematic evaluation of AI systems, or of broader systems or processes into which AI is directly integrated.	<i>“The Secretary of Defense shall establish performance objectives and accompanying metrics for the incorporation of artificial intelligence and digital readiness into such platforms, processes, and operations. . . ”</i>	<i>Audit, assessment, assess, impact assessment, red team, test, evaluate, metric, measure, monitor</i>
4. Incentives for compliance			
Access to business opportunities	Includes provisions that create or otherwise address advantages in obtaining business opportunities for people or organizations who comply with the instrument’s provisions (or withhold such advantages for those who do not comply), such as the right to do business in general, preferred or exclusive access to public tenders, etc.	<i>“Each unit shall structure procurement procedures, consistent with the purposes of this subtitle, to try to achieve an overall percentage goal of the unit’s total dollar value of artificial intelligence procurement contracts being made directly or indirectly to certified enterprises.”</i>	<i>Procure, public tender, request for proposals, public procurement, certified business, certified enterprise, conduct business, transact, license, approved provider, certified provider</i>
5. Application domains addressed			
Transportation	Includes autonomous vehicles, aerospace, avionics, and related components. This category also includes AI-enabled unmanned aerial vehicles; drones; and mobile robots for logistics and warehousing.	<i>“The Commandant, acting through the Blue Technology Center of Expertise, shall regularly assess available unmanned maritime systems and satellite vessel tracking technologies for potential use to support missions of the Coast Guard.”</i>	<i>Aircraft, vehicle, transportation, drone, train, transit, aerospace, flight, avionics, logistics, warehousing</i>

Table 1: Selected codes, definitions, examples, and keywords included in the AGORA codebook.

means to prevent (Altman, Wood, and Vayena 2018; NIST 2023; Slaughter, Kopec, and Batal 2021). Per AGORA’s definitions, “harms” are the consequences of “risks” (prior bullet). For example, an AI system that is insecure or biased (risk characteristics) might end up causing harm to physical health, human rights, or financial loss (harms). AGORA’s harm codes are adapted from categories of harm developed by CSET researchers for use with the AI Incident Database (Hoffmann and Frase 2023; McGregor 2021).

- **Governance strategies:** The means provided in the instrument to address, assess, or otherwise act with respect to the development, deployment, or use of AI. These essentially constitute the proposed solutions to the policy problems or goals articulated (Borrás and Edler 2014; Bullock et al. 2022). They comprise the largest group of codes in AGORA, reflecting the wide range of regulatory tools and tactics being used to address AI and its challenges (Schiff 2023; Johnson and Bowman 2021), from disclosure and evaluation requirements to conven-

ing, institutional creation, and pilot and testbed creation. AGORA’s governance strategy codes were developed iteratively by the AGORA leadership team over the first several months of annotation.

- **Incentives for compliance:** The types of incentives provided for people, organizations, etc., to comply with the requirements of the instrument (Gutierrez 2021). This small group of codes covers positive and negative incentives commonly seen in AI-related statutes and regulations, such as subsidies and fines, respectively.
- **Application domains addressed:** Any economic or social sectors, such as healthcare or defense, specifically addressed in the instrument as contexts for AI development, deployment, or use. AGORA’s application codes are adapted from the TINA industry taxonomy used in prior analyses of AI-related investment (Arcesati et al. 2023; Arnold, Rahkovsky, and Huang 2020) and are comparable in granularity to 2-digit NAICS codes (US Census Bureau 2022), with one exception: given significant attention to the use of AI by governments in particular (Medaglia, Gil-Garcia, and Pardo 2023; Wirtz, Weyerer, and Geyer 2019; White House 2020), the code for government applications of AI is divided into subcategories, allowing for more precise analysis.

As of publication, these categories encompass 77 different codes in total. To facilitate consistent annotation of abstract concepts, most codes are given detailed definitions. Annotators are also given examples of instrument text meeting more complex coding definitions as well as related keywords, exceptions, and considerations influencing interpretation. Selections from each category are provided in Table 1, with complete definitions and examples available online (see appendix).

Methodology and Workflow

This section describes the processes and concepts according to which instruments are collected, screened, and processed for inclusion in AGORA, as depicted in Figure 1.

Collection, Screening, and Compilation

Candidate documents for inclusion in AGORA are currently collected manually or using semi-automated means (e.g., saved queries against larger datasets) from a wide range of official and unofficial sources, reflecting the decentralized, largely ad-hoc status of current AI governance tracking. These sources include:

- Official, general-purpose regulatory compilations, such as the Congress.gov service for United States federal legislation and the Federal Register for United States federal regulation, and comparable subnational sources.
- Unofficial compilations of law and policy relevant to AI, digital issues, or related topics. Examples include the International Association of Privacy Professionals (IAPP) Global AI Law and Policy Tracker, the OECD AI Policy Observatory, and the OCEANIS collection of standards.

