



The accountability gap in AI-enabled welfare governance in India: State responsibility, human oversight and the right to challenge algorithmic decisions

Yogita

Department of Political Science, Chaudhary Charan Singh University Meerut, Uttar Pradesh, India

Abstract

The growing use of Artificial Intelligence (AI) in welfare governance is transforming the way governments identify beneficiaries, detect irregularities and deliver public services. While AI can improve efficiency and administrative capacity, it also raises concerns about transparency, human oversight and responsibility for decisions affecting citizens. This article examines the accountability gap in AI-enabled welfare governance in India, with particular attention to algorithmic decision-making, data-related errors, institutional responsibility and grievance redressal. Using a qualitative and exploratory approach based on policy documents, government reports and relevant academic literature, the study argues that technological efficiency alone cannot ensure accountable welfare governance. It highlights the need for meaningful human review, transparent decision-making, reliable data, clear institutional responsibility and accessible mechanisms through which citizens can challenge adverse decisions. The article concludes that India's transition towards AI-enabled welfare governance must be guided by a citizen-centred approach in which technological innovation develops alongside democratic accountability and social justice.

Keywords: Artificial Intelligence, welfare governance, algorithmic accountability, public policy, administrative accountability, human oversight, social justice, India

Introduction

Welfare governance is an important responsibility of the state, particularly in a country like India where millions of citizens depend on public programmes for healthcare, food security, social protection and financial assistance. In recent years, digital technologies, data analytics and Artificial Intelligence (AI) have increasingly become part of welfare administration. AI offers considerable possibilities for improving the identification of beneficiaries, detecting irregularities, managing large datasets and making service delivery more efficient. However, its growing role in public decision-making also raises an important question: who is accountable when an AI-assisted welfare decision is wrong? Traditionally, administrative responsibility could be traced through government officials, departments and established procedures. With algorithmic systems, this chain can become more complicated. A welfare decision may involve government authorities, data managers, technology providers and AI-based systems. Consequently, a citizen who is denied a welfare benefit may find it difficult to understand how the decision was reached, whether an algorithm influenced it, or which authority is responsible for correcting the error. This creates what can be described as an "accountability gap."

The issue becomes particularly significant for vulnerable sections of society. An inaccurate database, biased algorithm or incorrect classification can result in the exclusion of an eligible person from essential welfare benefits. Such exclusion is not merely a technological error; it can have serious social and economic consequences. Moreover, the presence of a human official in the decision-making process does not automatically guarantee accountability if officials simply accept algorithmic recommendations without meaningful scrutiny.

The Indian context presents an additional challenge because welfare governance involves multiple levels of government

and implementing agencies. Differences in technological capacity between States can also produce uneven outcomes. Therefore, the adoption of AI raises questions not only about efficiency and innovation but also about administrative responsibility, social justice, transparency and citizens' right to challenge state decisions.

Against this background, this study examines the accountability gap in AI-enabled welfare governance in India. It focuses on whether existing mechanisms of transparency, human oversight and grievance redressal are adequate when AI and data-driven systems influence welfare decisions. The study does not view AI as inherently beneficial or harmful; rather, it examines how institutional safeguards determine its consequences for citizens.

The central argument is that the success of AI-enabled welfare governance should not be measured only through efficiency, speed and accuracy. It must also be judged by whether citizens can understand, question and challenge decisions that affect their access to public welfare. India's challenge, therefore, is not simply to create a technologically advanced welfare system, but to ensure that technological innovation develops alongside democratic accountability and protection of citizen rights.

Background and Context of AI-Enabled Welfare Governance in India

The transformation of welfare governance in India has taken place gradually, beginning with the wider adoption of digital technologies in public administration and moving towards data-driven and increasingly intelligent systems. The introduction of digital identification, online databases, direct benefit transfer mechanisms and integrated service-delivery platforms has changed the manner in which governments identify beneficiaries and distribute public resources. These developments have sought to reduce leakages, improve administrative efficiency and make welfare delivery more

transparent. At the same time, they have increased the dependence of welfare administration on digital infrastructure and the quality of government data.

