IJIRSET
International Journal of Innovative Research in Science | Engineering | Technology

ISSN(O): 2319-8753
ISSN(P): 2347-6710



International Journal of Innovative Research in Science Engineering and Technology (IJIRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)



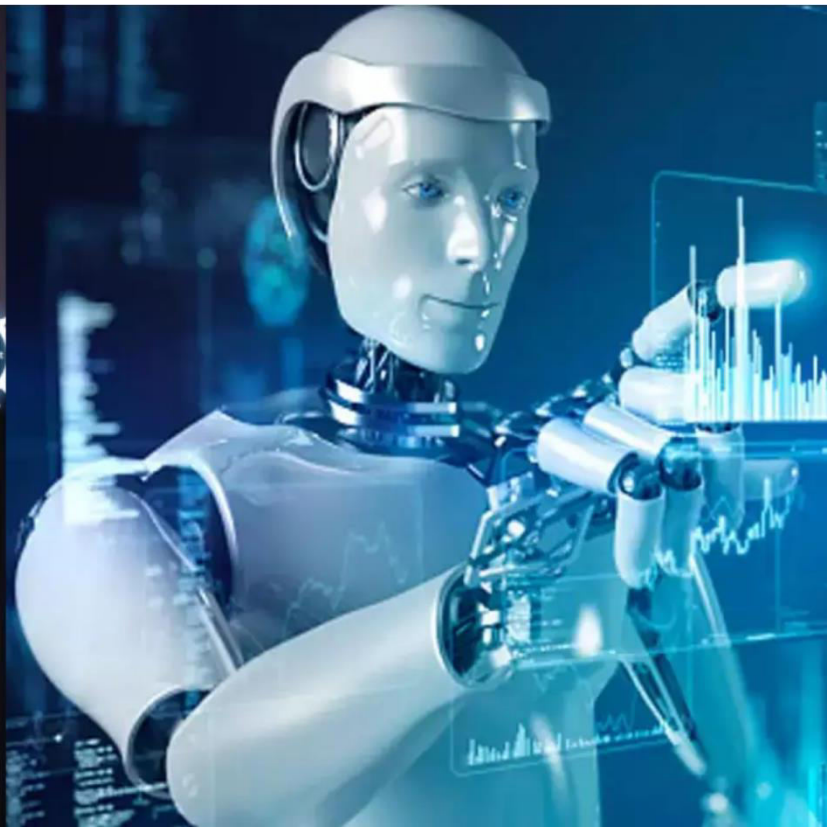

Impact Factor: 9.935

Volume 15, Issue 11, November 2026

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 9.935

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 11, November 2026

|DOI: 10.15680/IJIRSET.2026.1511040|

AI-Based Personalized Government Welfare Benefits Recommendation System

Dr. P. S. Deshmukh[1], Pratansha Sandip Chaudhary[2], Mansvi Vinod Kadu[3], Riya Bhimrao Ramteke[4], Pratik Sanjay Malve[5]

Guide, Department of Computer Science, G.H. Raisoni University, Amravati, India[1]

Student, Department of Computer Science and Engineering, Prof. Ram Meghe Institute of Technology and Research, Badnera, Amravati, India[2-5]

ABSTRACT: Government welfare programmes provide important support in education, agriculture, healthcare, housing, employment, entrepreneurship and social security. However, citizens may find it difficult to identify schemes relevant to their circumstances because scheme information is distributed across multiple portals and documents, while eligibility conditions and application requirements can be difficult to interpret. This paper presents an AI-based personalized government welfare benefits recommendation system intended to simplify scheme discovery through a unified, explainable recommendation process. The proposed framework uses two sequential stages. A rule-based eligibility engine first evaluates mandatory conditions such as age, income, occupation, residence, gender, social category and other scheme-specific requirements. Schemes that pass this screening are then ranked using content-based filtering and cosine similarity between structured citizen-profile features and scheme attributes. An explanation component presents matched, unmatched and unverified conditions so that users can understand why a scheme appears in the results and what information may still be required. The proposed application is designed around a Java and Spring Boot backend, MySQL database and a responsive web interface using HTML, CSS, JavaScript and Bootstrap; Python and Scikit-learn may support model development and evaluation. The paper describes the motivation, related work, architecture, methodology, modules, data design, expected outcomes and limitations. The framework is a proposed system design; its effectiveness must be established through implementation and evaluation using verified scheme data and representative citizen profiles.