- Informal lists compiled by researchers, typically focused on particular topics or scopes of interest, such as criminal law, frontier model governance, ethical frameworks, industry standards, or Chinese AI regulation.

Human screeners review these sources (periodically, in the case of sources that update) and assess the instruments in them against the AGORA scoping definition and screening instructions. Screening decisions are made by single annotators, who are encouraged to consult with the broader AGORA team in difficult cases.³

For instruments judged to be in scope, screeners locate the authoritative text of the instrument (for example, on the official website of the United States Congress or a state legislature) and use it to populate basic metadata such as title and date of introduction. Annotators also identify “packages,” or larger, thematically diverse instruments containing AI-related portions amidst other, AI-unrelated material. A typical example is the annual National Defense Authorization Act (NDAA) in the United States, a massive, largely AI-unrelated law with some diverse AI-related provisions sprinkled throughout in recent years. NDAAAs and other such packages are divided into conceptually discrete AGORA instruments, corresponding to sections, subsections, or other subdivisions in the packages, according to standing guidance (see appendix).

Annotation and Validation

Using the basic metadata and authoritative text compiled during the screening process, AGORA annotators generate summaries and thematic codes for each in-scope instrument using the AGORA taxonomy (see Table 1 above) and further instructions provided in the AGORA codebook (see appendix). A custom-built Airtable interface structures the annotator workflow and facilitates quick and accurate annotation.

AGORA’s summaries are meant mainly to help users skim and sift, rather than as an analytic resource in themselves; the codebook provides brief instructions for short- and long-form summaries, but significant discretion is left to annotators.

The thematic coding process is more constrained. Annotators read each instrument in full, and then decide whether each of the 77 codes in the AGORA taxonomy applies at any point in the instrument,⁴ based on the definition and (where available) examples and keywords provided in the

³If capacity permits in the future, we hope to move to double screening of prospective AGORA records, consistent with our double annotation/validation process, described below.

⁴As of July 2024, AGORA’s taxonomy applied to instruments as a whole. Longer or more thematically diverse instruments may have many associated tags, and it may not be immediately clear which parts of the instruments justify which tags. Annotators were tagging some of these instruments at the section or subsection level as of July 2024, but the resulting granular data were not yet integrated into the AGORA interface and public dataset. Over time, we plan to increase the proportion of instruments coded at the section or subsection level and integrate the outputs into the AGORA interface and dataset.

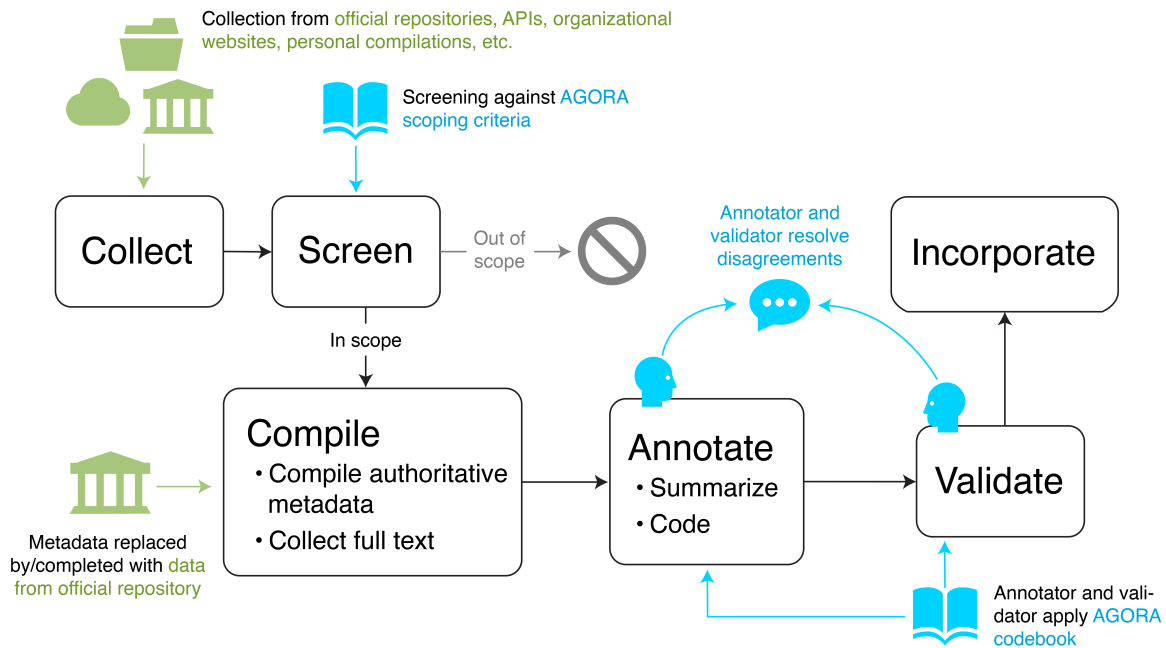


Figure 1: Data collection and analysis workflow

codebook. The codes in the taxonomy are described generally, and examples provided, above (see “Taxonomizing AI Governance”); full definitions of the 77 codes are available online (see appendix).