The growing availability of large datasets has created new possibilities for the use of Artificial Intelligence in public administration. AI-based tools can assist government agencies in identifying unusual patterns, detecting potential fraud, assessing applications and allocating administrative resources. In welfare programmes, such technologies can potentially help governments reach beneficiaries more efficiently and identify irregularities that may be difficult to detect through conventional administrative methods. The use of AI in healthcare-related welfare administration, for instance, demonstrates how data analytics and automated systems can support the detection of potentially fraudulent claims while allowing government officials to undertake further investigation.

However, the increasing use of AI also introduces a new layer of complexity into welfare administration. Conventional administrative decisions generally involve identifiable officials who can provide reasons for their actions. Algorithmic systems, in contrast, may involve complex models and large datasets whose functioning is not easily understood by ordinary citizens. Even when the final decision remains with a government official, the recommendation generated by an algorithm may influence the official's judgement. This can make it difficult to determine where administrative responsibility actually lies. The problem becomes more serious when AI systems rely on inaccurate, incomplete or outdated information. A citizen may be incorrectly classified because of an error in the underlying database rather than because of any deliberate administrative decision. If the individual does not know that an automated system contributed to the decision, the opportunity to challenge the error becomes even more limited. Thus, data accuracy, algorithmic transparency and accessibility of grievance mechanisms become essential components of accountable welfare governance.

Another important dimension is the changing relationship between the state and private technology providers. Governments may depend on external companies for software development, data management, algorithmic tools and technical expertise. While such partnerships can increase technological capacity, they can also complicate accountability. A public authority cannot simply transfer responsibility for a welfare decision to a private technology provider because the ultimate responsibility for ensuring fair and lawful welfare administration remains a public obligation.

India's federal structure adds another dimension to this issue. Welfare schemes may involve the Union government, State governments and local administrative institutions, each with different technological capabilities and administrative arrangements. If AI systems are designed centrally but implemented differently across States, questions may arise regarding responsibility for errors, monitoring standards and grievance redressal. Unequal technological capacity among States may also create differences in the quality and accessibility of AI-enabled welfare services.

The emerging policy environment therefore requires a shift from viewing AI merely as an instrument of administrative efficiency towards understanding it as part of a broader governance framework. The important question is not simply whether AI can improve welfare delivery, but

whether it can do so while preserving transparency, fairness and citizens' ability to seek remedies. An accountable system must ensure that individuals are not left without an effective remedy merely because a decision has been influenced by an automated process.

This background establishes the central concern of the present study: technological innovation in welfare administration must be accompanied by institutional mechanisms capable of assigning responsibility, explaining decisions and protecting citizens from unjust exclusion. The challenge is therefore to develop a model of AI-enabled welfare governance in which efficiency and innovation are balanced with accountability, human oversight and social justice.

Conceptual and Theoretical Framework

The use of Artificial Intelligence in welfare administration cannot be understood only as a technological development. It represents a change in the way public institutions collect information, exercise administrative authority and interact with citizens. Therefore, examining AI-enabled welfare governance requires a framework that connects technological decision-making with established ideas of accountability, administrative discretion and democratic governance.

Understanding Accountability in Public Governance

Accountability is a fundamental principle of democratic administration. In its simplest sense, it means that those exercising public authority must be able to explain and justify their decisions and should be answerable for their actions. In welfare administration, accountability is particularly important because government decisions directly affect citizens' access to public resources and essential services.

The introduction of AI complicates this traditional understanding. When an algorithm contributes to a decision, responsibility may become distributed among several actors. The government may own the welfare programme, a public official may implement it, a technology company may develop the system, and another agency may manage the underlying data. This makes it difficult to identify who should answer when an automated recommendation produces an incorrect or discriminatory outcome.

For the purpose of this study, accountability is therefore understood through four interconnected elements: answerability, transparency, responsibility and remedy. A welfare system cannot be considered genuinely accountable merely because an authority formally exists. Citizens must also be able to understand the basis of a decision, identify the responsible institution and obtain an effective remedy when an error occurs.

Algorithmic Governance and the Changing Role of the State

Algorithmic governance refers to the increasing use of computational systems, data and automated decision-making tools in the exercise of public authority. Algorithms can assist governments in identifying patterns, assessing applications, detecting irregularities and prioritising administrative action. Their use can reduce the burden on public officials and allow governments to process information at a scale that would otherwise be difficult to manage.