KEYWORDS: e-Governance, government welfare schemes, recommender system, rule-based eligibility, content-based filtering, cosine similarity, explainable AI, citizen profile.

I. INTRODUCTION

India has a broad welfare ecosystem covering agriculture, education, healthcare, housing, women empowerment, skill development, entrepreneurship and social security. Central and state governments publish schemes intended for different groups of citizens, including students, farmers, women, senior citizens, workers and economically weaker households. The challenge addressed in this work is not simply the availability of schemes; it is helping each citizen identify the programmes that may apply to their personal circumstances and understand how to proceed.

Information about schemes may be distributed across ministry websites, state portals, department pages and official documents. A citizen may need to search several sources, interpret administrative language, compare conditions and identify the documents required for an application. Conventional keyword-based search can be difficult when a person does not know the scheme name, department or correct terminology. As a result, potentially eligible citizens may remain unaware of support for which they could apply.

Digital platforms such as myScheme and UMANG provide access to government scheme and service information [9]. These platforms form important foundations for digital service delivery. The proposed research focuses on a complementary personalization layer: it uses a structured citizen profile to screen schemes against mandatory conditions, rank the schemes that pass the screening, and explain the result in understandable terms. The recommendation is intended to support discovery and should not be interpreted as an official approval or guarantee of benefit.

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 9.935

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 11, November 2026

|DOI: 10.15680/IJIRSET.2026.1511040|

The system follows a hybrid-by-function design. Deterministic rules are used where eligibility depends on explicit requirements, while content-based similarity is used to order the schemes that satisfy the mandatory checks. This separation is important because a high similarity score alone cannot establish legal or administrative eligibility. The explanation component further records which conditions were satisfied, not satisfied or could not be verified from the available profile.

II. PROBLEM STATEMENT AND RESEARCH GAP

The central problem is the gap between the availability of public welfare information and the ability of citizens to locate, understand and use the information relevant to them. A single scheme may include multiple conditions related to age, annual family income, occupation, residence, gender, social category, family circumstances, land ownership, education or employment. Checking these conditions manually across a large collection of schemes is time-consuming and may lead to missed opportunities or incorrect assumptions.

A. Limitations in Existing Approaches

Rule-based eligibility tools can check predefined conditions, but they may return a broad set of potentially eligible schemes without ranking them according to the citizen's interests or needs. General recommendation systems can rank items according to similarity or user behaviour, but recommendations made without a prior eligibility check may include schemes whose mandatory conditions are not met. Behaviour-based methods may also be limited when interaction history is sparse, particularly for new users or less frequently used public services. Chatbot-based assistance can improve access to information, but the quality of its answers depends on the accuracy, currency and traceability of the information supplied to it.

B. Identified Research Gap

The reviewed work indicates a need to integrate three functions in one process: (i) verifiable eligibility screening, (ii) relevance-based ranking of the schemes that pass screening, and (iii) explanations linked to the underlying eligibility criteria. The proposed system addresses this gap by keeping eligibility decisions separate from similarity ranking and by retaining a trace of the conditions used to produce each recommendation. Missing profile information is treated as unverified rather than silently assumed to be satisfied.

III. OBJECTIVES AND SCOPE

The primary objective is to design an explainable hybrid recommendation system that helps citizens discover government welfare schemes using structured personal and socioeconomic information. The specific objectives are:

1. Personalized recommendation: analyze relevant profile attributes, including age, income, occupation, education, family details, location, social category and other scheme-related information, to identify schemes relevant to the citizen.
2. Automated eligibility verification: compare profile values with mandatory scheme conditions and classify each condition as satisfied, unsatisfied or unverified. This makes the basis of the screening visible.
3. Similarity-based ranking: apply content-based filtering and cosine similarity to rank the schemes that pass the mandatory eligibility stage, rather than using similarity as a substitute for eligibility.
4. Explainable output: show the conditions that match, the conditions that do not match and the information that is missing or needs confirmation.
5. Accessible scheme information: present scheme benefits, required documents and application guidance in a structured interface. The proposed design can be extended to multilingual support, natural-language queries and updated scheme datasets.

