



SCIENTIFIC OASIS

Decision Making: Applications in Management and Engineering

Journal homepage: www.dmame-journal.org
ISSN: 2560-6018, eISSN: 2620-0104

Measuring Compliance of AI-Generated Content Platforms with Data Protection Regulations: A Computational Framework and Empirical Analysis of Privacy Policy Disclosures

Jiangfeng Hao^{1,*}¹ School of Law, Taiyuan University of Science and Technology, Taiyuan 030024, China

ARTICLE INFO

Article history:

Received 14 April 2026

Received in revised form 12 June 2026

Accepted 17 June 2026

Available online 30 June 2026

Keywords:

AI-Generated Content; Data Protection Compliance; Privacy Policy Analysis; Regulatory Technology; Natural Language Processing; GDPR; PIPL; CCPA

ABSTRACT

The rapid expansion of AI-generated content (AIGC) platforms, encompassing large language models, text-to-image generators, and code assistants, has intensified concerns regarding whether their privacy disclosures satisfy major data protection regimes, including the EU General Data Protection Regulation (GDPR), China's Personal Information Protection Law (PIPL), and California's Consumer Privacy Act (CCPA). Despite these concerns, there remains no systematic and reproducible approach for determining the completeness and accuracy of platform privacy policies against statutory disclosure requirements. To address this gap, this paper introduces PolicyLens, a computational framework for compliance assessment that: (1) develops a machine-readable regulatory ontology representing 127 distinct disclosure obligations derived from GDPR, PIPL, and CCPA provisions; (2) uses an LLM-based information extraction pipeline to transform privacy policy content into structured evidence of compliance; and (3) aligns the extracted evidence with the ontology to generate multi-dimensional compliance scores. PolicyLens is evaluated using AIGC-Privacy-200, a newly compiled dataset containing the complete privacy policies and terms of service of 200 major AIGC platforms collected between January and March 2026. The dataset covers services headquartered across the US, EU, China, and other jurisdictions. The results show that: (i) only 23% of the platforms fully satisfy the disclosure requirements of any one regulation; (ii) obligations concerning training data provenance and opt-out mechanisms are disclosed least frequently, with compliance rates of 11% and 17%, respectively; and (iii) compared with EU-based platforms, Chinese platforms demonstrate significantly greater compliance with PIPL but lower compliance with GDPR, indicating that privacy policies tend to be drafted around the regulatory requirements of their home jurisdictions. Validation against expert legal audits conducted on a 40-platform subset shows that PolicyLens achieves an F1-score of 0.89 for obligation-level compliance detection. To facilitate further research in regulatory technology, the framework, dataset, and regulatory ontology are released as open-source resources.

1. Introduction

The emergence of AIGC platforms has substantially reshaped the digital environment. LLM-

* Corresponding author.

E-mail address: haojiangfeng@tyust.edu.cn<https://doi.org/10.5281/zenodo.23207940>

powered services such as ChatGPT, Claude, and Gemini, alongside text-to-image generators including Midjourney, DALL-E, and Stable Diffusion and code assistants such as GitHub Copilot and Amazon CodeWhisperer, had collectively attracted more than two billion users worldwide by early 2026. Their operation involves processing substantial volumes of personal data, including user prompts, conversation histories, uploaded documents, and images. At the same time, their dependence on extensive training datasets has introduced new concerns surrounding the provenance of training data and its legal status [1].

Regulatory frameworks for data protection have faced considerable challenges in responding to this rapidly evolving environment. GDPR, PIPL, and CCPA impose extensive transparency requirements on data controllers and processors, covering matters such as the purposes of data collection, lawful bases for processing, data subject rights, cross-border data transfers, and automated decision-making practices [2]. Meeting these requirements is particularly complex for AIGC platforms because personal data may be involved at two distinct stages: during training, when information may be incorporated from web-scale datasets, and during inference, when interactions with users generate additional personal data [3]. Although regulatory compliance is increasingly important for AIGC platforms, its assessment continues to rely largely on manual procedures. Legal audits require substantial time and resources, making them difficult to implement at scale across numerous platforms operating under different regulatory regimes. Existing computational methods for analysing privacy policies have primarily addressed conventional websites and mobile applications. Such approaches commonly rely on keyword-based techniques, rule-driven systems, or traditional machine learning classifiers, which may be insufficient for interpreting the contextual and legally nuanced language found in AIGC-related policies.

More recent developments in LLM-based information extraction have shown considerable potential for analysing legal documents, including annotation-based approaches that enable structured knowledge to be derived from complex textual materials [4]. Nevertheless, these capabilities have yet to be systematically employed to assess the regulatory compliance of AIGC platform privacy policies. To address this gap, this paper introduces PolicyLens, a computational framework developed specifically to assess AIGC platform privacy policies across multiple data protection regimes. The framework offers three principal contributions:

- (1) Regulatory Ontology Construction:** We construct a machine-readable ontology comprising 127 distinct disclosure obligations derived from GDPR (Articles 12-22, 44-49), PIPL (Articles 17, 23-25, 38-39), and CCPA (Sections 1798.100-1798.199). Each obligation is represented according to its regulatory origin, obligation type (inform, provide mechanism, document), specificity level, and corresponding mappings across regulatory regimes.
- (2) LLM-based Policy Analysis Pipeline:** We develop a three-stage information extraction pipeline in which privacy policy documents are segmented, individual clauses are mapped to the obligation ontology, and structured evidence of compliance is extracted. The pipeline employs fine-tuned language models together with domain-specific prompting strategies designed to support the interpretation of legal text.
- (3) Empirical Analysis of AIGC-Privacy-200:** We compiled and released a comprehensive dataset containing the complete privacy policies and terms of service of 200 major AIGC platforms. PolicyLens is then applied to this dataset to conduct a large-scale cross-jurisdictional assessment of disclosure compliance across the AIGC industry.