However, algorithmic governance can also change the nature of administrative power. Decisions that were previously based primarily on human judgement may increasingly depend on data-driven classifications and recommendations. This does not necessarily eliminate bureaucracy; rather, it can reconfigure bureaucratic discretion. Officials may retain formal decision-making authority while becoming increasingly dependent on technological recommendations.

This raises an important question for welfare governance: does the introduction of AI reduce administrative discretion, or does it simply move discretion to a less visible technological layer? The answer has significant implications for accountability because decisions made through complex systems may be more difficult for citizens to question than conventional administrative decisions.

Street-Level Bureaucracy and Human Oversight

The theory of street-level bureaucracy provides another useful perspective. Welfare policies are ultimately implemented through officials who interact with citizens and interpret administrative rules in particular circumstances. These officials traditionally exercise a degree of discretion because policies cannot anticipate every situation encountered on the ground.

AI introduces a new dimension to this relationship. An official may receive an automated recommendation regarding a person's eligibility, risk level or potential irregularity. While the official may technically retain the authority to accept or reject the recommendation, institutional pressures, workload and confidence in technological systems may encourage reliance on the algorithm.

Therefore, the presence of a human decision-maker should not automatically be treated as evidence of meaningful human oversight. Genuine oversight requires the official to have the authority, information and institutional capacity to question an automated recommendation. Otherwise, the human official may merely become the final point in an automated chain rather than an independent safeguard for the citizen.

Social Justice and the Vulnerable Citizen

A welfare system must ultimately be evaluated in terms of its consequences for the people it is designed to serve. This makes social justice an important component of the theoretical framework.

AI systems are dependent on data, and data can contain errors or reflect existing patterns of inequality. If these problems are incorporated into automated decision-making, technologically sophisticated systems may unintentionally reproduce existing disadvantages. The consequences can be particularly serious for individuals who have limited digital literacy, inadequate documentation or restricted access to administrative institutions.

From a social-justice perspective, therefore, the relevant question is not simply whether an algorithm is technically accurate. It is whether its operation produces fair and accessible outcomes across different sections of society. A system that performs efficiently for the majority but repeatedly excludes vulnerable groups cannot be considered fully successful as a welfare-governance mechanism.

The Accountability Gap: A Conceptual Proposition

Drawing these perspectives together, this study conceptualises the accountability gap as the distance between the growing capacity of AI-enabled systems to influence welfare decisions and the institutional mechanisms available to explain, challenge and correct those decisions.

The gap becomes particularly visible when four conditions occur simultaneously:

First, an AI system influences an administrative decision.

Second, the basis of that decision is difficult for the affected citizen to understand.

Third, responsibility is distributed among several governmental and technological actors.

Fourth, the citizen lacks a simple and effective mechanism through which the decision can be reviewed.

The study therefore proposes that AI-enabled welfare governance should be considered accountable only when technological decision-making is accompanied by identifiable institutional responsibility, meaningful human oversight, transparency, accessible grievance redressal and the possibility of correction.

AI-Enabled Welfare Governance in India: Policy and Institutional Context

India's welfare administration has undergone a significant transformation with the expansion of digital technologies. Over the past decade, government programmes have increasingly relied on digital databases, online platforms, direct benefit transfers and data-based monitoring. These developments have improved the ability of public institutions to manage large welfare programmes, but they have also made the delivery of public services increasingly dependent on the accuracy and reliability of digital systems. The emergence of Artificial Intelligence represents a further shift in this process. AI and data analytics are increasingly being explored for identifying irregularities, improving beneficiary targeting, monitoring welfare programmes and supporting administrative decision-making. The IndiaAI Mission and related government initiatives indicate a broader policy interest in developing an ecosystem for responsible and large-scale use of AI. This suggests that AI is gradually moving from an experimental technology towards a potential component of mainstream public administration.

The use of AI in welfare governance can offer several advantages. Government agencies deal with extremely large and diverse datasets, making conventional methods of analysis time-consuming and resource-intensive. AI can assist in identifying patterns that may otherwise remain unnoticed, detecting potentially fraudulent transactions and helping officials prioritise cases requiring further examination. In principle, such applications can reduce administrative burdens and allow public resources to be used more effectively.

