

Artificial Intelligence Governance as Policy Discourse: A Computational Analysis of Official Indian Policy Documents

Kumar Sumit^a, Prince Buragohain^b, Manendra Kumar Singh^c

^a Research Scholar, Department of Library and Information Science, Mizoram University, Aizawl/ Mizoram, Pin 796004, India.
Email: aryan.sumit121@gmail.com

^b Research Scholar, Department of Library and Information Science, Mizoram University, Aizawl/ Mizoram, Pin 796004, India.
Email: buragohainprince@gmail.com

^c Assistant Professor, Department of Library and Information Science, Mizoram University, Aizawl/ Mizoram, Pin 796004, India.
Email: manebhu007@gmail.com

Received: 21 July 2026; Accepted: 11 August 2026

AI governance in India is usually examined through legal, institutional and strategic perspectives; however, limited attention has been paid to how governance priorities are communicated and how policy language shapes authority, responsibility, and inclusion. This study thus approaches governance as a policy discourse rather than implementation outcomes through an interpretive-computational research design that combines computational text mining with discourse-oriented interpretation. A dataset of national-level AI policy documents was curated through systematic retrieval and screening. Policy sections were analysed using semantic modeling, framing analysis, and lexical polarity analysis with emphasis on interpretive analysis and institutional context rather than predictive modeling or frequency-based measures. The findings show that AI governance discourse in India is institutionally coordinated and development oriented. Innovation, economic growth and institutional coordination emerged as dominant themes while ethics, risk, and data-protection concerns appear primarily within regulatory contexts. In framing analysis, opportunity-oriented domains dominate regulatory domains. Risk and caution are addressed primarily within regulatory domains. This study demonstrates how AI governance operates through discourse and rhetorical framing, and not only through regulation. It positions computational text mining as a research design, offering a scalable approach for analysing the social policy dimensions of emerging technology governance.

Keywords: AI Governance; Computational Text Mining; Policy Discourse; Topic Modelling; Framing Analysis; Digital Government

1. Introduction

Artificial intelligence (AI) governance has emerged as a major concern in contemporary public policy as governments increasingly integrate AI into administration, welfare delivery, digital infrastructure, economic planning, and public services. Across countries, AI governance frameworks attempt to balance innovation, economic competitiveness, ethical accountability, regulatory oversight, and public trust^{1,2}. In India, AI governance is closely associated with developmental policy objectives such as digital transformation, economic growth, technological self-reliance, public-sector modernisation, and workforce adaptation^{3,4}. The rapid expansion of digital governance systems has created opportunities for administrative efficiency and service delivery while simultaneously raising concerns related to privacy, algorithmic bias, exclusion, and institutional accountability. For administrative

efficiency and public services, the Indian government increasingly relies on digital systems such as the Aadhaar biometric identification system and the Aarogya Setu application⁵. India's technology governance ecosystem combines policy innovation, public-private collaboration, and international partnerships⁴. Government initiatives increasingly emphasise security, privacy, regulation, and ethical governance while simultaneously promoting economic growth and technological competitiveness^{3,6}. National-level initiatives also encourage the integration of digital technologies into sectors such as education and public administration⁷. Despite these developments, India's digital transformation efforts continue to face ethical, socio-economic, infrastructural, and governance-related challenges, highlighting the need for socially responsive and institutionally accountable AI governance frameworks^{8,9,10}.

To understand AI governance adequately, attention must be given not only to legal and institutional mechanisms but also to the language and narratives through which governance priorities are communicated and legitimised. Policy documents do not merely describe governance arrangements; they also construct institutional authority, define policy priorities, and frame inclusion, risk, responsibility, and public participation. Computational text mining offers a systematic approach for examining these discursive patterns across large policy corpora. In policy research, text mining is used to identify thematic structures, rhetorical patterns, and variations in policy preferences across institutions and sectors¹¹. Text-mining approaches involve data preparation, computational processing, and analytical modeling, thereby integrating information retrieval, computational linguistics, natural language processing, and machine learning techniques¹². Recent developments in computational text analysis have enabled researchers to address the growing challenge of information overload in governance and policy research¹³. Such approaches are increasingly applied across domains including public policy analysis, healthcare, governance studies, and digital communication research¹⁴.

1.1 AI Governance as Social Policy and Discursive Power

Existing literature on AI governance looks at governance mainly through formal structures and principles, including regulatory instruments, institutional arrangements, and ethical frameworks. The majority of research focuses on issues such as legal compliance, accountability mechanisms, and technical approaches to risk mitigation^{1,2}.