The findings indicate substantial compliance deficiencies throughout the AIGC industry. Among the 200 platforms examined, only 23% fully satisfy the disclosure requirements of any single regulatory framework. Compliance is particularly limited for emerging AIGC-specific obligations,

most notably disclosures concerning training data provenance and mechanisms enabling users to opt out. The analysis further reveals systematic jurisdiction-anchored drafting practices, whereby platforms place greater emphasis on satisfying the regulatory requirements of their home jurisdictions while giving comparatively less attention to applicable extraterritorial obligations. These findings carry important implications for regulators, platform developers, and users. The remainder of this paper is organised as follows. Section 2 examines prior research concerning privacy policy analysis, regulatory technology, and AI governance. Section 3 provides a detailed account of the PolicyLens framework, followed by the construction of the AIGC-Privacy-200 dataset in Section 4. The empirical findings are presented in Section 5, while Section 6 evaluates the framework against expert legal audits. Section 7 addresses the implications and limitations of the study, and Section 8 concludes.

2. Related Work

2.1 Computational Privacy Policy Analysis

Automated privacy policy analysis has received substantial attention within NLP and legal informatics. Early research primarily relied on keyword-based techniques to classify privacy policy paragraphs into predefined categories [5]. The subsequent introduction of the OPP-115 dataset established an important benchmark for privacy policy annotation and facilitated supervised learning using conventional classifiers, including support vector machines and random forests [6]. As computational methods advanced, transformer- and deep learning-based approaches became increasingly prominent. Polis combined convolutional neural networks with attention mechanisms to support automated privacy policy analysis [7]. PolicyQA adopted a different approach by framing compliance assessment as a question-answering task over privacy policy documents [8].

Despite these advances, existing approaches present two fundamental limitations when applied to AIGC platforms. First, they were primarily developed for conventional websites and mobile applications, whose data-processing activities differ substantially from those associated with generative AI systems [9]. AIGC platforms involve distinct data flows, including the acquisition of training material from web-scale corpora, retention of user prompts for model improvement, and generation of synthetic outputs that may contain personal information. These practices extend beyond the scope of existing privacy taxonomies [10]. Second, previous approaches generally assess compliance within a single regulatory framework [11]. Such an approach is insufficient for globally operating AIGC platforms, which may simultaneously be subject to multiple regulatory regimes with differing or potentially conflicting requirements.

2.2 Regulatory Technology and Legal AI

Regulatory technology (RegTech) has become an important application area for AI and NLP, with automated compliance systems increasingly used to assess requirements in financial regulation, healthcare standards, and environmental governance. Within legal AI, recent advances in LLM-based systems have strengthened capabilities for legal reasoning, document analysis, and regulatory interpretation. Domain-specific annotation platforms have further demonstrated that human expertise can be effectively integrated with machine learning to extract structured information from complex technical documents [4]. These established architectural principles provide a foundation for the approach adopted in this study to extract regulatory obligations. Rigorous evaluation is equally important for computational systems operating in high-stakes domains. Previous research has shown that weaknesses in computational evaluation metrics can

produce misleading assessments of system performance [12]. This issue is particularly relevant to compliance assessment, where false positives may incorrectly indicate that regulatory requirements have been satisfied and consequently carry significant legal implications. Accordingly, the methodological principles developed to address measurement bias in latency-sensitive language processing systems are incorporated into the design of the compliance scoring framework.

2.3 AI Governance and AIGC Regulation

The governance of AIGC has become an increasingly prominent concern across policy, industry, and academic communities. The EU AI Act (2024) introduced a risk-based framework for classifying AI systems and established specific transparency obligations applicable to generative AI [13]. In China, the Interim Measures for the Management of Generative AI Services (2023) set requirements concerning the legality of training data, content labelling, and user consent. Regulatory attention has also increased in the United States, where Executive Order 14110 on Safe, Secure, and Trustworthy AI (2023) instructed federal agencies to develop guidance addressing AI transparency [14].

Alongside these regulatory developments, the technical community has advanced frameworks intended to support responsible AI deployment. Research on safety-aware AI systems has demonstrated the value of embedding safety constraints within system architecture rather than addressing compliance only after deployment [15]. This principle of "compliance by design" underpins the present approach, in which automated compliance assessment is intended to facilitate continuous monitoring rather than reliance solely on one-time audits. Despite the expanding regulatory environment, empirical evidence concerning actual compliance among AIGC platforms remains limited. Existing survey-based research has largely examined users' privacy concerns, leaving platform-level regulatory compliance comparatively underexplored through systematic and data-driven methods. This study addresses that gap by providing a comprehensive computational assessment of AIGC platform privacy policy compliance across multiple jurisdictions and regulatory frameworks.