However, efficiency cannot be the sole measure of successful welfare governance. The consequences of an incorrect algorithmic assessment can be serious when it affects a person's access to a public benefit. A citizen may be incorrectly classified because of inaccurate data, an inappropriate model or an error in the system. The problem becomes more difficult when the individual does not know

that an automated system influenced the decision or does not have an accessible mechanism to challenge it.

India's institutional structure further complicates the issue. Welfare delivery frequently involves the Union government, State governments, local institutions and implementing agencies. AI systems may be developed or supported at one level while their consequences are experienced at another. Differences in technical expertise, infrastructure and administrative capacity among States can therefore influence how effectively AI is implemented and monitored. Another concern is the involvement of private technology providers. Government agencies may depend on external organisations for developing software, managing databases or providing AI-related services. Such partnerships can provide valuable expertise, but they can also create uncertainty regarding responsibility when a technological system produces an undesirable outcome. The use of private technology does not remove the public authority's responsibility towards the citizen receiving welfare services. For this reason, India's emerging AI governance framework needs to be examined not only from the perspective of innovation but also through the principles of public accountability and citizen protection. A responsible welfare system should provide clarity about when AI is being used, ensure appropriate human oversight, maintain mechanisms for correcting inaccurate data and provide citizens with meaningful opportunities to challenge adverse decisions. The central policy challenge is therefore to ensure that India's movement towards AI-enabled welfare governance does not create a situation in which technological complexity becomes a barrier to accountability. The objective should be to develop systems where AI supports public officials without making responsibility invisible. Technology should strengthen the capacity of the welfare state, but it should not weaken the state's obligation to remain answerable to its citizens.

Dimensions of the Accountability Gap in AI-Enabled Welfare Governance

The accountability gap in AI-enabled welfare governance does not arise from a single problem. It develops through several connected stages, beginning with the collection of data and extending to the final decision affecting a citizen. Understanding these stages is important because an error at any point can influence the outcome of welfare delivery.

Lack of Transparency

One of the most significant concerns is the limited transparency surrounding algorithmic decision-making. A citizen may be informed that an application has been rejected or a benefit has been withheld without being given a clear explanation of how the decision was reached. Complex algorithms can make it difficult for ordinary citizens to understand the factors influencing an outcome. Transparency therefore requires more than simply disclosing that AI is being used. Citizens should be able to know when an automated system has influenced a decision, what broad factors were considered, and which authority remains responsible for the final outcome.

Data Errors and Algorithmic Bias

The effectiveness of an AI system depends heavily on the quality of the data available to it. Inaccurate, incomplete or outdated government records can produce incorrect classifications. If such problems disproportionately affect

disadvantaged groups, technological systems may unintentionally reinforce existing inequalities.

This is particularly important in welfare administration because those who depend most heavily on public support may have fewer resources to correct errors. Therefore, regular data verification and mechanisms for citizens to update or challenge incorrect information are essential.

Diffusion of Responsibility

AI-enabled welfare decisions can involve several actors, including government departments, officials, database managers, technology providers and software developers. When an error occurs, this distribution of responsibility can make it difficult to determine who should be held accountable.

A government department may attribute the problem to the technology provider, while the technology provider may argue that the system only followed the data or instructions supplied by the government. Such situations can create a responsibility vacuum in which the affected citizen struggles to obtain a clear answer or remedy.

Limits of Human Oversight

Human involvement is often presented as an important safeguard against algorithmic errors. However, meaningful oversight requires more than simply placing an official at the end of an automated process.

Officials must have the authority and ability to question an AI-generated recommendation. They should also have sufficient information to understand why a recommendation was produced. If officials routinely accept algorithmic outputs because of workload, institutional pressure or excessive trust in technology, human oversight may become largely symbolic.

Grievance Redressal and the Right to Challenge

An accountable welfare system must provide citizens with a realistic opportunity to challenge decisions affecting their entitlements. Traditional grievance mechanisms may become inadequate if they do not address decisions influenced by automated systems.