The scope is a recommendation and information-support system. Final eligibility, application acceptance and benefit disbursement remain the responsibility of the relevant government department. The system requires verified and regularly updated scheme data; it cannot compensate for inaccurate source information.

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 9.935

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 11, November 2026

|DOI: 10.15680/IJIRSET.2026.1511040|

IV. LITERATURE REVIEW

Raza et al. [1] review the evolution of recommender systems, including content-based filtering, collaborative filtering, deep learning, graph-based recommendation, reinforcement learning and large language models. Their discussion highlights that recommendation quality depends on the application, available data and the need for transparency. Hossain et al. [2] survey recommendation techniques in the e-commerce domain and describe how filtering methods help users identify relevant information from large collections. These general approaches provide a technical foundation, but public welfare applications also require explicit handling of mandatory eligibility conditions.

Mandalageri [3] discusses a hybrid recommendation approach for government services using citizen demographics, service information and behavioural data. Combining profile and behavioural information can improve personalization, although the system must account for the availability and quality of interaction data. Naresh [4] considers eligibility matching between citizen profiles and government schemes. This direction is relevant because scheme requirements are often expressed as structured conditions that can be checked against profile attributes.

Jai Akash et al. [5] discuss AI-driven government scheme recommendation based on citizen attributes and scheme information. Thiyagarajan [6] describes welfare assistance that combines citizen information, eligibility guidance and chatbot support. Asif Ahamed [7] examines beneficiary-to-scheme mapping together with grievance-related information, while Rokade [8] discusses the use of NLP and AI for processing citizen information and scheme descriptions. These studies indicate opportunities for improving access, but structured eligibility rules remain necessary when a scheme contains mandatory conditions.

The proposed framework draws on these themes and combines them in a controlled sequence. Unlike a system that ranks all schemes only by profile similarity, the proposed method first applies explicit rules. Unlike a rule-only system, it then ranks qualifying schemes according to profile relevance. Finally, it links the output to the conditions that generated it, supporting transparency and subsequent verification.

V. PROPOSED SYSTEM ARCHITECTURE

The proposed architecture contains a citizen-facing web interface, profile management, a scheme database, an eligibility rule engine, a content-based recommendation component, an explanation component and an administrative module. The citizen provides profile information through the interface. The backend validates and stores the information, retrieves the relevant scheme records and invokes the rule engine. The eligible or potentially eligible records are passed to the ranking component. The explanation component combines the rule outcomes and ranking output to generate a readable result.

A. Technology Stack

The proposed backend uses Java and Spring Boot to implement application services and business logic. MySQL stores citizen profiles, scheme descriptions, eligibility criteria, recommendation records, documents and feedback. The user interface is designed using HTML5, CSS3, JavaScript and Bootstrap. Python and Scikit-learn may be used to develop and test the content-based similarity component. These technologies describe the proposed implementation framework; the final deployment configuration may depend on testing and institutional requirements.

B. Main Processing Flow

The processing sequence is: citizen registration and profile entry; validation of required profile fields; retrieval of scheme records; rule-based evaluation of mandatory conditions; separation of satisfied, unsatisfied and unverified cases; similarity scoring of schemes that pass the eligibility screening; ranking and presentation; and generation of an explanation for each displayed scheme. The citizen can then review the scheme details, documents and application guidance.