2.4 Natural Language Processing for Legal Text Analysis

The use of NLP for legal text analysis has advanced considerably over the past decade. Earlier methods typically relied on bag-of-words representations and manually engineered features. Although these approaches performed adequately on clearly defined classification tasks, their ability to capture the semantic complexity of legal language remained limited. The emergence of transformer-based architectures, particularly BERT and its variants, substantially improved this capability. LEGAL-BERT, a domain-specific pre-trained model, achieved consistent performance improvements over general-purpose alternatives across legal NLP benchmarks [16]. The practical relevance of NLP to consumer protection was further demonstrated through CLAUDETTE, which automated the identification of potentially unfair clauses in online terms of service [17].

The subsequent development of LLMs has further expanded the analytical possibilities for legal documents. These models can support zero-shot and few-shot classification of legal provisions, produce structured representations of complex documents, and identify relationships and cross-references across legal texts. Their use for regulatory compliance assessment, however, creates distinct methodological challenges. Compliance analysis requires outputs that are deterministic and reproducible rather than dependent on stochastic generation, assessments must remain sufficiently explainable to withstand legal scrutiny, and safeguards are necessary to prevent hallucinated evidence from being interpreted as proof of compliance. PolicyLens addresses these concerns

through constrained decoding, ontology-guided extraction, and multi-stage validation. Cross-lingual legal NLP introduces further complexity that is particularly relevant to AIGC platforms. Because these platforms operate across linguistic and jurisdictional boundaries, their privacy policies may differ not only in language but also in legal traditions and terminological conventions. Effective multilingual transfer therefore depends on accounting for language-specific characteristics and applying appropriate domain adaptation strategies [18]. These considerations are incorporated into the pipeline through language-specific pre-processing and domain-adapted models for Chinese-language policies.

3. The PolicyLens Framework

PolicyLens employs a three-stage computational architecture to evaluate privacy policy documents against data protection requirements. As illustrated in Figure 1, the framework receives privacy policy documents and relevant regulatory texts as inputs and systematically evaluates their alignment. The resulting assessments are expressed as multi-dimensional compliance scores at the obligation, category, and overall levels.




Fig.1: Overview of the PolicyLens computational compliance assessment framework.

3.1 Regulatory Ontology Construction

PolicyLens is underpinned by a machine-readable regulatory ontology that formally represents the disclosure obligations established by data protection regulations. Its construction follows a systematic legal analysis in which automated extraction is combined with expert validation. The process begins by identifying provisions within each regulatory framework that establish disclosure or transparency requirements for data controllers/processors. For GDPR, the analysis covers Articles 12-22 (data subject rights and information obligations), Articles 44-49 (cross-border transfer safeguards), and Articles 35-36 (data protection impact assessments). The PIPL analysis includes Articles 17, 23-25 (individual consent and notification), Articles 38-39 (cross-border provisions), and Articles 55-56 (special processing rules). For CCPA, Sections 1798.100-1798.199 (consumer rights and business obligations) are examined.

Each obligation identified through this process is represented as a structured record containing the following attributes: (a) *obligation_id*, providing a unique identifier; (b) *regulation*, specifying the applicable framework (GDPR/PIPL/CCPA); (c) *article*, identifying the corresponding legal provision;

(d) category, assigning the obligation to one of 12 thematic categories (see Table 1); (e) obligation_type, classified as inform, provide_mechanism, or document; (f) specificity, designated as high, medium, or low according to the precision with which the required disclosure is defined; and (g) cross_mappings, linking equivalent obligations across the other regulatory frameworks. The resulting ontology contains 127 obligations in total, comprising 52 from GDPR, 41 from PIPL, and 34 from CCPA.

Table 1
Obligation Categories and Distribution Across Regulations

Category	GDPR	PIPL	CCPA	Total
Data Collection Purpose	5	4	3	12
Legal Basis for Processing	6	4	3	13
Data Subject Rights	8	6	5	19
Cross-border Transfer	5	4	2	11
Data Retention Period	3	3	2	8
Third-party Sharing	5	4	4	13
Automated Decision-Making	4	3	2	9
Training Data Provenance	4	3	3	10
Opt-Out Mechanism	3	3	4	10
Data Breach Notification	3	3	2	8
DPO Contact Information	3	2	2	7
Cookie/Tracking Disclosure	3	2	2	7

3.2 LLM-Based Policy Analysis Pipeline

The second stage of PolicyLens applies a three-step information extraction pipeline to privacy policy documents, systematically identifying and extracting evidence relevant to regulatory compliance.

3.2.1 Text Segmentation

Privacy policies are generally lengthy legal documents organised into sections, subsections, and hierarchically nested clauses. The segmentation module initially reconstructs this structure through heading detection and hierarchical parsing before dividing the content into individual clauses, where each clause represents the smallest textual unit capable of independently expressing a single privacy-related obligation or commitment. To accommodate clauses extending across paragraph boundaries, the module applies an overlapping sliding-window approach. Across the dataset, each privacy policy contains an average of 87.3 clauses, with a standard deviation of 34.2, indicating considerable variation in policy length and level of granularity among platforms.

3.2.2 Clause Classification

Each segmented clause is mapped to the 12 obligation categories specified in the regulatory ontology. Because a single clause may correspond to more than one category, this process is formulated as a multi-label classification task. The classifier uses a fine-tuned LegalBERT model further adapted through domain-specific pre-training on 15,000 privacy policy documents obtained from web archives. Evaluation on a held-out validation set comprising 2,000 manually annotated clauses yields a macro-averaged F1-score of 0.91.