These perspectives give less attention to how policy language itself influences governance priorities, because the tone of language matters in terms of governance. Policy texts define social objectives of the government. They assign responsibilities and set boundaries for inclusion, protection and risk^{15,16}. Previous social policy research highlights that government priorities go beyond welfare services or the distribution of resources. The government priorities work through labour rules, skill development, risk management, and the way in which institutions frame their actions^{17,18}. In this framework, AI governance positions technological change as an opportunity requiring adjustment by individuals, institutions, and labour

markets. Current approaches to technology governance emphasise preparation for future change, adopting principles such as transparency, fairness, accountability, and trust^{2,19}. Governance arrangements affect how individuals and institutions are expected to adapt. Research on governance frequently highlights ethical principles. A considerable body of AI governance research takes a normative approach, emphasising ethical concepts such as transparency, accountability, fairness, and trust. These principles are widely recognised; however, critical policy scholarship has pointed out that ethical language does not necessarily translate into redistributive results or strengthened social safeguards²⁰. Ethical commitments may coexist with governance approaches that emphasise innovation, efficiency, and competitiveness. Reviewing how ethical and regulatory discourse is embedded within policy texts is important for understanding the societal consequences of AI.

Policy discourse analysis provides a framework for examining these dynamics. Discourse-oriented approaches conceptualise policy texts as sites where authority, power and legitimacy are created and negotiated, instead of neutral expressions of policy intent^{21,22}. Through discourse, governments define policy issues, institutionalise priorities, and determine which risks and social groups are worthy of attention. In the domain of technological governance, managerial and optimistic policy language often presents technological changes as inevitable and controllable. This framing can narrow the space for political contestation and public deliberation^{15,23}.

In the Global South, AI governance is frequently embedded within developmental state narratives emphasising economic growth, digital transformation, technological sovereignty, and global competitiveness²⁴. National AI strategies increasingly position artificial intelligence as an instrument for accelerating productivity, modernising governance, and integrating citizens into digital economies. Recent scholarship argues that AI governance in developing economies often combines innovation-oriented developmentalism with selective inclusion frameworks, where participation in technological systems is closely tied to employability, skilling, and economic integration rather than redistributive welfare protections. Studies on India's AI governance ecosystem further demonstrate how AI policy²⁵ discourse balances developmental aspirations with

emerging concerns regarding ethics, accountability, data governance, and public trust^{4,6}. Rather than functioning solely as technical governance instruments, AI policies increasingly operate as institutional narratives that legitimise particular models of growth, labour adaptation, and state-led digital transformation.

Computational text mining analysis helps to analyse large amounts of policy texts, and it helps to identify recurring themes, semantic structures, and rhetorical patterns^{26,27}. However, existing applications often rely on techniques such as topic modeling or sentiment analysis. There is no doubt that these techniques provide descriptive results, but they lack theoretical integration. Such approaches are insufficient for explaining how governance discourse functions as a social policy mechanism. Integrating computational methods with interpretive policy analysis allows examination of both thematic content and discursive structure. This enables us to analyse how policy meanings are prioritised, stabilised, and institutionalised across contexts²⁸.

Building on this perspective, the present study conceptualises AI governance as a discursively enacted social policy domain. Rather than assessing regulatory adequacy or policy effectiveness, it examines how national AI policy texts in India construct governance priorities through language. By combining section-level semantic modeling with framing and governance-lexical polarity analysis, the study situates computational text mining within a discourse-analytical framework that preserves institutional context while enabling systematic interpretation of social policy implications. Accordingly, the study adopts a section-level, interpretive-computational approach to analyse AI governance as a discursively structured social policy domain.

1.2 Analytical Orientation: Welfare Rationalities in AI Policy

Governance structure and discourse in terms of AI policy should not be understood only as a regulatory design but as a part of broader welfare state restructuring, as it shapes social outcomes. Present-day social policy has shifted from redistribution to

activation-oriented policy emphasising skills, employability, and productivity. Social policy has been linked to labour market participation, and welfare is framed through economic contribution^{17,18}. Policies oriented towards technological advancements reshape access to social participation and define who is included in development through technological capability. In such situations, economic participation becomes a condition for inclusion. Policy language defines problems and priorities, and these texts legitimise certain policy choices over others^{21,22}. The policy discourse reveals implicit welfare logic and focuses on development, labour adaptation, and institutions.

1.3 Operationalising Social Policy Dimensions in AI Governance

In addition to viewing AI governance as a technical or regulatory space, this research defines AI governance as a process of social-policy making that influences inclusion, adjustment, accountability, and opportunities for development. In order to explain the rationality of welfare within the framework of AI governance discourse, the discursive operationalisation model based on social policy and governance literature is adopted in this research.

Instead of quantifying the welfare outcome itself, this study focuses on analysing how policy documents emphasise certain governance logic. Social-policy orientations were therefore interpreted through discursive indicators embedded within policy narratives. The four governance orientations and their discursive indicators are summarised in Table 1.

