

From Documents to Signals:

Toward an AI-powered Policy Intelligence Framework

Masato Oda
OSINTech Inc.,
Kobe Institute of Computing
Kobe Japan
masato@osintech.net

Masamichi Kobayashi
OSINTech Inc.,
Graduate School of Societal
Safety Sciences, Kansai
University
k.masamichi@osintech.net

Kazue Oda
OSINTech Inc.
kazue@osintech.net

Tomoyuki Fuse
OSINTech Inc.
tomoyuki@osintech.net

Sho Taya
OSINTech Inc.
taya@osintech.net

Abstract

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

Information Society 2026, 5–9 October 2026, Ljubljana, Slovenia

© 2026 Copyright held by the owner/author(s).

DOI : <https://doi.org/10.70314/is.2026.ai4sdg.157>

Conventional policy research has struggled to identify weak and dispersed policy concerns systematically and at scale when they emerge across countries, languages, institutions, and policy domains. This study examines whether an AI-assisted approach can identify policy trends using primary policy sources collected and structured across national, linguistic, and organizational boundaries.

We used RuleWatcher, which covers multilingual primary policy sources published by governments in approximately 150 countries, as well as communications from organizations involved in rulemaking, including United Nations bodies and non-governmental organizations (NGOs). For the empirical analysis, we constructed a dataset of 2,753 records concerning AI ethics.

We divided the data into an earlier period (1 January 2019 to 31 December 2024) and a later period (1 January 2025 to 31 July 2026). From expressions found in the earlier-period documents, we extracted 132 keywords after consolidating synonymous terms. We then used AI-based semantic similarity assessment and author review to organize these keywords into 26 "Concept Cards." We counted the occurrence of expressions associated with each Concept Card by policy actor and quarter, and examined whether concepts detected in the earlier period could be treated as early indications of policy agenda formation when their frequency increased in the later period.

Of the 2,753 records analyzed, 1,116 contained expressions associated with at least one Concept Card.

Counting documents assigned to multiple concepts in each applicable category produced 2,671 concept-document assignments. Of the assignments observed in the earlier period, 65.6% came from non-governmental actors, including NGOs. The quarterly distribution also showed that references by non-governmental actors continued from the early part of the observation period. Across complete quarters, excluding the partial third quarter of 2026, the mean quarterly count for all actors increased from 77.3 in the earlier period to 127.3 in the later period. Comparison of the 26 concepts by actor and quarter revealed a temporal difference for multiple concepts: early references by non-governmental actors were followed by increased references from UN bodies and governments.

These findings suggest that AI-based classification and temporal analysis of multilingual, multi-actor primary policy sources can provide evidence for observing the formation of policy agendas. However, the observed temporal differences do not demonstrate causal diffusion between actors or policy adoption.

Based on these findings, this study proposes a "Governance Observatory" as an AI-assisted intelligence framework for continuously observing policy agendas as they emerge.

Humanity increasingly faces challenges that no single country can address alone. Perspectives dispersed across underrepresented regions and linguistic minorities may contribute to addressing such challenges. Yet the rapid adoption of AI in policymaking may introduce centralization bias and further marginalize communications from these regions and communities. To mitigate this risk, policymakers should consider using AI to monitor a broad range of primary policy sources continuously, including communications from non-governmental actors.

Keywords

AI-assisted policy analysis, policy agenda formation, Governance Observatory, horizon scanning, RuleWatcher

1. Background and Research Objectives

Policy is one of the principal means by which public institutions address societal challenges. However, existing policy frameworks may not respond adequately to complex global problems such as climate change and social instability caused by inequality. This limitation is particularly pronounced for rapidly advancing technologies such as AI: regulation and rulemaking often lag behind social change, while concerns that are currently weak and dispersed but may become important to future policymaking or the resolution of global challenges remain insufficiently recognized.

To respond to increasingly complex challenges in advance, policymakers need systematic ways to identify concerns that have not yet been clearly articulated as policy issues but are appearing sporadically, and to present them in a form that can inform policy decisions.

Hajer [1] and Ingram et al. [2] have shown how policy issues are formed through discourse and social meaning-making. Building on this perspective, this study focuses on dispersed references that appear before an issue becomes clearly established on the policy agenda.