3.2.3 Evidence Extraction

For every classified clause, the evidence extraction module generates structured compliance evidence using a fine-tuned LLM with constrained decoding. The extraction process captures: (a)

the obligation(s) addressed by the clause; (b) its compliance status, classified as fully addressed, partially addressed, or not addressed; (c) the corresponding textual span; and (d) any conditions or qualifications associated with the disclosure. This design builds on structured information extraction methods that have demonstrated strong performance in identifying domain-specific relationships and annotations within complex documents [4]. The resulting structured compliance evidence record is subsequently passed to the ontology alignment module.

3.3 Compliance Assessment via Ontology Alignment

The final stage computes compliance scores by aligning the extracted evidence against the regulatory ontology. For each obligation o_i in the ontology, the alignment module determines whether the privacy policy provides adequate disclosure by matching extracted evidence records against the obligation's requirements. We define three compliance levels: (1) Full compliance: the policy explicitly and completely addresses the obligation; (2) Partial compliance: the policy addresses the obligation but with insufficient specificity or missing elements; (3) non-compliance: the policy does not address the obligation. The obligation-level compliance assessments are aggregated into category-level scores (averaging across obligations within each of the 12 categories) and an overall compliance score (weighted average across categories, with weights reflecting the regulatory significance of each category as determined by expert consultation). Formally, the overall compliance score S for a platform p under regulation r is computed as:

$$S(p, r) = \sum_i w_i \cdot (1/|C_i|) \cdot \sum_{o_j \in C_i} s(p, o_j)$$

Where w_i is the weight of category i , C_i is the set of obligations in category i , and $s(p, o_j) \in \{0, 0.5, 1\}$ represents the compliance score of platform p for obligation o_j (corresponding to non-compliance, partial compliance, and full compliance, respectively).

The category weights w_i are established through a Delphi process involving six data protection law experts. The experts independently evaluated the importance of all 12 categories using a 1-10 scale, followed by two rounds of discussion to achieve convergence. Final weights are obtained by normalising the resulting mean scores. These weights range from 0.062 for Cookie/Tracking Disclosure to 0.112 for Data Subject Rights, representing differences in the regulatory importance assigned to individual disclosure categories. Sensitivity analysis confirms that the principal findings remain robust under reasonable variations in the weight vector (see Section 6).

3.4 Implementation Details

PolicyLens is implemented in Python 3.10, with the Hugging Face Transformers library supporting model inference and spaCy handling text pre-processing. Clause classification is performed using a fine-tuned LegalBERT-base model containing 110M parameters. The model is trained for 5 epochs on the annotated corpus of 15,000 privacy policies using a learning rate of $2e-5$, a batch size of 32, and the AdamW optimiser. For evidence extraction, a fine-tuned Llama-3-8B model applies constrained decoding with JSON schema enforcement to generate structured outputs. Ontology alignment integrates three complementary mechanisms: exact matching for obligation-specific keywords, semantic matching based on cosine similarity between sentence embedding with a threshold of 0.85, and rule-based logic for resolving cross-references. Using a machine equipped with a single NVIDIA A100 GPU, one privacy policy requires approximately 45 seconds for end-to-end processing, allowing the complete AIGC-Privacy-200 dataset to be processed in batches within three hours.

4. The AIGC-Privacy-200 Dataset

4.1 Platform Selection

The AIGC-Privacy-200 dataset was developed through a systematic selection procedure intended to represent the diversity of the AIGC ecosystem. Candidate platforms were identified from four sources: (1) AI-related rankings on major app stores (Apple App Store, Google Play, Huawei AppGallery); (2) industry reports published by CB Insights, Pitchbook, and 36Kr; (3) academic surveys covering AIGC services; and (4) regulatory filings (China's algorithm registry, EU AI Act notifications). This process initially yielded 412 candidate platforms. Eligibility was then determined using three inclusion criteria: (a) direct provision of generative AI services to end users; (b) availability of a publicly accessible privacy policy in English, Chinese, or both; and (c) at least 100,000 monthly active users as of January 2026, estimated from app store data, SimilarWeb traffic data, or company disclosures. Following the application of these criteria, 200 platforms were retained in the final dataset.

4.2 Data Collection

For each platform, the complete privacy policy and terms of service were collected between January 15 and March 15, 2026. Automated retrieval was performed using Playwright, with a headless browser configured to render JavaScript-intensive webpages, while platforms employing non-standard document structures were handled through manual collection. To support reproducibility, every document was time-to-day and archived using the Wayback Machine API. Where policies were available in multiple languages, both English and Chinese versions were collected. The English version was prioritised for analysis when available; otherwise, the Chinese version was processed through professional translation.

4.3 Dataset Statistics

Table 2 presents the principal characteristics of the AIGC-Privacy-200 dataset, which comprises 200 platforms spanning five AIGC categories, four jurisdiction groups, and 23 countries of incorporation. Privacy policies contain an average of 4,832 words, with a median of 3,956 and a range of 847-18,423 words, demonstrating considerable variation in the extent of disclosure across platforms.

Table 2
Summary Statistics of the AIGC-Privacy-200 Dataset

Attribute	Value	Notes
Total Platforms	200	
LLM Chabot's	68	34%
Text-to-Image Generators	42	21%
Code Assistants	35	17.5%
Video Generators	27	13.5%
Multi-Modal Platforms	28	14%
US-Based Platforms	62	31%
EU-Based Platforms	45	22.5%
China-Based Platforms	58	29%
Other Jurisdictions	35	17.5%

5. Empirical Results

PolicyLens is applied to the AIGC-Privacy-200 dataset to conduct a comprehensive assessment of regulatory compliance. The findings are presented across four dimensions: overall compliance rates, obligation-level performance, cross-jurisdictional patterns, and variations among platform types.