A citizen should not be required to understand complex algorithms in order to challenge an adverse welfare decision. The responsibility should remain with the public authority to provide a comprehensible explanation, conduct a human review and correct demonstrable errors.

Unequal Technological Capacity

The accountability problem also has a federal dimension. Indian States differ considerably in administrative capacity, digital infrastructure, technical expertise and financial resources. Consequently, the adoption of AI may not produce uniform outcomes across the country.

States with stronger technological capacity may be better equipped to audit systems, monitor vendors and establish effective grievance mechanisms, while weaker administrative environments may become more dependent on external technology providers. This creates a possibility that technological inequality could translate into unequal standards of welfare governance.

The Emerging Policy Challenge

These dimensions demonstrate that the accountability gap is not simply a problem of whether an algorithm is accurate. It is fundamentally an institutional problem. A technically

efficient system can still produce an unfair outcome if citizens cannot understand, question or correct its decisions. AI-enabled welfare governance therefore requires a shift from a narrow focus on technological performance towards a broader framework based on transparency, responsibility, human oversight, auditability and accessible grievance redressal. The central challenge for India is to ensure that technological innovation strengthens welfare administration without making the exercise of public power less visible or less answerable to citizens.

Research Methodology

The present study follows a qualitative and exploratory approach to understand how the growing use of Artificial Intelligence may be changing accountability within welfare governance in India. Since the study is concerned with questions of responsibility, transparency and citizen rights, it focuses more on understanding institutional processes than on measuring the technical accuracy of individual AI systems.

The research primarily uses secondary sources, including government reports, policy documents, parliamentary committee reports, official publications, academic books and research articles. Documents relating to India's emerging AI governance framework and selected welfare programmes are examined to understand how technology is being incorporated into public administration and what safeguards are available to citizens.

The study uses a purposive case-study approach. Selected examples of AI or advanced data-driven systems in welfare administration are examined because they provide useful insights into issues such as beneficiary identification, fraud detection, automated recommendations and human decision-making. The purpose is not to compare every welfare programme in India, but to identify recurring accountability concerns across relevant cases.

The information collected is examined through thematic analysis. The study looks particularly at five themes: transparency, responsibility, human oversight, data-related risks and grievance redressal. These themes help in understanding where accountability may become weakened when technological systems become involved in welfare decisions.

An interpretive perspective is also adopted because accountability cannot be understood only as a technical or administrative requirement. It is closely connected with the relationship between the state and its citizens. A welfare system may be technologically efficient, yet still raise serious governance concerns if an affected person cannot understand a decision, identify who is responsible for it, or obtain an effective remedy.

The study is subject to certain limitations. AI-based welfare systems are developing rapidly, and detailed information about their algorithms, datasets and internal processes is not always publicly available. The research therefore relies on information that can be accessed and examined through official and credible secondary sources. Despite this limitation, the approach provides a useful basis for examining the emerging accountability gap and for suggesting institutional safeguards for more citizen-centred AI-enabled welfare governance in India.

Emerging Applications of AI in Indian Welfare Governance

The use of Artificial Intelligence in Indian welfare governance is still developing, but its growing presence can already be seen in areas such as healthcare, fraud detection, beneficiary management and data-based service delivery. These developments are important because they demonstrate that AI is gradually moving from a purely technological discussion into the practical functioning of public administration.

One useful example is the use of AI and data analytics in healthcare-related welfare programmes. Large public health schemes generate substantial volumes of claims and beneficiary information. Analysing this information manually can be difficult for administrative agencies. AI-based systems can assist in identifying unusual patterns and flagging claims that may require further examination. Such applications can potentially reduce fraudulent practices and help public authorities use limited welfare resources more effectively.

However, an important distinction needs to be maintained between identifying a case for further investigation and making a final decision against a beneficiary. An algorithmic alert should not automatically be treated as evidence of wrongdoing. If a citizen's access to a welfare service is restricted solely because a system has classified a transaction as suspicious, the possibility of an incorrect exclusion becomes significant. Human verification is therefore essential, particularly when the decision can affect an individual's access to an important public service.