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 9.935

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 11, November 2026

|DOI: 10.15680/IJIRSET.2026.1511040|

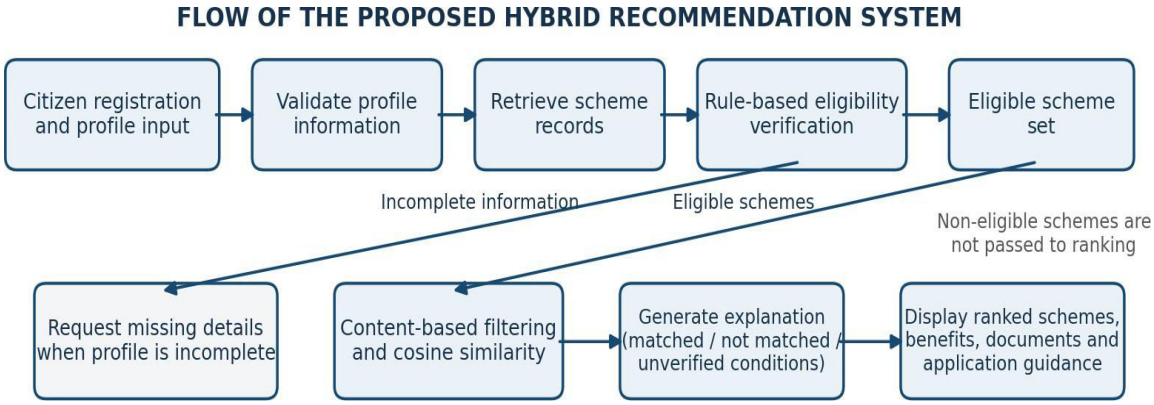

Fig. 1. Flow diagram of the proposed hybrid government welfare scheme recommendation system

C. Design Principle

Eligibility and relevance are treated as different questions. The rule engine answers whether the known profile meets the scheme's defined mandatory conditions. The ranking component answers which of the schemes that pass the screening are more relevant to the citizen's profile. The explanation layer links both outputs to the stored conditions. This separation reduces the risk of presenting a high similarity score as proof of eligibility.

VI. METHODOLOGY

A. Citizen Profile Collection

The system collects profile attributes that are relevant to scheme conditions. These may include age, annual income, gender, social category, occupation, educational qualification, family size, marital status, state, district, employment status, disability status, housing condition and agricultural status. Only attributes required for the intended eligibility checks should be collected. Input validation is required to identify missing, invalid or inconsistent values before processing.

B. Scheme Representation

Each scheme is stored with structured descriptive information, including scheme name, department, purpose, benefits, target group, eligibility conditions, required documents and application guidance. Eligibility conditions are represented as attributes, operators and required values. For example, a condition may specify a minimum age, an income ceiling, a permitted occupation or a residence requirement. Scheme information should be drawn from authoritative sources and reviewed whenever official rules change.

C. Rule-Based Eligibility Engine

For each scheme, the rule engine compares the citizen's profile with its mandatory conditions. A condition is marked satisfied when the available profile value meets the defined requirement, unsatisfied when it conflicts with the requirement, and unverified when the relevant profile value is missing or cannot be evaluated. A scheme with an unsatisfied mandatory condition is excluded from the eligible recommendation ranking. A scheme with unverified mandatory information should be presented, if at all, as requiring further verification rather than as confirmed eligible.
This approach is designed to avoid treating missing information as a positive match. It also allows the interface to explain why a scheme was excluded or why more information is needed. The precise rule logic must be defined for each scheme and reviewed against the latest official eligibility guidelines.

D. Content-Based Similarity Ranking

After eligibility screening, content-based filtering ranks the remaining schemes by comparing the citizen profile with the structured features of each scheme. Profile and scheme attributes are transformed into compatible feature vectors. Categorical attributes can be represented using suitable encoded values, while numeric attributes may be normalized where required. Cosine similarity can then be used to estimate the directional similarity between the vectors:

$$\text{Cosine Similarity}(A, B) = (A \cdot B) / (\|A\| \times \|B\|)$$

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 9.935

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 11, November 2026

|DOI: 10.15680/IJIRSET.2026.1511040|

Here, A represents the citizen feature vector and B represents the scheme feature vector. A higher similarity score indicates greater profile-to-scheme relevance within the chosen representation. The score is used for ordering only; it does not override mandatory rules, determine official acceptance or prove that the citizen will receive a benefit. Feature selection, encoding and weighting must be tested to ensure that the score is meaningful for the welfare domain.