5.1 Overall Compliance Rates

Figure 2 illustrates obligation-level GDPR compliance profiles across the jurisdiction groups using radar charts. A clear jurisdiction-anchored pattern emerges, with EU-based platforms demonstrating consistently higher compliance across almost all obligation categories. Their strongest performance is observed for data collection purpose (82%) and data subject rights (78%). By comparison, China-based platforms display a more limited GDPR compliance profile, with rates remaining below 42% across all categories. US-based platforms fall between these groups, performing relatively strongly in cookie/tracking disclosure (61%) while showing substantially lower compliance for training data provenance (11%) and opt-out mechanisms (17%).




Fig.2: Radar charts of obligation-level GDPR compliance profiles by platform jurisdiction.

Across the 200 platforms examined, only 46 (23%) fully satisfy the complete disclosure obligations of at least one regulatory framework. None achieves full compliance simultaneously across all three regulations. The highest overall compliance score recorded across the regulatory frameworks is 78.3%, attained by a major EU-based LLM provider. At the opposite end of the distribution, 12 platforms (6%) obtain scores below 10% under all three regulations, indicating an almost complete absence of meaningful privacy disclosure.

5.2 Obligation-Level Analysis

Figure 3 illustrates the distribution of GDPR and PIPL compliance scores across jurisdiction groups using violin plots overlaid with box plots. By displaying the complete score distributions rather than only mean values, these visualisations capture differences in skewness, multimodality, and outlier patterns within each jurisdiction. For GDPR, EU-based platforms display a right-skewed distribution with a comparatively high median, whereas China-based platforms are concentrated towards the lower end and exhibit a long tail. The opposite pattern emerges for PIPL, where China-based platforms demonstrate a substantially higher median compliance score together with greater variance.




Fig.3: Violin plots with embedded box plots showing compliance score distributions by jurisdiction: (a) GDPR compliance; (b) PIPL compliance. Diamond markers indicate means.

Complementing the distributional analysis, Figure 4 provides a detailed heatmap of obligation-level compliance rates for every jurisdiction-regulation combination. This representation facilitates fine-grained comparison across regulatory and jurisdictional contexts, revealing interaction patterns that may not be apparent from aggregate compliance statistics alone.




Fig.4: Heat map of obligation-level compliance rates (%) across jurisdiction-regulation combinations.

Two categories exhibit particularly low levels of compliance. Training data provenance, which requires disclosure of the sources, types, and legal basis of data used for model training, records a global average compliance rate of only 11.3%. This limited disclosure is especially significant given the substantial public and regulatory attention directed towards AI training data practices. Meaningful information concerning training data sources is provided by only 18% of EU-based platforms and 8% of China-based platforms, despite transparency requirements under both GDPR and PIPL concerning the purposes and sources of data processing. Opt-out mechanisms similarly

demonstrate limited compliance, with a global average rate of 17.1%. These mechanisms allow users to prevent their data from being used for model training or to request the removal of their contributions from training datasets. Some form of opt-out mechanism is available on 25% of EU-based platforms, compared with 12% of China-based platforms and 19% of US-based platforms. This low level of provision is particularly important because relevant opt-out rights are explicitly established under GDPR Article 21 (right to object), PIPL Article 15 (right to refuse automated decision-making), and CCPA Section 1798.120 (right to opt-out of sale or sharing).

5.3 Cross-Jurisdictional Patterns

Our analysis identifies systematic differences in the strategies adopted by AIGC platforms to address multi-jurisdictional compliance. Three dominant approaches emerge:

- (1) Home-Jurisdiction Anchoring (62% of Platforms):** Privacy policies following this strategy are structured primarily around the regulatory requirements of the platform's home jurisdiction, with little or no adjustment for extraterritorial obligations. This represents the most prevalent approach and is especially common among China-based platforms, 78% of which adopt this strategy.
- (2) Multi-Jurisdictional Layering (24% of Platforms):** Under this approach, a core privacy policy is accompanied by jurisdiction-specific addenda (e.g., "GDPR Supplement," "CCPA Notice at Collection"). The strategy is more frequently observed among larger platforms operating internationally and, on average, is associated with higher cross-jurisdictional compliance scores.
- (3) Minimal Compliance (14% of Platforms):** Platforms adopting this strategy provide broad privacy disclosures without substantively addressing the requirements of any particular regulatory framework. Their privacy policies are generally shorter, averaging 1,247 words compared with 5,631 for platforms employing multi-jurisdictional layering, and this approach occurs disproportionately among smaller or more recently established AIGC services.

Figure 5 visualizes the progression of platforms from jurisdiction to compliance strategy and subsequent compliance outcome through an alluvial diagram. Home-jurisdiction anchoring clearly dominates across all jurisdictions, accounting for 124 of the 200 platforms, but is disproportionately associated with low compliance outcomes. By comparison, only 48 platforms adopt multi-jurisdictional layering, yet 28 of these achieve high compliance levels (>60%).




Fig.5: Alluvial diagram showing the flow from platform jurisdiction through compliance strategy to compliance outcome level.