A second emerging area is the use of AI and advanced analytics for social-welfare delivery at the State level. Recent initiatives demonstrate that State governments are beginning to explore data-driven tools to improve the identification of beneficiaries, monitor programmes and strengthen accountability. Such developments reflect an important shift: governments are no longer viewing digital technology simply as a means of transferring benefits electronically, but increasingly as a tool for analysing welfare systems themselves.

This shift, however, creates new administrative responsibilities. When AI is used to identify patterns or make recommendations, officials need to understand the limitations of the technology. A system may identify a statistical pattern without understanding the circumstances of an individual household. Welfare decisions, in contrast, often involve social and economic conditions that cannot be completely captured through numerical data.

The use of AI also raises questions about data dependency. Welfare decisions are increasingly influenced by information contained in government databases. If that information is incomplete or outdated, the resulting technological assessment may also be unreliable. The affected citizen may then face the difficult task of proving that the underlying data are incorrect. This is particularly problematic where there is no simple process for reviewing or correcting information.

Another concern is the increasing involvement of technology companies and external technical agencies in public-service delivery. Their expertise can be valuable for developing sophisticated systems, but public authorities must retain clear responsibility for decisions affecting citizens. A technology provider may design or maintain a system, but responsibility for the fairness and legality of

welfare administration cannot simply be transferred to a private actor.

These emerging applications demonstrate that the accountability challenge is not necessarily created by the use of AI itself. Rather, it depends on how AI is incorporated into the existing administrative structure. Where AI is used as a supporting tool, with careful human verification and clear avenues for appeal, it may strengthen welfare administration. Where automated recommendations become difficult to question or where responsibility is dispersed among several actors, the risk of an accountability gap increases.

The Indian experience therefore points towards the need for a balanced approach. AI can contribute to more efficient welfare governance, but technological efficiency must remain connected to administrative responsibility, human judgement and citizen protection. The crucial issue for future policy is not whether AI should be used in welfare administration, but under what institutional conditions its use can remain consistent with the principles of fairness and accountable governance.

Analysis and Discussion

The preceding discussion indicates that the use of AI in welfare governance presents both an opportunity for administrative improvement and a new challenge for public accountability. The technology can help governments process information more quickly and identify patterns that may be difficult to detect through conventional methods. However, efficiency alone cannot determine whether a welfare system is functioning fairly. The more important question is whether technological decision-making remains subject to human judgement, institutional responsibility and citizen scrutiny.

One of the central concerns is the shift from visible administrative discretion to less visible algorithmic influence. Traditionally, an official making a welfare decision could be questioned about the reasons behind it. When an AI system generates a recommendation, the decision may appear objective simply because it is based on data. Yet data-driven decisions are not necessarily neutral. They reflect the quality of the information used, the assumptions incorporated into the system and the way particular categories have been defined.

This becomes especially significant when welfare beneficiaries belong to socially or economically vulnerable groups. A person who is incorrectly excluded from a welfare programme may not have the knowledge or resources required to identify whether the problem originated in the database, the algorithm or the administrative process. Consequently, the same technology that is intended to make welfare delivery more efficient can, in certain circumstances, make exclusion more difficult to detect and challenge.

The question of human oversight is therefore central. Simply retaining a government official in the final stage of decision-making does not necessarily solve the accountability problem. Meaningful oversight requires officials to examine algorithmic recommendations critically rather than treating them as unquestionable outcomes. They must have access to sufficient information and the authority to reconsider a decision when circumstances indicate that the technological assessment may be incorrect.

Another important finding concerns the distribution of responsibility. AI-enabled welfare systems often involve several institutions, making it possible for responsibility to become fragmented. A government department may rely on an external technology provider, while officials depend on databases managed by another agency. When an adverse decision occurs, this complexity can make it difficult for citizens to identify the appropriate authority. An accountable governance structure must therefore establish responsibility before problems occur rather than attempting to assign blame after an error has already affected a citizen.

The analysis also highlights the importance of grievance redressal. A welfare system cannot be considered citizen-centred if individuals have no practical means of challenging an AI-assisted decision. The right to seek review should not depend on the citizen's ability to understand complex technological processes. Instead, public institutions should provide a clear route through which individuals can request human review, correct inaccurate information and receive a reasoned response.