The study operationalises four broad governance orientations within AI policy discourse:

These categories were not treated as rigid coding variables but as interpretive analytical dimensions guiding thematic consolidation and discourse interpretation. The framework enabled examination of how AI governance discourse positions citizens, workers, institutions, and vulnerable populations within narratives of technological transformation and national development. This operationalisation further supported interpretation of how governance priorities

Table 1 — Governance orientation and discursive indicators within policy texts

Governance orientation	Discursive indicators within policy texts
Productivist orientation	Innovation, competitiveness, economic growth, productivity, industrial transformation
Activation-based welfare logic	Skilling, employability, workforce adaptation, participation in digital economy
Protective welfare orientation	Rights, safeguards, regulation, accountability, privacy, risk mitigation
Participatory and inclusive governance	Inclusion, accessibility, public trust, citizen participation, equitable access

are distributed across policy sections and institutional contexts, particularly in relation to labour-market adaptation, responsible AI, economic modernisation, and inclusion-oriented developmental discourse.

The objective of this study is to investigate whether the AI policy framework reflects a protective welfare logic, a productivist orientation, or an activation-based model. Building on this orientation, this study formulates the following research questions:

RQ1. How is AI governance discursively structured across sections of India's national policy framework?

RQ2. How are social policy themes related to AI distributed and institutionalised across different policy sections?

RQ3. How do Indian AI policy texts frame AI in terms of developmental opportunity and precautionary risk?

RQ4. What rhetorical polarity characterises AI policy discourse across institutional and thematic contexts?

RQ5. Which social groups are foregrounded within India's AI policy discourse, and how are they positioned within narratives of development and governance?

2. Methodology

2.1 Research Design

This study combines computational analysis with interpretive reading to examine the discursive construction of artificial intelligence (AI) governance within Indian public policy. It treats policy documents as governance texts expressing institutional authority and responsibility. The analysis focuses on a discursive approach through which AI policy is framed, governed and operationalised by the state. Computational analysis, together with discourse-oriented interpretation, identifies patterns in framing and rhetorical orientation across policy texts. This approach enables us to examine governance discourse across multiple documents while retaining its institutional context.

2.2 Corpus Construction and Document Retrieval

Policy documents were collected in several stages, from automated retrieval to manual selection of policy texts. Documents concerning digital governance and artificial intelligence were obtained from the official websites of GOI and associated government websites. The documents selected included policy documents,

strategy documents, guidelines and white papers on artificial intelligence. 260 policy documents were found, which comprised a wide array of AI policy texts. Basic metadata on these documents such as issuer of the document, year of issuance, type of document, and website where they can be found was noted down for screening.

2.3 Document Screening, Corpus Curation, and Policy Selection

The initially retrieved corpus consisted of 260 AI-related policy and governance documents collected from official Government of India portals and institutionally affiliated repositories. Following preliminary screening, approximately 40 documents were retained based on substantive relevance to artificial intelligence governance, public-sector deployment, regulatory frameworks, ethical principles, workforce transformation, or national development strategy.

The final analytical corpus was intentionally restricted to seven national-level policy documents that collectively represented the principal institutional and discursive dimensions of India's contemporary AI governance framework. The selection followed a purposive policy-curation strategy rather than exhaustive inclusion. Documents were retained based on five criteria: (1) national institutional authority, (2) explicit engagement with AI governance or deployment, (3) policy-level strategic orientation, (4) substantive discussion of governance, ethics, inclusion, skilling, regulation, or public-sector implementation, and (5) influence on national AI discourse across sectors.

The chosen texts reflected several but interrelated dimensions of the AI governance discourses in India. The dimension of foundational principles and governance were captured through Principles for Responsible AI and Operationalizing Principles for Responsible AI, documents published by NITI Aayog within the context of "Responsible AI for All" initiative where AI governance was defined in terms of constitutional values, accountability, inclusion, and public trust^{29,30}. The dimension of development economics was captured through AI for Viksit Bharat where AI is depicted as a driver of productivity, innovation, industrial revolution, and rapid growth³¹. The dimension of transition and adaptation was captured through the Roadmap for Job Creation in the AI Economy, where the issue of AI-driven changes to

Table 2 — List of Documents

Document	Institutional focus	Governance relevance
National Strategy for Artificial Intelligence (NITI Aayog, 2018)	National AI strategy	Foundational national AI vision and governance framing
Principles for Responsible AI (NITI Aayog, 2021)	Ethical AI governance	Responsible AI principles grounded in constitutional and governance frameworks
Operationalizing Principles for Responsible AI (NITI Aayog, 2021)	Governance implementation	Institutional operationalisation of responsible AI governance
AI for Viksit Bharat (NITI Frontier Tech Hub, 2025)	Economic growth and productivity	AI-led development, industrial productivity, and innovation strategy
Roadmap for Job Creation in the AI Economy (NITI Frontier Tech Hub, 2025)	Workforce transformation	AI employment transition, reskilling, and labour-market adaptation
AI for Inclusive Societal Development (NITI Frontier Tech Hub, 2025)	Inclusive development	AI discourse relating to informal workers, inclusion, and social participation
Report of Committee-D on Cyber Security, Safety, Legal and Ethical Issues (MeitY)	Cyber security and ethics	Regulatory, ethical, and cyber-governance dimensions of AI

Table 3 — Corpus composition

Source type	Count	Description
Crawled policy documents	260	Retrieved from official Government of India portals and websites
Policy-relevant documents after screening	40	Retained based on relevance and institutional authority
Final curated documents	7	Policy reports, strategies, and guidelines
Total sections (unit of analysis)	37	Segmented policy sections

employment was extensively discussed³². Social inclusion and informal-sector discourse were captured through *AI for Inclusive Societal Development*, which foregrounded informal workers, digital inclusion, accessibility, and equitable participation within AI-driven development³³. Regulatory, ethical, and cyber-governance dimensions were further represented through the Ministry of Electronics and Information Technology report *Committee-D on Cyber Security, Safety, Legal and Ethical Issues*, which addressed cybersecurity, AI ethics, explainability, privacy, and trust within emerging AI ecosystems³⁴.