In deciding which tags apply, annotators are instructed to consider only the operative text of each instrument; to focus on what the instrument explicitly states or clearly and directly implies; and to ignore material unrelated to artificial intelligence. (For elaboration on each of these constraints, refer to the relevant portions of the codebook, available online (see appendix)).

After initial annotation, a second annotator (designated the “validator”) reviews each instrument in full and discusses any disagreements with the initial annotator. Note that these are not fixed roles; each AGORA annotator serves as initial annotator on some instruments and as validator on others, minimizing the chance of certain systematic biases.

To assess the reliability of the annotation process, 21 AGORA instruments, covering a range of instrument types and jurisdictions, were assigned to two different and non-overlapping annotator-validator pairs, resulting in two different sets of all 77 unique thematic codes for each instrument. For this test set (highly skewed distributions, nominal data, two coders), Guilford’s *g* (Holley and Guilford 1964) was 91.4 (ranging from 71.4 to 1 for different codes) and Gwet’s AC1 (Gwet 2008) was .83 (ranging from .43 to 1), indicating variable but generally substantial interrater agreement across different codes in the AGORA taxonomy.

Output and Distribution

Once all issues identified in validation have been resolved, the instrument’s record, consisting of validated metadata,

short and long summaries, and thematic codes, is marked complete. The schema for a complete, public AGORA record is included in the appendices.

Complete records are periodically added to the public AGORA dataset and web interface using an automated script. The web interface, depicted in Figure 2, is designed for efficient browsing and includes full text of instruments.

Current State of AGORA

As of July 2024,⁵ AGORA included 168 fully annotated instruments, plus 19 screened but unannotated “packages” (discussed above). An additional 278 instruments had been screened, determined to be in scope, and were in the annotation process. New instruments are added regularly as they clear the screening process.

Initial collection efforts focused on recent, adopted United States federal laws and regulations, making those instruments overrepresented initially; 140 of the 168 fully annotated instruments fully annotated as of July 2024 were U.S. federal laws or regulations, largely from 2021 or later, or sections of such documents. The other instruments included regulations, statutes, and policy documents issued by U.S. states and localities, the Chinese central government, and the European Commission, indicating the range of source material amenable to AGORA’s methodology and taxonomy. The median number of thematic codes applied to a fully annotated AGORA instrument was 8 (range: 0 to 43; average: 10.2; calculated May 2024).

⁵For current data, visit <https://agora.eto.tech>.

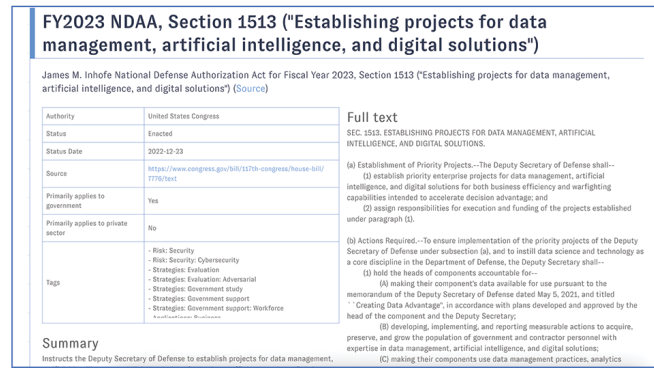
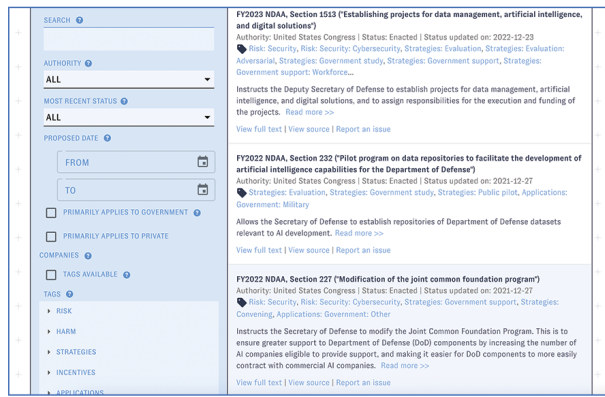


Figure 2: Screenshots from the AGORA web interface as of July 2024: browsing view (left) and instrument detail view (right)

Using AGORA: Implications for Academics, Policymakers, and Developers

Though it may be useful to anyone with an interest in AI, AGORA was developed with three primary user populations in mind: scholars, governance practitioners, and AI developers.