Practical approaches to future-oriented policy analysis include strategic foresight [3], scenario planning [4], and horizon scanning. Horizon scanning, in particular, is regarded as useful for science, technology, and innovation (STI) policy and governance because it identifies "weak signals," functions as an early-warning system for emerging issues, and provides insights to policymakers [5].

Conventional horizon-scanning exercises often select observation targets through expert nomination or predefined information sources and classification frameworks. Depending on the research design, such approaches may overlook early concerns that do not fit established problem definitions or that originate from peripheral actors and languages.

This study uses AI to address a question that remains insufficiently examined: how can perceptions of policy problems dispersed across policy domains be observed before they are explicitly recognized as policy issues? We report the results of this investigation.

2. Research Questions

This study examines a general analytical approach for identifying policy agendas as they emerge from multilingual primary policy sources published by diverse actors. AI ethics is used as an empirical case to assess the feasibility of this approach. Accordingly, we address the following research questions:

RQ1. To what extent can AI-assisted conceptual classification and temporal analysis of multilingual primary policy sources reveal the early formation of policy agendas?

RQ2. When and how consistently do policy agendas appear in primary policy sources published by non-governmental actors, including NGOs and specialist organizations, UN bodies, and national governments? How do their frequency and temporal trajectories differ across policy actors?

RQ3. How can AI-based candidate detection be combined with human expert review to continuously identify indications of emerging policy agendas in multilingual primary policy sources published by diverse actors?

3. Method

We used RuleWatcher [6], which standardizes multilingual primary policy sources into a common analytical language and enables comparison across countries, regions, and organization types. Taking AI ethics as an empirical case, we assessed an analytical method for identifying early policy agenda formation and considered its broader applicability.

Overview of the RuleWatcher Data

Number of records: 1,742,102

Coverage: news releases from governments in 151 countries, 59 UN bodies, 904 NGOs and other non-governmental actors, and other sources

Publication period: 1 January 2019 to 31 July 2026

Multilingual primary policy sources standardized into English

17 themes and approximately 1,500 topic tags

The AI Governance theme uses a search framework based on UNESCO's Recommendation on the Ethics of Artificial Intelligence [7]

Dataset Construction

From 1,742,102 RuleWatcher records (a), we extracted 96,751 records (b) published by governments, UN bodies, and non-governmental actors and classified under the AI Governance theme.

From dataset (b), we extracted records assigned the predefined RuleWatcher topic tag "AI Ethics." After duplicate and relevance screening, 2,753 records (c) remained in the empirical dataset.

We classified dataset (c) by organization type into government (d), UN bodies (e), and other non-governmental actors (f).

For each organization type, we divided the records into the earlier period, 1 January 2019 to 31 December 2024 (g, h, i), and the later period, 1 January 2025 to 31 July 2026 (j, k, l).

Figure1: Dataset construction

4. Data Analysis

We divided corpora (g)-(l) into quarterly subsets and conducted a two-stage analysis using OpenAI's GPT-5.6-sol.

RuleWatcher retained the original-language records and standardized their textual content into English through its machine-translation pipeline. The present analysis was conducted on the standardized English descriptions. GPT-5.6-sol was applied using a fixed prompt across all quarterly subsets. Because generative-model outputs may vary between runs and model versions, the final list of 132 expressions and their mappings to the 26 Concept Cards were preserved as fixed analytical artifacts.

In the first stage, we extracted candidate expressions from the earlier-period data using AI” or “we used AI to extract candidate expressions from the earlier-period data. Expressions appearing at least twice within a document were evaluated using the following four criteria. After removing duplicates from the 500 highest-ranked expressions, 132 expressions were retained.

- Novelty: expressions rarely used in the past
- Growth: expressions whose frequency increased from the previous quarter
- Burst: expressions showing an anomalous increase over a short period
- TF-IDF: expressions frequent within individual documents but not widely distributed across the full dataset

The 132 expressions included synonymous and related terms that differed in wording but referred to a common problem concept. For example:

- AI environmental impact
- energy consumption of AI
- carbon footprint of AI
- data center emissions
- environmental cost of foundation models

Although phrased differently, these expressions can all be understood as referring to the common problem concept of the environmental impact of AI. We therefore combined AI-based semantic similarity assessment with author review to organize the 132 expressions into 26 semantic groups, which we define as “Concept Cards.”