Together, these documents enabled analysis of AI governance not merely as technical regulation but as a broader developmental, institutional, and social-policy discourse extending across ethics, innovation, labour transformation, inclusion, and governance coordination. The seven documents included in the final corpus are summarised in Table 2.

The corpus intentionally excluded documents that were primarily technical manuals, implementation reports, pilot studies, procurement guidelines, sector-specific feasibility assessments, or documents containing only peripheral references to AI. For example, pilot implementation studies such as the *Healthcare Artificial Intelligence Catalyst (HAIC) Pilot Study* were excluded because their primary focus was operational feasibility and implementation

assessment rather than national governance discourse³⁵.

Finally, the seven selected documents were segmented into analytically substantive policy sections, enabling section-level examination of governance framing, rhetorical orientation, thematic distribution, and institutional discourse across India's evolving AI governance landscape. The corpus composition at each stage of screening is presented in Table 3.

2.4 Text Indexing, Cleaning, and Normalisation

Selected policy documents were converted from PDF to machine-readable texts and were later extracted page by page to preserve document order and section structure. The extracted text was indexed to maintain alignment between textual content and document metadata. To ensure that the analysis is focused on substantive policy content rather than formatting artefacts, non-substantive elements such as headers, footers, table of contents and reference lists were removed.

2.5 Section-Level Parsing and Unit of Analysis

In the next step, cleaned texts were segmented into functional policy sections such as strategic statements, governance frameworks, ethics provisions, and implementation guidelines. Section-level parsing of

Table 4 — Topic–theme mapping

Topic ID	Representative keywords	Assigned analytical theme
0	governance, regulation, framework, oversight	Governance & Regulatory Frameworks
1	data, privacy, protection, consent	Data Protection & Privacy
2	innovation, industry, growth, startups	Innovation, Industry & Economic Growth
3	ethics, responsibility, trust, fairness	Ethical & Responsible AI
4	security, misuse, risk, cyber	Security & Risk Governance
5	skilling, workforce, capacity, education	Skilling & Capacity Building

the policy texts was conducted as different sections performed different governance functions within policy documents. Section boundaries were identified using explicit headings and numbering where available. A conservative approach was applied, and only clearly defined sections were retained to avoid artificial segmentation. These sections were then chosen as the unit of analysis, enabling examination of both policy themes and their institutional location within documents.

2.6 Section Audit and Corpus Refinement

Before analysis, extracted sections were reviewed based on proper headings, section counts and title-content match. During review, analytically non-relevant sections such as annexures, references, and front matter were identified and removed. These sections were then refined to ensure structural clarity and substantive relevance. After removal of structurally irrelevant sections, 37 analytically substantive section-level texts were retained as the final unit of analysis.

2.7 Semantic Modeling, Framing, and Rhetorical polarity

Recurring meaning patterns were identified from the policy documents by conducting topic modeling at the section level using BERTopic. Formed clusters were interpreted and themes that were relevant to the governance of AI were constructed. Framing of the policy documents was conducted to establish whether the discussion on AI focused more on opportunities for development or whether the focus was on regulating AI and addressing the challenges it posed. The basis for framing was based on whether AI policy was seen as an enabler for national development or whether AI needed regulation. Polarity of language was assessed from affirmative versus cautious tones in the policy documents. The use of semantic topic modeling involved the application of the BERTopic framework in conjunction with pre-trained Sentence-BERT embeddings via the all-MiniLM-L6-v2 transformer. A custom CountVectorizer was used to

enhance semantic specificity in the analysis of the policy text by employing stopword removal and phrase extraction based on domain expertise. The vectorizer was parameterized using bi- and tri-gram extraction (`ngram_range=(2,3)`), minimum document frequency (`min_df=5`), and maximum document frequency (`max_df=0.6`). Since the research utilized an interpretive discourse analytical approach as opposed to predictive text classification, BERTopic was set to use a comparatively high minimum topic size (`min_topic_size=10`) to facilitate semantic coherence in identified governance topics and minimize fragmentation within the small section-level corpus. Topic reduction (`nr_topics='auto'`) was applied to similar semantic clusters. Standard English stopwords, punctuation marks, and other non-substantive policy artifacts were excluded from the preprocessing procedure. Dimensionality reduction and density-based clustering were carried out through default UMAP and HDBSCAN parameterizations available in the BERTopic implementation.