The federal nature of Indian governance adds another layer to the discussion. States differ in administrative capacity, technological infrastructure and access to technical expertise. A uniform technological model may therefore produce different outcomes in different administrative environments. Stronger States may be able to monitor AI systems more effectively, whereas weaker institutions may become highly dependent on external providers. This suggests that responsible AI governance also requires capacity-building at the State and local levels.

Overall, the analysis suggests that the accountability gap is not an unavoidable consequence of AI. It is largely a consequence of institutional design. If AI is introduced without clear responsibility, transparency and effective safeguards, it can make administrative decisions more difficult to understand and challenge. If, however, technological systems are accompanied by meaningful human oversight, independent monitoring, accessible grievance mechanisms and clearly assigned responsibility, AI can support rather than weaken accountable welfare governance.

The central implication is therefore that India should not evaluate AI-enabled welfare systems only by asking whether they are faster or more accurate than conventional methods. The more important measure is whether they remain fair, explainable and answerable to the citizens whose lives they affect. This requires placing accountability at the centre of technological innovation rather than treating it as an issue to be addressed after the technology has already been implemented.

Policy Recommendations

The growing use of AI in welfare governance does not require the State to choose between technological innovation and citizen protection. The two can develop together if appropriate institutional safeguards are built into welfare systems from the beginning. Based on the issues discussed above, several measures can help reduce the accountability gap.

Make Human Review Mandatory for Adverse Decisions

AI should primarily support administrative decision-making rather than replace it, particularly where a decision may result in the denial or withdrawal of an important welfare

benefit. Every significant adverse decision influenced by an automated system should be subject to meaningful human review. The reviewing official should have the authority to reconsider the recommendation rather than merely approve it.

Establish Clear Institutional Responsibility

Responsibility should be clearly assigned before an AI system is introduced into a welfare programme. Citizens should not have to determine whether a government department, software provider or data-management agency is responsible for an error. The concerned public authority should remain accountable for the final decision, even when technological systems are developed or operated by external organisations.

Strengthen Transparency and Explainability

Citizens should be informed when AI has materially influenced a welfare decision. They should also receive a simple explanation of the principal factors behind an adverse outcome. Complete disclosure of complex technical details may not always be practical, but meaningful explanation should be possible without requiring citizens to possess specialised technological knowledge.

Introduce Regular Algorithmic Audits

AI systems used in welfare administration should be reviewed periodically for accuracy, bias, exclusion and unintended consequences. Such assessments should not be limited to technical performance. They should also examine whether particular social or economic groups are disproportionately affected by errors. Independent or external auditing can strengthen public confidence in these systems.

Create Accessible Grievance Mechanisms

An effective grievance mechanism should be simple, affordable and available through both digital and non-digital channels. Citizens should be able to request human review, correct inaccurate information and receive a reasoned response. Vulnerable beneficiaries should not be forced to navigate complicated technological procedures merely to challenge an administrative decision.

Improve Data Quality and Citizen Control

Since AI systems depend heavily on government data, mechanisms for correcting inaccurate or outdated information are essential. Citizens should have reasonable opportunities to verify information. Responsibility should be clearly assigned before an AI system is introduced into a welfare programme. Citizens should not have to determine whether a government department, software provider or data-management agency is responsible for an error. The concerned public authority should remain accountable for the final decision, even when technological systems are developed or operated by external organisations.

Strengthen Transparency and Explainability

Citizens should be informed when AI has materially influenced a welfare decision. They should also receive a simple explanation of the principal factors behind an adverse outcome. Complete disclosure of complex technical

details may not always be practical, but meaningful explanation should be possible without requiring citizens to possess specialised technological knowledge.

Introduce Regular Algorithmic Audits

AI systems used in welfare administration should be reviewed periodically for accuracy, bias, exclusion and unintended consequences. Such assessments should not be limited to technical performance. They should also examine whether particular social or economic groups are disproportionately affected by errors. Independent or external auditing can strengthen public confidence in these systems.

Create Accessible Grievance Mechanisms

An effective grievance mechanism should be simple, affordable and available through both digital and non-digital channels. Citizens should be able to request human review, correct inaccurate information and receive a reasoned response. Vulnerable beneficiaries should not be forced to navigate complicated technological procedures merely to challenge an administrative decision.