In the second stage, we mapped the 132 expressions to their corresponding Concept Cards and identified documents containing at least one expression associated with each card. We then counted documents by Concept Card, policy actor, and quarter.

Within each combination of Concept Card, policy actor, and quarter, each document was counted only once. A document associated with multiple Concept Cards was counted once under each applicable card.

This analysis examined whether the frequency of problem concepts detected in the earlier period increased in the later period, and whether their timing, persistence, and temporal trajectories differed across policy actors.

Reproducibility and Validation Protocol

To ensure traceability, the analytical workflow was documented at each stage, including corpus extraction, quarterly segmentation, candidate-expression ranking, semantic consolidation, document matching, and aggregation. The reproducibility materials specify the dataset inclusion and exclusion criteria, the 132 candidate expressions and their mappings to the 26 Concept Cards, the prompts used for AI-assisted classification, the model configuration, and the rules used to count documents by concept, actor, and quarter. The Concept Card dictionary was fixed before its application to the later-period data, thereby preventing later observations from changing the classification framework.

AI-generated groupings were not accepted automatically. The authors reviewed whether the expressions assigned to each Concept Card referred to the same or a closely related policy problem, revised ambiguous groupings, and confirmed the final mappings. Exact keyword matching was then applied to the standardized English descriptions. Within each Concept Card, policy actor, and quarter, a document was counted only once, although a document could be assigned to multiple Concept Cards.

5. Results

Applying the 26 Concept Cards developed from the earlier-period data, we detected keywords associated with at least one card in 1,116 of the 2,753 documents (40.5%). Counting documents assigned to multiple concepts under each applicable card yielded 2,671 concept-document assignments.




Figure 2: Quarterly concept-document assignments by policy actor across 26 concepts

Trends in the Number of Documents. For 2026, the data cover only January–July.

Across the full observation period, non-governmental actors accounted for 1,705 concept-document assignments, governments for 557, and UN bodies for 409 (Figure 2). Excluding the partial third quarter of 2026, the mean quarterly count across all actors increased by 64.7%, from 77.3 in the earlier period to 127.3 in the later period. Counts increased for all three actor groups; the mean for UN bodies approximately doubled, from 10.7 to 22.0.




Figure 3: Heatmap of 26 concepts by policy actor and quarter

Figure 3 presents the frequency of the 26 concepts by policy actor and quarter. Aggregated across all actors, five concepts - Privacy; Human Oversight and Appeals; Biometrics and Remote Surveillance; Children and Youth; and AI, Elections, and Democracy - appeared in all 31 observed quarters. Responsible and Trustworthy AI and Algorithmic Discrimination appeared in 30 of the 31 quarters. These patterns indicate that the concepts were addressed consistently throughout the observation period rather than appearing as temporary topics.

By contrast, Generative AI and Foundation Models, Deepfakes and Content Authenticity, AI Readiness Assessment, and AI Literacy increased in frequency particularly after 2023. The AI ethics dataset thus contains both long-standing concerns and concerns that became more visible in association with recent technological and institutional changes.

We next compared the first quarter in which each concept was detected within the dataset across policy actors. Of the 21 concepts observed for both non-governmental actors and UN bodies, detection by non-governmental actors occurred earlier for 17 concepts, in the same quarter for three, and later for one. Of the 21 concepts observed for both non-governmental actors and governments, detection by non-governmental actors occurred earlier for 15 concepts, in the same quarter for five, and later for one.

These findings do not mean that non-governmental actors preceded other actors for every problem concept. Nevertheless, among the 21 comparable concepts, non-governmental actors were detected earlier for 17 concepts relative to UN bodies and for 15 relative to governments. Multiple problem concepts were therefore

observable in the primary policy sources of non-governmental actors at a comparatively early stage.