In this regard, while the BERTopic pipeline initially produced more clusters of semantically distinguishable topics based on the corpus comprising section-level texts, a few of these topics had overlapping conceptual content, and reflected similar issues of governance through distinct lexical expressions. Considering that the present research involved interpretive discourse analysis, the semantically similar clusters were grouped together into higher order categories or themes through interpretive analysis. In effect, this ensured a theoretical fit of computational results within the context of governance. Thus, the six themes outlined in Table 4 emerge from a greater number of semantic clusters detected computationally.

2.8 Analytical Integration

Results from semantic modeling, structural mapping, framing analysis, and governance-tone assessment were compared to produce an integrated interpretation of policy discourse. Rather than

Table 5 — Analytical pipeline

Stage	Technique	Purpose
Data acquisition	Web crawler	Build policy corpus
Curation	Manual screening	Ensure relevance and authority
Segmentation	Section parsing	Preserve institutional context
Semantic modeling	BERTopic	Identify semantic clusters for interpretive consolidation
Framing analysis	Lexicon-based	Risk–opportunity positioning
Lexical polarity analysis	VADER	Governance rhetoric
Integration	Multi-layer synthesis	Discourse interpretation

Note: The six analytical themes represent interpretive consolidation of semantically related clusters generated during topic modeling.




Fig. 1 — Institutional Location of Governance Narratives

analysing each method in isolation, findings were examined together to assess how semantic structure, institutional location, ideological framing, and rhetorical tone shape the governance discourse. Visualisations were interpreted relationally, serving as analytical aids rather than measurements of policy intent or effectiveness. The analytical pipeline is summarised in Table 5.

3. Results

The following results are based on analysis of 37 section-level units derived from seven national AI policy documents.

3.1 Institutional Location of Governance Narratives

Figure 1 represents the policy themes against document sections of the AI policy framework,

indicating the sections where specific themes appear across different parts of the policy texts. Some sections reflect higher thematic density, such as governance, introduction and implementation, while security, summary and conclusion show fewer themes. This reflects that each section has different roles, and policy writing is highly structured. Introduction and governance sections contain the widest range of themes, as these sections are the foundational and coordinating parts of policy documents. Objectives and implementation sections show selective but repeated themes. Ethics-related themes appear mainly in the ethics and governance sections. The positioning of innovation and economic growth themes can be seen mainly in strategic and governance sections and is less evident in protection, security and summary sections. This suggests that




Fig. 2 — Theme Salience Bar Chart

discourse in innovation and economic growth is positioned early and structurally in the policy texts. There is a limited appearance of themes related to risk, safety and security. These are mostly limited to governance, regulatory and protection-related sections. Ethics-related themes are presented in multiple sections with lower intensity. Overall, ethical discourse complements governance and strategy. Themes related to skills and capacity building are largely visible in service-related sections. Policy documents distribute themes according to sections and are hence structured. The analysis reflects a governance-centric social policy orientation where social protection and participatory considerations have peripheral positions within the AI policy framework of India.

3.2 Discursive Salience of Policy Themes

Figure 2 examines the frequency of themes across sections. Frequency is measured by the section-level units in which a theme occurs. Strategic vision and national-level framing-related themes appear across multiple sections and themes, including governance and regulatory framework, which recur frequently. Innovation and economic-growth-related themes appear in fewer sections but are noticeable. Public-

sector adoption and service-delivery-based themes show moderate recurrence. These themes are unevenly spread across the policy documents. Risk management, cybersecurity and safety themes appear in fewer sections.

Overall, the policy documents show uneven thematic recurrence; only a small number of themes appear widely across the documents, while others are limited and in specific contexts. Figure 2 highlights the differences in the discussion of policy themes across the AI policy documents.

3.3 Ideological Framing of AI: Risk versus Opportunity

Figure 3 presents the risk-opportunity framing done on a lexicon-based approach. Each policy section was scored separately for risk-related language and opportunity-related language. The relative intensity of each framing within a section is represented by the scores, and values are normalised to account for differences in section length. The scores reflect linguistic emphasis, rather than sentiment or policy outcomes. The horizontal axis represents risk framing intensity, while opportunity framing intensity is reflected on the vertical axis, with each point corresponding to a section of a policy document. The policy sections are distributed across all four




Fig. 3 — Risk–Opportunity Framing Map

quadrants of the plot. Several sections, such as introductions, forewords and strategic sections, score higher on opportunity framing, indicating growth, development and public value. In contrast, some sections related to regulatory, security and protection score higher in risk framing and lower in opportunity framing. A number of sections are spread near the centre of the plot, reflecting balance in framing. Overall, aspirational themes lean towards opportunity framing, and regulatory and protection-oriented sections lean towards risk framing. Figure 3 indicates that risk and opportunity framing coexist across policy sections.