These findings are consistent with the broader tendency for regulatory compliance within the

technology sector to develop reactively rather than proactively. Navigating multiple data protection regimes imposes considerable compliance costs, particularly on smaller platforms without dedicated legal teams. Comparable challenges have been observed in AI safety and deployment, where effective incorporation of safety constraints depends on deliberate integration into system architecture rather than their addition after system development [15].

5.4 Platform Type Analysis

Figure 6 presents box plots overlaid with jittered individual data points to examine compliance distributions across AIGC platform types and according to policy update recency. Panel (a) shows that LLM chatbots have the highest median compliance scores, with their overall distribution positioned noticeably further to the right than those of other platform categories. Video generators display the most compressed and left-skewed distribution, with a median below 20% and several outliers corresponding to platforms with minimal privacy disclosures. The dispersion of the jittered observations within each category further demonstrates substantial variation in compliance scores, indicating that platform type alone does not account for differences in compliance behaviour. Panel (b) indicates a clear monotonic association between the recency of policy updates and compliance. Platforms whose privacy policies were updated within three months of data collection record a mean compliance rate of 41.3%, with an interquartile range of 25-55%. By comparison, platforms whose policies have remained unchanged for more than 12 months achieve a mean of only 18.2%, with scores concentrated within a narrow range near the lower end of the distribution. This pattern indicates a strong relationship between active policy maintenance and higher compliance, although the direction of causality cannot be established.




Fig.6: Box plots with jittered data points: (a) compliance distribution by AIGC platform type; (b) compliance vs. privacy policy update recency. Diamond markers indicate means.

6. Framework Validation

The accuracy of PolicyLens was validated through expert legal audits conducted on a subset of 40 platforms, comprising 10 from each jurisdiction group. Framework-generated assessments were compared directly with evaluations performed by six legal professionals specialising in data protection law, including two experts for each of GDPR, PIPL, and CCPA. Each platform was assessed independently by two experts, with any disagreements subsequently resolved through discussion.

6.1 Validation Results

Figure 7 summarises the validation results. Panel (a) reports the F1-score achieved by PolicyLens across individual obligation categories. Overall, the framework attains an F1-score of 0.89, while category-specific performance ranges from 0.78 for Training Data Provenance to 0.95 for DPO Contact Information. Higher accuracy is observed for categories based on clearly defined and formulaic disclosures, such as DPO contact and cookie tracking, whereas categories involving greater legal interpretative complexity, including training data provenance and automated decision-making, produce comparatively lower but still acceptable scores. Panel (b) benchmarks PolicyLens against five baseline methods. Keyword matching produces the weakest performance ($F1 = 0.54$), indicating the limitations of purely lexical techniques for compliance assessment. Performance improves progressively with TF-IDF with SVM ($F1 = 0.68$), BERT-base ($F1 = 0.76$), and LegalBERT ($F1 = 0.80$), with the latter demonstrating the benefit of domain-specific pre-training. GPT-4 zero-shot achieves a comparatively strong F1-score of 0.82 but remains below the fine-tuned PolicyLens pipeline ($F1 = 0.89$). The principal advantage of PolicyLens is evident in recall (0.87 vs. 0.81), suggesting that ontology-guided extraction is better able to capture subtle or implicit compliance evidence that may be overlooked by general-purpose models.




Fig.7: Validation results: (a) F1-score by obligation category; (b) comparison with baseline methods.

6.2 Error Analysis

A detailed error analysis was performed on the 40-platform validation subset to identify the main sources of PolicyLens misclassification. False negatives, where the framework failed to detect existing compliance evidence, represented 62% of all errors, whereas false positives, involving non-compliant disclosures incorrectly identified as compliant, accounted for the remaining 38%. False negatives occurred most frequently when obligations were addressed indirectly or implicitly through cross-references to external documents (e.g., "see our Data Processing Agreement") rather than through explicit disclosures within the policy itself. By contrast, false positives primarily resulted from aspirational or hedged statements (e.g., "we may provide" or "we aim to ensure") being interpreted as definitive compliance commitments. These patterns illustrate a fundamental difficulty in automated legal text analysis: distinguishing substantive compliance from performative compliance, in which compliance-oriented language is used without establishing concrete commitments. Such distinctions depend heavily on contextual interpretation and therefore remain challenging for current NLP systems. Similar limitations arise in other NLP evaluation settings, where

surface-level metrics may not adequately reflect meaningful system performance [12]. Future research should therefore investigate more sophisticated methods for recognising hedged, conditional, and non-binding commitments in legal texts.

7. Discussion

7.1 Implications for Regulators

The findings carry several implications for data protection regulators. First, the widespread reliance on jurisdiction-anchored compliance indicates that existing extraterritorial enforcement mechanisms may not be sufficient to achieve consistent compliance across jurisdictions. Addressing this limitation may require greater regulatory coordination, including mutual recognition arrangements or harmonised compliance standards tailored specifically to AIGC services. Second, the particularly low compliance rates observed for AIGC-specific obligations, notably training data provenance and opt-out mechanisms, suggest a need for supplementary regulatory guidance that more directly addresses the distinctive data flows associated with generative AI systems. Third, the strong association between policy update recency and compliance provides a potential basis for more targeted regulatory monitoring. Prioritising platforms with outdated privacy policies could enable regulators to allocate enforcement resources more efficiently. PolicyLens could support this process as an automated screening mechanism, helping regulatory agencies identify platforms that warrant more detailed manual review.