Improve Data Quality and Citizen Control

Since AI systems depend heavily on government data, mechanisms for correcting inaccurate or outdated information are essential. Citizens should have reasonable opportunities to verify important information used in welfare decisions and request corrections where necessary. This can reduce the possibility of technological systems reproducing administrative errors.

Conclusion

The growing use of Artificial Intelligence in welfare governance marks an important change in the way the Indian state delivers public services. AI can help governments manage large volumes of information, identify irregularities and improve administrative efficiency. These possibilities are significant, particularly for a country where welfare programmes serve millions of people. Yet technological progress by itself cannot guarantee better governance.

This study has argued that the most important challenge lies in the accountability gap that can emerge when AI becomes involved in welfare decisions. When responsibility is distributed among government departments, officials, data systems and technology providers, it can become difficult for citizens to understand who is responsible for an adverse decision. This problem becomes more serious when the decision is based on inaccurate data, an opaque algorithm or an automated recommendation that receives insufficient human scrutiny.

The analysis also demonstrates that human oversight, transparency and grievance redressal must be treated as essential components of AI-enabled welfare governance. The presence of a human official at the end of an automated process is not sufficient unless that official has the authority and capacity to question the technological recommendation. Similarly, transparency should mean more than simply announcing that AI is being used; citizens should be able to understand when technology has influenced a decision and have a meaningful opportunity to challenge it.

For India, the challenge is particularly significant because welfare governance operates across different levels of

government and technological capacity varies among States. The increasing involvement of private technology providers adds another layer of responsibility. These developments make it necessary to establish clear institutional arrangements so that technological complexity does not become a means of avoiding public accountability.

The study therefore concludes that AI should strengthen the capacity of the welfare state without weakening its responsibility towards citizens. Future AI-enabled welfare systems should be designed around clear responsibility, meaningful human review, reliable data, independent auditing and accessible grievance mechanisms. Most importantly, no citizen should be left without an effective remedy simply because a decision has been influenced by an automated system.

The real measure of India's progress in AI-enabled welfare governance will consequently not be how much technology the state adopts, but how responsibly that technology is used. A technologically advanced welfare system will have lasting legitimacy only when efficiency is accompanied by fairness, transparency and the continuing ability of citizens to question the exercise of public power.

Consultation on Strengthening Accountability Systems, 2026.

References

1. Chaudhuri B. Programmed welfare: An ethnographic account of algorithmic practices in the public distribution system in India. *Social Media + Society*, 2022, 8(4).
<https://doi.org/10.1177/14614448221079034>
2. Garg M, Lanka SN. Automating the state: Public sector AI in India and the challenge of rights-based regulation. *Indian Journal of Law and Technology*, 2025, 21(1).
3. Ministry of Electronics and Information Technology. India AI governance guidelines. Government of India, 2025.
4. Ministry of Electronics and Information Technology. IndiaAI Mission and AI governance initiatives. Government of India, 2026.
5. NITI Aayog. Approach document for India: Part 1—Principles for responsible AI. Government of India, 2021.
6. NITI Aayog. Approach document for India: Part 2—Operationalizing principles for responsible AI. Government of India, 2021.
7. NITI Aayog. Responsible AI: AI for all. Government of India, 2021.
8. NITI Aayog. Towards responsible AI: Enforcement of principles. Government of India, 2021.
9. Press Information Bureau. Artificial Intelligence (AI) transforming rural India. Government of India, 2026.
10. Sharma S. Automating the state: Promise, peril, and the politics of AI in public administration. SSRN, 2026.
<https://doi.org/10.2139/ssrn.6391458>
11. Sharma R. Legal accountability in algorithmic governance and public decision-making in administrative law. *Legality*, 2026, 1(1).
12. Algorithmic solutions, subjectivity and decision errors: A study of AI accountability. *Digital Policy, Regulation and Governance*, 2025;27(5):523–552.
<https://doi.org/10.1108/DPRG-05-2024-0090>
13. Government of Kerala. AI and data analytics for social welfare delivery. Kerala State IT Mission, 2026.
14. National Law School of India University. Algorithmic accountability in welfare and identity systems.