3.4 Theme-Specific Rhetorical Differentiation

Figure 4 shows the tone of governance across AI policy themes. There is variation in tone across

different policy domains and governance communication. This difference in tone is due to the different governance roles of policy sections. Strong affirmative rhetorical polarity can be observed in themes such as innovation, economic growth, infrastructure and governance frameworks. These domains show confidence and certainty in description, conveying institutional readiness and control. The consistency in affirmative rhetorical polarity implies a stable and confident institutional register, leaving less space for uncertainty in discourse. In contrast, domains like privacy and data protection show greater variation in tone, displaying cautious and conditional policy language. It indicates the sensitivity and regulatory complexity of data governance within policy documents. The variability in tone suggests that only specific sections of policy, such as




Fig. 4 — Rhetorical Polarity Across Key Policy Themes

regulatory and rights-oriented domains, permit the acknowledgement of constraints and risks rather than the whole document. Themes related to skilling, workforce and capacity building show a broader range of tone, indicating that the discourse recognises the challenges of transition while emphasising adaptation and capacity enhancement. The range in tone emphasises skill development and readiness, rather than workforce change as a matter of social protection or redistribution. The differential tonal range across themes reflects their functionally different institutional roles within the AI policy of India. Affirmative rhetoric dominates the AI policy documents associated with strategic direction and institutional control, while caution is exercised in limited themes related to rights and regulations.

3.5 Target Population of AI policy.

Figure 5 analyses the distribution of the target population referenced within India's AI policy discourse, revealing how different social groups are positioned in relation to AI-driven development and governance. The distribution lays emphasis on

disproportionate addressing of different social groups, resulting in uneven prioritisation. Informal workers dominate the discourse as they appear most frequently. Their high frequency in appearance can be attributed to their framing as beneficiaries of productivity or service delivery rather than as subjects of protection or rights-based governance. As the policy discourse is linked to economic integration, the form of inclusion of the informal workers is largely instrumental, emphasising participation in economic and technological processes while limiting attention to issues of vulnerability or social security. The policy emphasis appears oriented toward the economic participation of the informal workers rather than vulnerability. Women and the rural population group are also visible but not central. These groups are framed through access, empowerment and service reach; however, the emphasis remains largely instrumental, framing informal workers as beneficiaries of technological inclusion rather than as subjects of protection or collective rights within AI governance frameworks. Their emphasis is on developmental framing of these sections rather than




Fig. 5 — Target Population of AI Policy

policy responsibility through safeguards against harm. The orientation of policy formation is focused on present integration rather than future transitions, as groups such as youth, MSMEs, and students are underarticulated. The emphasis on human capital is narrowly framed. The analysis reflects that specific groups are being targeted and not society as a whole. Inclusivity is tied to productivity, participation and service delivery. Social impacts are handled but indirectly through skills, participation and adaptation.

4. Discussion

4.1 Implications for Policy and Government Practice

The analysis of India's AI policy documents provides policymakers with a reflective lens on how the government communicates governance. The discursive patterns reveal the themes where optimism dominates and where caution is exercised. This can be used by policymakers to balance the alignment between vision and safeguard by identifying the existing gaps. The rhetorical consistency varies across the policy sections, where aspirational sections dominate strategic texts. In the rhetorical structure, safeguards occur in a limited context and in limited sections. This alignment between aspirations and safeguards creates a platform for public trust. The governance discourse is distributed across institutions as inferred from theme-section mapping. These patterns clarify overlaps and gaps in discourse, paving

the way for coordination and role clarity among different ministries. AI policy discussion in India is characterized by balance between growth and regulation; and such balance allows one not to go to extreme positions while enhancing the country's credibility in international circles. Responsibility commitment is important in making policies, but in practice, they are quite abstract concepts; discursive analysis of the policy framework will make them tangible. The analysis patterns will show which issues were included or excluded from the responsibility theme. It will be helpful for policymakers for an internal assessment of their policies' consistency. Such method is helpful in assessing the evolution of a policy field temporally.

4.2 Illustrative Discursive Excerpts from AI Policy Texts

The findings can be better understood when contextualized within select snippets from the policy corpus. In several policy documents, AI tends to be portrayed in terms of narratives of development, productivity, and institution modernization. For instance, innovation emerges as an important part of policy discourse on AI for Viksit Bharat where AI is viewed as an instrument of "accelerated economic growth" and innovation³¹. Productivist governance tends to see technological change through the lens of growth, productivity, and economic modernization. Furthermore, employment-centred policy discourse sees inclusion largely in terms of adaptation and

employability, and not social protection. The roadmap for job creation in the AI economy focuses heavily on “reskilling”, “workforce readiness,” and “participation in the AI economy” as strategies for addressing changes brought about by technology³². Thus, citizens tend to be seen as economic actors responsible for incorporating themselves into emerging digital labour structures. Ethical and regulatory discourse seems to be prominent in documents concerned with responsible AI governance. Principles of responsible AI, for example, tend to focus heavily on concepts like “trust,” “fairness,” “accountability,” and “responsible innovation”. However, these ethical commitments frequently coexist alongside developmental and innovation-oriented policy language rather than replacing them. Regulatory and ethical governance therefore function within a broader discourse that continues to prioritise innovation, institutional coordination, and technological growth. These excerpts support the broader computational findings of the study by illustrating how AI governance discourse in India combines developmental ambition, institutional coordination, ethical regulation, and workforce adaptation within a predominantly growth-oriented governance narrative.