Here, "first detection" refers to the first quarter within the dataset's observation period in which at least one keyword associated with a Concept Card was detected. It does not identify when a problem concept was first proposed in practice, when it was institutionally adopted, or whether one policy actor directly influenced another. The results nevertheless indicate that continuous observation of primary policy sources from diverse actors may identify problem concepts before they become widely recognized as policy agendas.

6. Conclusion

This study examined whether AI-based candidate detection, conceptual classification, and temporal analysis of multilingual primary policy sources published by diverse actors could identify early indications of emerging policy agendas. In the empirical case of AI ethics, differently worded concerns were organized into common conceptual units, enabling their timing, frequency, and persistence to be visualized across actors. The findings demonstrate the potential of primary policy sources for observing temporal patterns in policy agenda formation.

Non-governmental actors, including NGOs, accounted for 65.6% of the concept-document assignments observed in the earlier period. Many problem concepts were detected in their primary policy sources before being detected in those of UN bodies or governments, after which counts increased across all actors. This temporal ordering does not demonstrate direct influence between actors or policy adoption. The findings may also be affected by differences in publication frequency, digital publishing infrastructure, RuleWatcher's coverage, machine-translation accuracy, and communication strategies across policy actors. Consequently, the observed counts do not directly measure the societal importance or actual prevalence of each problem concept.

Proposal: A Governance Observatory to Support Policymaking

Based on these findings, this study proposes the "Governance Observatory," an intelligence framework that continuously identifies indications of emerging policy agendas and presents them in a form that can inform policymakers' decisions.

As the use of AI in policy research expands, bias in the information selected for analysis may reproduce and amplify dominant regional, linguistic, and institutional values. The Governance Observatory therefore encompasses primary policy sources not only from governments and major international organizations, but also from NGOs, specialist organizations, actors in underrepresented regions, and linguistic minorities. By presenting weak and dispersed concerns in a form that

can be examined, the framework aims to deepen understanding of complex policy challenges and provide a foundation for more inclusive and well-grounded policymaking.

Future research should assess the reproducibility and generalizability of the framework by applying it to policy fields beyond AI ethics.

Data and Materials Availability

RuleWatcher is a proprietary platform, and the record-level corpus used in this study cannot be publicly redistributed because it contains materials collected from external primary sources. The authors have retained the corpus-selection criteria, the 132-expression dictionary, the mappings to the 26 Concept Cards, the AI prompt and model configuration, the validation protocol, and the aggregated quarterly counts underlying Figures 2 and 3. Requests for access to these derived research materials for scholarly verification may be directed to the corresponding author and will be considered subject to applicable third-party rights and platform terms. Information about RuleWatcher and its access options is available at <https://web.rulewatcher.com/>.

References / Literatura

- [1] Maarten A. Hajer. 1995. *The Politics of Environmental Discourse: Ecological Modernization and the Policy Process*. Oxford University Press, Oxford, UK.
- [2] Helen Ingram, Anne L. Schneider, and Peter deLeon. 2007. *Social Construction and Policy Design*. In Paul A. Sabatier (Ed.), *Theories of the Policy Process*, 2nd ed., 93-126. Westview Press, Boulder, CO.
- [3] OECD. 2019. *Strategic Foresight for Better Policies: Building Effective Governance in the Face of Uncertain Futures*. OECD Publishing, Paris, France.
- [4] Kees van der Heijden. 2005. *Scenarios: The Art of Strategic Conversation*, 2nd ed. John Wiley & Sons, Chichester, UK.
- [5] OECD. 2026. *Building Capacity for Technology Horizon Scanning: A Guide for Policymakers*. OECD Science, Technology and Industry Working Papers, No. 2026/06. OECD Publishing, Paris, France.
- [6] OSINTech Inc. 2026. RuleWatcher: A global policy data tool. It offers extensive coverage of developing countries and includes a wealth of information from non-governmental entities. By providing local-language content in English, it directs users to primary sources. Users can also download structured data categorized by over 1,600 expert-selected tags. Access at: <https://web.rulewatcher.com/> (Accessed September 29, 2026).
- [7] UNESCO. 2022. *Recommendation on the Ethics of Artificial Intelligence*. United Nations Educational, Scientific and Cultural Organization, Paris, France.