7.2 Implications for Platform Developers

For AIGC platform developers, the findings emphasise the value of incorporating compliance considerations early in platform development. Although multi-jurisdictional layering requires greater resources, it produces substantially stronger compliance outcomes than home-jurisdiction anchoring. Particular attention should be directed towards the identified weaknesses in training data provenance and opt-out mechanisms, especially as these areas may attract greater regulatory scrutiny. PolicyLens and AIGC-Privacy-200 also provide practical resources for benchmarking privacy policies against industry practices and identifying specific compliance deficiencies. Their open-source availability allows these resources to be incorporated into compliance workflows, facilitating continuous monitoring rather than reliance solely on periodic audits. Such an approach is consistent with safety-by-design principles in responsible AI, whereby compliance constraints are integrated into development processes instead of being introduced retrospectively [15].

7.3 Implications for NLP and Legal AI Research

From a technical perspective, the findings demonstrate both the capabilities and remaining limitations of LLM-based compliance assessment. PolicyLens achieves strong performance ($F1 = 0.89$), supporting the feasibility of automated compliance checking. However, error analysis identifies continuing difficulties in recognising implicit compliance, resolving cross-document references, and interpreting hedged language. These limitations reflect broader challenges in legal NLP, where the precision and structure of legal language can facilitate automated analysis while its contextual and interpretative complexity can constrain performance. Effective information extraction for legal compliance can also benefit from methodological advances across the wider NLP field. Approaches developed for structured knowledge extraction, including biological pathway annotation, offer architectural principles that can be adapted to compliance analysis [4]. Likewise, advances in multilingual and cross-domain language processing provide useful foundations for

addressing the linguistic diversity of cross-jurisdictional regulatory analysis [18], particularly where privacy policies differ across languages and legal traditions.

7.4 Limitations

Several limitations warrant consideration. First, the analysis relies on publicly available privacy policies and therefore cannot determine whether platforms' actual data practices correspond with their stated disclosures, commonly referred to as the "privacy policy gap." Second, although the dataset provides broad coverage, it represents a specific period (January-March 2026) and does not account for subsequent policy revisions resulting from regulatory developments. Third, the binary and categorical compliance scoring inevitably simplifies legal assessment, as compliance may vary in degree depending on regulatory interpretation and enforcement precedents. Fourth, the framework currently supports English and Chinese texts; application to additional languages would require further validation of the NLP pipeline's cross-linguistic performance. A further limitation concerns the expert validation process. Although the framework demonstrates strong agreement with expert assessments, these evaluations may themselves involve subjective legal judgements on which reasonable practitioners could differ. More broadly, variation in how regulatory provisions is interpreted across jurisdictions and courts creates inherent uncertainty in compliance assessment, thereby constraining the level of accuracy that any automated system can ultimately achieve.

7.5 Ethical Considerations

Several ethical considerations arise from this research. First, releasing platform-specific compliance scores could affect the commercial interests of the platforms assessed. To reduce this risk, individual platform results are anonymised in the publicly released dataset, while findings are reported only in aggregate by jurisdiction and platform type. Access to platform-identified results is restricted to researchers who sign a data use agreement outlining responsible disclosure requirements.

Second, automated compliance assessment may encourage "compliance theater", whereby platforms modify privacy policy wording to perform well under automated evaluation without making corresponding changes to their underlying data practices. PolicyLens should therefore be regarded as a screening and benchmarking tool rather than a replacement for comprehensive legal audits or regulatory enforcement. Since the framework evaluates disclosed practices rather than actual platform behaviour, its findings should be considered alongside technical audits, analysis of user complaints, and regulatory investigations. Third, constructing the regulatory ontology requires interpretative judgements concerning the scope and specificity of disclosure obligations. Legal experts may reasonably differ over whether provisions establish binding requirements or represent recommended practices. Although the Delphi-based expert consultation process reduces reliance on individual judgement, the resulting ontology reflects a consensus interpretation that may not correspond with every judicial or regulatory interpretation of the relevant law.

8. Conclusion

This paper presents PolicyLens, a computational framework for evaluating the compliance of AIGC platform privacy policies with data protection regulations. By integrating a regulatory ontology comprising 127 obligations, an LLM-based information extraction pipeline, and AIGC-Privacy-200, the study provides a comprehensive cross-jurisdictional assessment of privacy compliance across the AIGC industry. The findings reveal substantial gaps in current privacy compliance. Only 23% of platforms achieve full compliance with any single regulation, while jurisdiction-anchored policy

drafting remains widespread. Compliance is particularly limited for AIGC-specific obligations, including training data provenance (11%) and opt-out mechanisms (17%), indicating that meaningful transparency in these areas remains insufficient across much of the industry. Validation against expert legal audits demonstrates that PolicyLens can support scalable and reproducible compliance assessment, achieving an F1-score of 0.89. The framework therefore offers practical value for regulatory screening and platform-level compliance improvement. To facilitate further research in RegTech for AI governance, PolicyLens, AIGC-Privacy-200, and the regulatory ontology are released as open-source resources. Future research should expand PolicyLens to additional regulatory frameworks, including Brazil's LGPD and India's DPDP Act, and introduce longitudinal monitoring to examine changes in compliance over time. Further work could also integrate compliance assessment more directly into platform development workflows, supporting a transition from periodic auditing towards continuous automated monitoring of AIGC services.

Acknowledgments

The author thanks the legal professionals who participated in the expert validation study and the anonymous reviewers for their constructive feedback.