A core goal of the project is to support scholars and other analysts in understanding trends in, features of, and impacts of AI governance. The current version of AGORA already covers some analytically meaningful groupings in their entirety, including all AI-relevant United States federal laws and regulations enacted from 2020 to 2023. Over time, AGORA's ongoing, cumulative regulatory data collection process will allow increasingly general insights into the pace, diffusion, substantive focus, and impact of AI governance efforts globally, as well as comparisons across jurisdictions or topics. Beyond direct trends then, the data enable insights into deeper theoretical and ethical questions such as the nature of power and policy influence (Khanal, Zhang, and Taihagh 2024), the role of actors in the policy process (Arts 2003), the translation of social or ethical values into governance (Ulnicane and Aden 2023), and unique features of the governance of emerging technology (Leung 2020; Schiff 2022).

Critically, AGORA provides both structured data (instrument-level thematic codes according to the AGORA concept taxonomy) and unstructured data (document text) as raw material for such analyses, supporting a wide range of analytic methods. These include quantitative methods (regression analyses, time series analyses, quasi-experimental designs, quantitative content analysis (e.g., Neill et al. 2017; Parinandi et al. 2024; Schiff 2024) and qualitative approaches (process tracing, case studies, thematic analysis, discourse analysis (e.g., Voltolini and Eising 2017; DeLeon 1998; af Malmberg and Trondal 2023)). As a result, AGORA is well suited to serve as a keystone tool for a broad range of research questions, such as those listed in Table 2.

AGORA is also designed for AI governance practitioners,

such as legislators, regulatory agencies, and state and local officials grappling with the opportunities and challenges of increasingly powerful AI. This is an especially important use given the technical, conceptual, and interpretive complexity associated with AI. While policymakers at all scales of government have been trying to build capacity and expertise in AI (Lee 2022; Shrum et al. 2019), the challenge is likely to remain, and applied and accessible tools will be indispensable for evidence-based policymaking (Frisch and Kelly 2003; Umbach 2020). For these users, AGORA provides an archive of precedents and potential source material for new AI regulations and a way to monitor regulatory developments in peer jurisdictions or at different levels of government.

Finally, AGORA offers valuable data for AI researchers, developers, and deployers subject to regulation, as well as second- and third-party auditors, journalists, and members of civil society interested in evaluating the practical and ethical limitations of existing and proposed AI guardrails (Raji 2022; Samoili et al. 2020). These actors can use AGORA to identify regulations that apply to them or affect their work, assess regulatory trends and strategies emerging elsewhere that could eventually spread to their local jurisdictions, or engage in analysis or advocacy related to perceived gaps in AI governance.

Conclusion

Limitations

AGORA has important limitations. When practicable, mitigations have been added to AGORA guidance and processes.

First, AGORA's focus on instruments that directly address AI means that many AI-relevant laws, regulations, and norms are excluded. In particular, laws of general applicability—for example, securities regulation (FINRA 2024), civil rights law (Andrews and Bucher 2022; Kim 2022), or even common-law tort doctrines (?Sharkey 2024)—have implications for modern AI. However, this body of often older, AI-relevant general law is a potentially

Research Domains	Example Questions
Trends and basic descriptive questions	<ul style="list-style-type: none"> • Do AI governance instruments typically cover the public or private sector? • Which social and economic sectors (e.g., healthcare, military) are commonly addressed or neglected? • What are the prevalent and neglected harm categories, risk categories, and governance instruments reflected in AI governance documents? • How have priorities reflected in governance instruments changed over time?
Policy formulation, legitimation, adoption, and diffusion	<ul style="list-style-type: none"> • To what extent are policy problems coupled with relevant policy solutions, including monitoring, enforcement, or incentive strategies? • What factors are associated with the diffusion of problem framings (e.g., risks or harms) or policy solutions over time and across levels of government or regions? • What might explain the adoption, or lack thereof, of proposed policies?
Policy implementation and evaluation	<ul style="list-style-type: none"> • Are policy instruments with monitoring, evaluation, or enforcement provisions successfully implemented? • What accounts for the success or failure of policy implementation?
Comparative questions	<ul style="list-style-type: none"> • How does the underlying understanding of AI policy problems and governance solutions vary across sectoral or institutional contexts, including different levels of government or across countries and regions? • How do different types of documents (e.g., standards versus laws) vary in terms of their formality, specificity, or normative force? • To what extent is there convergence or divergence across governance instruments, and why?