4.3 Implications for Researchers

Text mining is often applied in a tool-driven manner; this study demonstrates its research-design-oriented approach. It does not consider the computational tools as ends in themselves. It first defines what kind of policy object AI governance is and then chooses section-level analysis intentionally. The analytical design progresses across multiple levels from structure to framing to lexical polarity analysis and then to target populations. Emphasis is laid on the structural arrangement of the texts as policy meanings emerge from how the texts are arranged and not just what words appear. A comparable computational research-design approach has recently been applied to AI-governance policy in the domain of scholarly publishing³⁶. Computational analysis can support theory and not replace it, as these methods can reveal patterns that can later be interpreted, grounded in theory. Patterns are systematically identified and then theorised. Computation and interpretation go together; this approach helps social policy scholars analyse complexity without abandoning theory by extending discourse analysis into a new policy domain using computational means.

5. Conclusion

This study developed and applied a computational text-mining framework to analyse how India constructs AI governance through policy discourse. Going further than descriptive topic modeling, the study incorporated semantic, structural, ideological, and rhetorical analysis to uncover a nuanced narrative of governance that strikes a balance between innovation and regulation. The study makes a significant contribution both to AI governance literature and computational social science by showing how text mining can shed light on the discourse underlying national technology policies. Future studies could use the same methodology to conduct comparative and longitudinal studies of AI governance in different nations.

References

- 1 Cath C, Wachter S, Mittelstadt B, Taddeo M and Floridi L, Artificial intelligence and the 'good society': the US, EU, and UK approach, *Science and Engineering Ethics*, **24** (2) (2018) 505–528. <https://doi.org/10.1007/s11948-017-9901-7>
- 2 Floridi L, Cowls J, Beltrametti M, Chatila R, Chazerand P, Dignum V, Luetge C, *et al.*, AI4People—an ethical framework for a good AI society: opportunities, risks, principles, and recommendations, *Minds and Machines*, **28** (4) (2018) 689–707. <https://doi.org/10.1007/s11023-018-9482-5>
- 3 Chatterjee S, AI strategy of India: policy framework, adoption challenges and actions for government, *Transforming Government: People, Process and Policy*, **14** (5) (2020) 757–775. <https://doi.org/10.1108/TG-05-2019-0031>
- 4 Harmon S, Wilschmann M, Joshi G, Ballesteros A and Baitinger P, Decoding India's AI governance strategy and its implications for the U.S.–India bilateral relationship, *Indian Public Policy Review*, **5** (4) (2024) 51–82. <https://doi.org/10.55763/ippr.2024.05.04.003>
- 5 Lakshitha C P, Manoj A and Jeevanandhan L, AI in Indian governance, *Electronic Journal of Social and Strategic Studies*, **6** (Special Issue VII) (2025) 122–138. <https://doi.org/10.47362/EJSSS.2025.6607>
- 6 Barooah S P, Mohandas S and Hickok E, The AI task force report: the first steps towards India's AI framework, *SSRN Electronic Journal*, (2025). <https://doi.org/10.2139/ssrn.5337308>
- 7 Tambat P, Harnessing artificial intelligence to transform secondary education in India: innovations, challenges, and future directions, *International Journal for Multidisciplinary Research*, (2024). <https://doi.org/10.36948/ijfmr.2024.v06i06.30145>
- 8 Bansal S and Jain N, A comprehensive study assessing the transformative role of artificial intelligence in India's governance policy framework, *International Journal for Research in Applied Science and Engineering Technology*, **11** (7) (2023) 1748–1756. <https://doi.org/10.22214/ijraset.2023.54973>
- 9 Biswal I and Verma R, Artificial intelligence in legal practice: implications for the legal profession's future, *International Journal for Research in Applied Science and*