References

- [1] Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M. S., Bohg, J., Bosselut, A., Brunskill, E., Brynjolfsson, E., Buch, S., Card, D., Castellon, R., Chatterji, N., Chen, A., Creel, K., Davis, J. Q., Demszky, D., . . . Liang, P. (2021). *On the opportunities and risks of foundation models*. arXiv. <https://doi.org/10.48550/arxiv.2108.07258>
- [2] Greenleaf, G. (2021). *Global data privacy laws 2021: despite COVID delays, 145 laws show GDPR dominance*. Elsevier BV. <https://doi.org/10.2139/ssrn.3836348>
- [3] Voigt, P., & von dem Bussche, A. (2017). *The EU general data protection regulation (GDPR): A practical guide*. Springer International Publishing. <https://doi.org/10.1007/978-3-319-57959-7>
- [4] Amaral, O., Abualhaija, S., Torre, D., Sabetzadeh, M., & Briand, L. C. (2022). AI-enabled automation for completeness checking of privacy policies. *IEEE Transactions on Software Engineering*, 48(11), 4647-4674. <https://doi.org/10.1109/tse.2021.3124332>
- [5] Costante, E., Sun, Y., Petković, M., & den Hartog, J. (2012). *A machine learning solution to assess privacy policy completeness*. In Proceedings of the 2012 ACM workshop on Privacy in the electronic society, <https://doi.org/10.1145/2381966.2381979>
- [6] Wilson, S., Schaub, F., Dara, A. A., Liu, F., Cherivirala, S., Giovanni Leon, P., Schaarup Andersen, M., Zimmeck, S., Sathyendra, K. M., Russell, N. C., B. Norton, T., Hovy, E., Reidenberg, J., & Sadeh, N. (2016). *The creation and analysis of a website privacy policy corpus*. In Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), <https://doi.org/10.18653/v1/p16-1126>
- [7] Harkous, H., Fawaz, K., Lebre, R. C. A. m., Schaub, F., Shin, K. G., & Aberer, K. C. A. v. (2018). *Polis: Automated analysis and presentation of privacy policies using deep learning*. In 27th USENIX Security Symposium (USENIX Security 18), <https://www.usenix.org/conference/usenixsecurity18/presentation/harkous>
- [8] Ahmad, W., Chi, J., Tian, Y., & Chang, K.-W. (2020). *PolicyQA: A reading comprehension dataset for privacy policies*. In Findings of the Association for Computational Linguistics: EMNLP 2020, <https://doi.org/10.18653/v1/2020.findings-emnlp.66>
- [9] Sarne, D., Schler, J., Singer, A., Sela, A., & Bar Siman Tov, I. (2019). *Unsupervised topic*

- extraction from privacy policies* In Companion Proceedings of The 2019 World Wide Web Conference, <https://doi.org/10.1145/3308560.3317585>
- [10] Zimmeck, S., Goldstein, R., & Baraka, D. (2021). *PrivacyFlash Pro: automating privacy policy generation for mobile apps* In Proceedings 2021 Network and Distributed System Security Symposium, <https://doi.org/10.14722/ndss.2021.24100>
- [11] Torre, D., Abualhaija, S., Sabetzadeh, M., Briand, L., Baetens, K., Goes, P., & Forastier, S. (2020). *An AI-assisted approach for checking the completeness of privacy policies against GDPR* In 2020 IEEE 28th International Requirements Engineering Conference (RE), <https://doi.org/10.1109/re48521.2020.00025>
- [12] Xu, X., Xu, W., Ouyang, S., & Li, L. (2025). *CA*: Addressing evaluation pitfalls in computation-aware latency for simultaneous speech translation* In Findings of the Association for Computational Linguistics: NAACL 2025, <https://doi.org/10.18653/v1/2025.findings-naacl.393>
- [13] European Parliament, & Council of the European Union. (2024). Regulation (EU) 2024/2847 of the european parliament and of the council, (2024). <https://eur-lex.europa.eu/eli/reg/2024/2847/oj>
- [14] Migliorini, S. (2024). China's interim measures on generative AI: Origin, content and significance. *Computer Law & Security Review*, 53, 105985. <https://doi.org/10.1016/j.clsr.2024.105985>
- [15] Zhang, Z., Shu, K., & Mo, M. (2026). Safemind: A risk-aware differentiable control framework for adaptive and safe quadruped locomotion. *IEEE Access*, 14, 50246-50269. <https://doi.org/10.1109/access.2026.3679520>
- [16] Chalkidis, I., Fergadiotis, M., Malakasiotis, P., Aletras, N., & Androutsopoulos, I. (2020). *LEGAL-BERT: the muppets straight out of law school* In Findings of the Association for Computational Linguistics: EMNLP 2020, <https://doi.org/10.18653/v1/2020.findings-emnlp.261>
- [17] Lippi, M., Pałka, P., Contissa, G., Lagioia, F., Micklitz, H.-W., Sartor, G., & Torroni, P. (2019). CLAUDETTE: An automated detector of potentially unfair clauses in online terms of service. *Artificial Intelligence and Law*, 27(2), 117-139. <https://doi.org/10.1007/s10506-019-09243-2>
- [18] Xu, X., Ouyang, S., Yan, B., Fernandes, P., Chen, W., Li, L., Neubig, G., & Watanabe, S. (2024). *Cmu's IWSLT 2024 simultaneous speech translation system* In Proceedings of the 21st International Conference on Spoken Language Translation (IWSLT 2024), <https://doi.org/10.18653/v1/2024.iwslt-1.20>