Table 2: Candidate research facilitated by AGORA

unbounded set, and the applicability of any particular instrument may be debated. Accordingly, we exclude these instruments from AGORA’s scope. Though this makes AGORA a necessarily incomplete guide to the overall AI governance landscape, it reduces uncertainty over the bounds of the resource. Also, AI-focused implementations of more general laws—such as agency regulations or guidance explaining how broad existing authorities will be applied to AI—are in scope for AGORA, reducing the impact of the exclusion. Nonetheless, this limitation means AGORA cannot substitute for careful region-, context-, and sector-specific analysis of regulations relevant to particular groups, organizations, or

AI use cases, though it can enhance such analysis.

Second, the AGORA screening and annotation processes require substantial human judgment and effort, and are thus relatively time-intensive, limiting throughput. AGORA’s nominal scope is far broader than the set of documents collected to date. This includes important limitations in terms of AGORA’s regional and linguistic coverage. We plan to experiment with automation (discussed below) and seek more resources to increase processing volume and multi-language coverage. In the meantime, we have focused annotators’ efforts on especially high-profile or consequential instruments (e.g., adopted rather than proposed policies) and on sets of instruments of particular interest to current AGORA stakeholders.

Third, annotators exercise judgment in applying AGORA’s scope and taxonomy, raising the risk of inconsistency emerging from researcher (or more fundamental) subjectivity. To improve reliability and rigor, we provide detailed conceptual definitions, along with examples and decision heuristics. Further, we implement a dual annotation process (initial annotation followed by validation and reconciliation) for each instrument, with disagreements and other issues elevated to AGORA leadership and resolved according to defined procedures. Finally, we maintain a searchable central repository of prior questions and answers and make it accessible to all AGORA annotators online.

Extending and Expanding AGORA

Moving forward, AGORA’s annotation team will continue to screen and annotate instruments for the dataset. Based on their experience, we will continue to refine AGORA’s scoping criteria and taxonomy. Priorities for annotation include achieving, and thereafter maintaining, comprehensive coverage of all enacted and proposed U.S. federal and state statutes within AGORA’s scope, even as we increase the number and diversity of instruments from other jurisdictions.

In parallel, to improve throughput and potentially annotation consistency, we plan to continue experimenting with large language models to summarize and annotate instruments according to the AGORA taxonomy, building on recent research into using LLMs for complex concept identification and similar processing tasks involving law and social science texts (Ziems et al. 2024; Shen et al. 2023). We may also integrate AGORA screening and metadata collection with regulatory APIs as available, including the Congress.gov API for federal legislation and the Plural API for state laws.

AI governance and regulation are still in their early days, with growing volume, complexity and variation posing challenges to public awareness and understanding. The AGORA database and public interface aim to help minimize some of these challenges, and will adapt in turn to meet the needs of scholars, developers, policymakers, and other stakeholders interested in responsible and ethical governance of AI systems.

Ethics Statement

The AGORA dataset is derived entirely from public documents produced by governments and other large organizations. It does not include any data specific to individual people or vulnerable groups. As such, the AGORA team does not foresee substantial ethical harms from the development and sharing of AGORA beyond those usually associated with policy repositories and datasets.

However, we are mindful of several potential ethical concerns. First, AGORA’s coverage is not comprehensive. Although data compilation is ongoing, at present (and for the foreseeable future) the content of the dataset skews toward statutory documents from the United States. Although we have applied the AGORA taxonomy and data collection workflow to documents from other jurisdictions and communities and are confident these processes will be compatible with more diverse AI governance data, we have yet to fully test this hypothesis. In the meantime, it will be critical to stress to AGORA users that the resource is not comprehensive, and in what ways.

Moreover, many informal and all unwritten mechanisms and practices of AI governance are omitted from AGORA altogether, as the dataset includes only written instruments. Relying *solely* on AGORA to understand the AI governance landscape would therefore give an incomplete picture that could, for example, neglect essential aspects of governance ranging from implicit norms to educational practices, exclude AI governance activity in communities with less formalized practices, or distort understanding of how AI governance works in organizations or situations where unwritten rules and power dynamics tend to supplant formal codes. We stress that AGORA cannot account for the entirety of AI governance, and caution that overreliance on formal and written instruments may undermine robust understanding, harming vulnerable groups and neglecting essential aspects of ethical AI practice.

Another key consideration is that the AGORA interface is provided in English and currently focuses on English-language documents, potentially leading to unequal access to this resource for some stakeholders (representing a known challenge in academia and in AI governance and ethics specifically). We intend to expand the regional and linguistic coverage of AGORA, provided sufficient partnerships and resources. We also welcome outreach from diverse stakeholders on how AGORA is or is not meeting their needs, any inaccuracies, or harms that the team has not anticipated so that we may address them.