- Engineering Technology*, **12** (4) (2024) 1974–1988. <https://doi.org/10.22214/ijraset.2024.60199>
- 10 Pandey A C and Shakya M P, Transforming governance: the impact of artificial intelligence on Indian public administration, *ShodhKosh: Journal of Visual and Performing Arts*, **5** (4) (2024). <https://doi.org/10.29121/shodhkosh.v5.i4.2024.6301>
 - 11 Cogburn D, Shikako-Thomas K and Lai J, Using computational text mining to understand public priorities for disability policy towards children in Canadian national consultations, *In Proceedings of the 53rd Hawaii International Conference on System Sciences*, (2020), p. 696–705. <https://doi.org/10.24251/HICSS.2020.096>
 - 12 Soriano J, Au T and Banks D, Text mining in computational advertising, *Statistical Analysis and Data Mining: The ASA Data Science Journal*, **6** (4) (2013) 273–285. <https://doi.org/10.1002/sam.11197>
 - 13 Pasupathi S, Selvan K S, Dedgaonkar S G, Vishnoi M, Rampal S and Yadav D, Unlocking the power of text mining for information understanding, *In Proceedings of the 15th International Conference on Computing, Communication and Networking Technologies (ICCCNT 2024)*, IEEE, Kamand, India, (2024), p. 1–7. <https://doi.org/10.1109/ICCCNT61001.2024.10724258>
 - 14 Şahin D, The use of text mining in new media studies: a systematic review, *Yeni Medya*, (2023). <https://doi.org/10.55609/yenimedya.1340590>
 - 15 Fischer F, *Reframing public policy: Discursive politics and deliberative practices* (Oxford University Press; Oxford), 2003.
 - 16 Stone D, *Policy paradox* (W.W. Norton; New York), 2012.
 - 17 Esping-Andersen G, *The three worlds of welfare capitalism* (Princeton University Press; Princeton), 1990.
 - 18 Taylor-Gooby P, *New risks, new welfare: The transformation of the European welfare state* (Oxford University Press; Oxford), 2004.
 - 19 Jobin A, Ienca M and Vayena E, The global landscape of AI ethics guidelines, *Nature Machine Intelligence*, **1** (9) (2019) 389–399. <https://doi.org/10.1038/s42256-019-0088-2>
 - 20 Yeung K, Algorithmic regulation, *Regulation and Governance*, **12** (4) (2018) 505–523. <https://doi.org/10.1111/rego.12108>
 - 21 Foucault M, Governmentality, *In The Foucault Effect: Studies in Governmentality*, edited by Burchell G, Gordon C and Miller P (University of Chicago Press; Chicago), 1991, p. 87–104.
 - 22 Hajer M, *The politics of environmental discourse* (Oxford University Press; Oxford), 1995.
 - 23 Jasanoff S, *The ethics of invention: Technology and the human future* (W.W. Norton & Company; New York), 2016.
 - 24 Mazzucato M, *The value of everything: Making and taking in the global economy* (Allen Lane; London), 2018.
 - 25 Joshi D, The Paris AI Summit: is geopolitical rivalry derailing AI governance? Available at <https://odi.org/en/insights/the-paris-ai-summit-is-geopolitical-rivalry-derailing-ai-governance/>
 - 26 Grimmer J and Stewart B M, Text as data: the promise and pitfalls of automatic content analysis methods for political texts, *Political Analysis*, **21** (3) (2013) 267–297. <https://doi.org/10.1093/pan/mps028>
 - 27 Roberts M E, Stewart B M, Tingley D, Lucas C, Leder-Luis J, Gadarian S K, Albertson B, *et al.*, Structural topic models for open-ended survey responses, *American Journal of Political Science*, **58** (4) (2014) 1064–1082. <https://doi.org/10.1111/ajps.12103>
 - 28 Jacobs T and Tschötschel R, Topic models meet discourse analysis, *Discourse Studies*, **21** (4) (2019) 469–491. <https://doi.org/10.1177/1461445619832254>
 - 29 NITI Aayog, *Approach document for India: Part 1 – Principles for Responsible AI* (Government of India), 2021. <https://www.niti.gov.in/sites/default/files/2023-02/Responsible-AI-22022021.pdf>
 - 30 NITI Aayog, *Approach document for India: Part 2 – Operationalizing Principles for Responsible AI* (Government of India), 2021. <https://www.niti.gov.in/sites/default/files/2023-02/Responsible-AI-22022021.pdf>
 - 31 NITI Frontier Tech Hub, *AI for Viksit Bharat: The opportunity for accelerated economic growth* (NITI Aayog), 2025. <https://www.niti.gov.in/whats-new/ai-viksit-bharat-opportunity-accelerated-economic-growth>
 - 32 NITI Frontier Tech Hub, *Roadmap for job creation in the AI economy* (NITI Aayog), 2025. <https://www.niti.gov.in/whats-new/roadmap-job-creation-ai-economy>
 - 33 NITI Frontier Tech Hub, *AI for inclusive societal development* (NITI Aayog), 2025. <https://niti.gov.in/whats-new/roadmap-ai-inclusive-societal-development>
 - 34 Ministry of Electronics and Information Technology (MeitY), *Report of Committee-D on cyber security, safety, legal and ethical issues* (Government of India). https://www.meity.gov.in/writereaddata/files/Committee_D-Cyber-n-Legal-and-Ethical.pdf
 - 35 NITI Aayog and UK Department for International Trade, *Healthcare Artificial Intelligence Catalyst (HAIC) pilot study: Project report* (Government of India), 2021. <https://www.niti.gov.in/node/342>
 - 36 Sumit K, Buragohain P and Singh M K, AI governance in scholarly publishing: a computational analysis of policy structures and gaps, *Annals of Library and Information Studies*, **73** (2) (2026) 195–204. <https://doi.org/10.56042/al.v73i2.31022>