On balance, it is our hope that the tool, developed by an interdisciplinary team with significant focus on ethical issues, can raise awareness of possible ethical risks, harms, and enable study of the successes or failures of proposed governance solutions and make high-quality data available to a broader range of AI governance stakeholders. To further mitigate potential user misunderstandings that could lead to harm, we document AGORA’s limitations, and discuss use cases and research questions for which it is *inappropriate*, in the AGORA documentation, available at <https://eto.tech/tool-docs/agora>.

Positionality Statement

While the authors come from a range of backgrounds and have diverse regional focus areas, all authors are currently affiliated with United States academic institutions and several have focused on United States legal and political issues in relevant prior work and training. This perspective shaped the AGORA taxonomy and initial data collection efforts, which skewed heavily toward documents from the United States. We have attempted to design AGORA to capture data and insight from a broad range of jurisdictions and institutional types. Nonetheless, we acknowledge that AI governance scholars and stakeholders from different backgrounds and disciplines may have focused on different types of governance documents, different jurisdictions, or different AI-related themes and concerns in building a similar dataset. We plan to solicit and incorporate feedback from a broad range of AGORA users as the resource develops.

Adverse Impact Statement

Because AGORA only includes public-domain, non-experimental data of a non-personal nature, we foresee little adverse impact from compiling or sharing the dataset. Our primary adverse impact concerns relate to potential misunderstanding of the dataset and its limitations, as explored above.

Acknowledgements

We are grateful for the contributions of our talented team of undergraduate and graduate researchers, including students at Georgetown University and Purdue University. As of July 2024, they included Zoe Borden, Eileen Chen, Ogadinma Enwereazu, Ari Filler, Diya Kalavala, Mallika Keralavarma, Kieran Lee, Sophia Lu, Anjali Nookala, Maya Snyder, Jayan Srinivasan, Alina Thai, and Julio Wang. Lilan Huang and Mina Naryanan also provided helpful input.

Appendix

Visit AGORA’s public documentation, available at <https://eto.tech/tool-docs/agora>, to access:

- **AGORA screening instructions:** AGORA annotators use this document to determine whether documents are within AGORA’s scope and to guide collection of basic metadata for in-scope documents.
- **AGORA codebook:** The codebook defines and explains each of the 77 thematic codes in the AGORA taxonomy.
- **AGORA schema:** The schema for fully annotated, published AGORA records, including field descriptions.

References

- af Malmborg, F.; and Trondal, J. 2023. Discursive Framing and Organizational Venues: Mechanisms of Artificial Intelligence Policy Adoption. *International Review of Administrative Sciences*, 89(1): 39–58.
- Altman, M.; Wood, A.; and Vayena, E. 2018. A Harm-Reduction Framework for Algorithmic Fairness. *IEEE Security & Privacy*, 16(3): 34–45.

- Andrews, L.; and Bucher, H. 2022. Automating Discrimination: AI Hiring Practices and Gender Inequality. *Cardozo Law Review*, 44: 145.
- Arcesati, R.; Chang, W.; Hmaid, A.; and von Carnap, K. 2023. AI Entanglements: Balancing Risks and Rewards of European-Chinese Collaboration. Technical report, Mercator Institute for China Studies. <https://merics.org/en/report/ai-entanglements-balancing-risks-and-rewards-european-chinese-collaboration>.
- Arnold, Z.; Rahkovsky, I.; and Huang, T. 2020. Tracking AI Investment: Initial Findings from the Private Markets. Technical report, Center for Security and Emerging Technology, Washington, D.C. <https://cset.georgetown.edu/research/tracking-ai-investment>.
- Arts, B. 2003. Non-State Actors in Global Governance: Three Faces of Power. Working Paper 2003/4, Preprints aus der Max-Planck-Projektgruppe Recht der Gemeinschaftsgüter. <https://www.econstor.eu/handle/10419/85112>.
- Barrett, A. M.; Hendrycks, D.; Newman, J.; and Nonnecke, B. 2023. Actionable Guidance for High-Consequence AI Risk Management: Towards Standards Addressing AI Catastrophic Risks. arXiv:2206.08966.
- Bengio, Y.; Hinton, G.; Yao, A.; Song, D.; Abbeel, P.; Harari, Y. N.; Zhang, Y.-Q.; Xue, L.; Shalev-Shwartz, S.; Hadfield, G.; Clune, J.; Maharaj, T.; Hutter, F.; Baydin, A. G.; McIlraith, S.; Gao, Q.; Acharya, A.; Krueger, D.; Dragan, A.; Torr, P.; Russell, S.; Kahneman, D.; Brauner, J.; and Mindermann, S. 2023. Managing AI Risks in an Era of Rapid Progress. arXiv:2310.17688.
- Borrás, S.; and Edler, J., eds. 2014. *The Governance of Socio-Technical Systems: Explaining Change*. EU-SPRI Forum on Science, Technology and Innovation Policy Series. Cheltenham, UK: Edward Elgar. ISBN 978-1-78471-018-7.
- Bullock, J. B.; Chen, Y.-C.; Himmelreich, J.; Hudson, V. M.; Korinek, A.; Young, M. M.; and Zhang, B., eds. 2022. *The Oxford Handbook of AI Governance*. Oxford University Press, 1 edition. ISBN 978-0-19-757932-9.
- Cihon, P.; Maas, M. M.; and Kemp, L. 2020. Fragmentation and the Future: Investigating Architectures for International AI Governance. *Global Policy*, 11(5): 545–556.
- DeLeon, P. 1998. Models of Policy Discourse: Insights versus Prediction. *Policy Studies Journal*, 26(1): 147–161.
- European Commission. 2021. Proposal for a Regulation on a European Approach for Artificial Intelligence. Technical Report COM(2021) 206 final, European Commission, Brussels, Belgium. <https://digital-strategy.ec.europa.eu/en/library/proposal-regulation-european-approach-artificial-intelligence>.
- FINRA. 2024. 2024 FINRA Annual Regulatory Oversight Report. Technical report, FINRA, Washington, D.C. <https://www.finra.org/sites/default/files/2024-01/2024-annual-regulatory-oversight-report.pdf>.
- Frisch, S. A.; and Kelly, S. Q. 2003. Don't Have the Data? Make Them Up! Congressional Archives as Untapped Data Sources. *PS: Political Science & Politics*, 36(2): 221–224.
- Gutierrez, C. I. 2021. Uncovering Incentives for Implementing AI Governance Programs: Evidence from the Field. SSRN Scholarly Paper 3897486, Social Science Research Network, Rochester, NY.
- Gwet, K. L. 2008. Computing Inter-Rater Reliability and Its Variance in the Presence of High Agreement. *British Journal of Mathematical and Statistical Psychology*, 61(1): 29–48.
- Hoffmann, M.; and Frase, H. 2023. Adding Structure to AI Harm. Technical report, Center for Security and Emerging Technology. <https://cset.georgetown.edu/publication/adding-structure-to-ai-harm>.
- Holley, J.; and Guilford, J. 1964. A Note on the G Index of Agreement. *Educational and Psychological Measurement*, 24(4): 749–753.
- Johnson, W. G.; and Bowman, D. M. 2021. A Survey of Instruments and Institutions Available for the Global Governance of Artificial Intelligence. *IEEE Technology and Society Magazine*, 40(4): 68–76.
- Khanal, S.; Zhang, H.; and Taeihagh, A. 2024. Why and How Is the Power of Big Tech Increasing in the Policy Process? The Case of Generative AI. *Policy and Society*, pae012.
- Kim, P. T. 2022. Race-Aware Algorithms: Fairness, Nondiscrimination and Affirmative Action. *California Law Review*, 110: 1539.
- Lee, N. 2022. Do Policy Makers Listen to Experts? Evidence from a National Survey of Local and State Policy Makers. *American Political Science Review*, 116(2): 677–688.
- Leung, J. 2020. *Who Will Govern Artificial Intelligence?* Ph.D. thesis, Oxford University, Oxford, England. <https://ora.ox.ac.uk/objects/uuid:ea3c7cb8-2464-45f1-a47c-c7b568f27665>.
- Maas, M. 2023. International AI Institutions: A Literature Review of Models, Examples, and Proposals. SSRN Scholarly Paper 4579773, Social Science Research Network, Rochester, NY.
- McGregor, S. 2021. Preventing Repeated Real World AI Failures by Cataloging Incidents: The AI Incident Database. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(17): 15458–15463.
- Medaglia, R.; Gil-Garcia, J. R.; and Pardo, T. A. 2023. Artificial Intelligence in Government: Taking Stock and Moving Forward. *Social Science Computer Review*, 41(1): 123–140.
- Murdick, D.; Dunham, J.; and Melot, J. 2020. AI Definitions Affect Policymaking. Technical report, Center for Security and Emerging Technology, Washington, D.C. <https://cset.georgetown.edu/publication/ai-definitions-affect-policymaking>.
- Neill, J. O.; Robin, C.; Brien, L. O.; and Buitelaar, P. 2017. An Analysis of Topic Modelling for Legislative Texts. In *Proceedings of the Second Workshop on Automated Semantic Analysis of Information in Legal Text*. London, UK.
- NIST. 2023. AI Risk Management Framework: AI RMF (1.0). Technical Report error: NIST AI 100-1, National

- Institute of Standards and Technology, Gaithersburg, MD. <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf>.
- OECD. 2022. OECD Framework for the Classification of AI Systems: A Tool for Effective AI Policies. Technical Report 323, OECD. <https://oecd.ai/en/p/classification>.
- OECD. 2023. Stocktaking for the Development of an AI Incident Definition. OECD Artificial Intelligence Papers 4, OECD, Paris, France. https://www.oecd-ilibrary.org/science-and-technology/stocktaking-for-the-development-of-an-ai-incident-definition_c323ac71-en.
- Parinandi, S.; Crosson, J.; Peterson, K.; and Nadarevic, S. 2024. Investigating the Politics and Content of US State Artificial Intelligence Legislation. *Business and Politics*, 26(2): 240–262.
- Raji, I. D. 2022. From Algorithmic Audits to Actual Accountability: Overcoming Practical Roadblocks on the Path to Meaningful Audit Interventions for AI Governance. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '22, 5. New York, NY, USA: Association for Computing Machinery.
- Samoili, S.; López Cobo, M.; Gómez, E.; De Prato, G.; Martínez-Plumed, F.; Delipetrev, B.; European Commission; and Joint Research Centre. 2020. AI Watch: Defining Artificial Intelligence: Towards an Operational Definition and Taxonomy of Artificial Intelligence. Technical report, Publications Office of the European Union, Luxembourg. https://op.europa.eu/publication/manifestation_identifier/PUB_KJNA30117ENN.
- Schiff, D. S. 2022. Setting the Agenda for AI: Actors, Issues, and Influence in United States Artificial Intelligence Policy. <https://doi.org/10.17605/OSF.IO/KW8XD>.
- Schiff, D. S. 2023. Looking through a Policy Window with Tinted Glasses: Setting the Agenda for U.S. AI Policy. *Review of Policy Research*, 40(5): 729–756.
- Schiff, D. S. 2024. Framing Contestation and Public Influence on Policymakers: Evidence from US Artificial Intelligence Policy Discourse. *Policy and Society*, puae007.
- Sharkey, C. M. 2024. A Products Liability Framework for AI. SSRN Scholarly Paper 4773874, Social Science Research Network, Rochester, NY.
- Shen, S.; Liu, J.; Lin, L.; Huang, Y.; Zhang, L.; Liu, C.; Feng, Y.; and Wang, D. 2023. SsciBERT: A Pre-Trained Language Model for Social Science Texts. *Scientometrics*, 128(2): 1241–1263.
- Shrum, K.; Gordon, L.; Regan, P.; Maschino, K.; Shark, A. R.; and Shropshire, A. 2019. AI and Its Impact on Public Administration. Technical report, National Academy of Public Administration, Washington, D.C. https://www.napawash.org/uploads/Academy_Studies/9781733887106.pdf.
- Slaughter, R. K.; Kopec, J.; and Batal, M. 2021. Algorithms and Economic Justice: A Taxonomy of Harms and a Path Forward for the Federal Trade Commission. Technical report, Yale Journal of Law & Technology, New Haven, CT. https://law.yale.edu/sites/default/files/area/center/isp/documents/algorithms_and_economic_justice_master_final.pdf.
- Ulnicane, I.; and Aden, A. 2023. Power and Politics in Framing Bias in Artificial Intelligence Policy. *Review of Policy Research*, 40(5): 665–687.
- Umbach, G. 2020. Of Numbers, Narratives and Challenges: Data as Evidence in 21st Century Policy-Making. *Statistical Journal of the IAOS*, 36(4): 1043–1055.
- US Census Bureau. 2022. Economic Census: NAICS Codes & Understanding Industry Classification Systems. <https://www.census.gov/programs-surveys/economic-census/year/2022/guidance/understanding-naics.html>.
- Voltolini, B.; and Eising, R. 2017. Framing Processes and Lobbying in EU Foreign Policy: Case Study and Process-Tracing Methods. *European Political Science*, 16(3): 354–368.
- White House. 2020. Executive Order on Promoting the Use of Trustworthy Artificial Intelligence in the Federal Government. <https://www.govinfo.gov/content/pkg/DCPD-202000870/pdf/DCPD-202000870.pdf>.
- Wirtz, B. W.; Weyerer, J. C.; and Geyer, C. 2019. Artificial Intelligence and the Public Sector—Applications and Challenges. *International Journal of Public Administration*, 42(7): 596–615.
- Ziems, C.; Held, W.; Shaikh, O.; Chen, J.; Zhang, Z.; and Yang, D. 2024. Can Large Language Models Transform Computational Social Science? *Computational Linguistics*, 50(1): 237–291.