E. Explanation Generation

For every recommendation, the explanation component presents the relevant profile attributes and the corresponding scheme conditions. The output can include satisfied conditions, unsatisfied conditions and unverified requirements. A "Why This Matches" section may explain, for example, that the citizen's recorded residence and occupation match the scheme's stated target group. Where income or another required value is missing, the system should explicitly request verification. Explanations are generated from rule outcomes and stored scheme data, rather than from an unsupported similarity score alone.

F. Recommendation Output

The final result displays ranked schemes with their purpose, benefits, eligibility summary, required documents and application instructions. Users may compare schemes and review the conditions behind each result. The interface should distinguish clearly between a scheme that passed the available rule checks and a scheme for which some requirements remain unverified. The system should also display the date on which scheme information was last reviewed, where that metadata is available.

VII. SYSTEM MODULES AND DATA DESIGN

The application is divided into functional modules so that profile management, scheme maintenance, eligibility checking and recommendation generation can be developed and tested independently.

Module	Purpose
User registration and profile	Create and maintain the citizen profile used for eligibility analysis.
Scheme listing and details	Display scheme purpose, benefits, eligibility, documents and application guidance.
AI eligibility checking	Apply structured rules to the citizen profile and record each condition outcome.
Benefit scoring and ranking	Prioritize schemes that pass screening using content-based similarity.
Explainable AI	Show matched, unmatched and unverified conditions behind each result.
Scheme comparison	Allow users to review relevant alternatives side by side.
Benefit wallet	Track availed, pending and potentially missed benefits, as proposed in the system design.
Family benefit dashboard	Consolidate welfare opportunities for family members, where supported by profile data.
Age-based benefit timeline	Highlight schemes that may become relevant at different life stages.
Feedback and grievance support	Collect feedback and service-related issues.
Admin module	Manage scheme records, eligibility rules, users, recommendations and grievances.

A. Database Design

The relational database stores citizens, welfare schemes, eligibility criteria, recommendations, documents and feedback. Primary and foreign keys maintain the relationships between records. The scheme table is linked to its eligibility criteria and documents; recommendation records connect a citizen with a scheme and store the match score, rank, status and explanation.

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 9.935

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 11, November 2026

|DOI: 10.15680/IJIRSET.2026.1511040|

Table	Important columns	Key
Citizen	citizen_id, name, age, income, occupation, location	PK: citizen_id
Welfare_Scheme	scheme_id, scheme_name, department, benefits, application_url	PK: scheme_id
Eligibility_Criteria	criteria_id, scheme_id, attribute, operator, required_value	PK: criteria_id; FK: scheme_id
Recommendation	recommendation_id, citizen_id, scheme_id, match_score, rank, status, explanation	PK: recommendation_id; FK: citizen_id, scheme_id
Document	document_id, scheme_id, document_name, description	PK: document_id; FK: scheme_id
Feedback	feedback_id, citizen_id, scheme_id, rating, comments	PK: feedback_id; FK: citizen_id, scheme_id

VIII. ALGORITHM AND WORKFLOW

The overall procedure can be expressed as the following logical steps:
Step 1: Receive the citizen profile and validate required fields.
Step 2: Retrieve active scheme records and their latest available eligibility criteria.
Step 3: For each scheme, evaluate every mandatory condition against the profile.
Step 4: Record each condition as satisfied, unsatisfied or unverified.
Step 5: Exclude schemes that fail a mandatory condition from the eligible ranking; flag records with missing mandatory information for verification.
Step 6: Convert the citizen profile and qualifying scheme attributes into feature vectors.
Step 7: Calculate cosine similarity and sort qualifying schemes by score.
Step 8: Generate an explanation using the rule outcomes and relevant profile attributes.
Step 9: Display ranked scheme information, documents and application guidance.
Step 10: Store recommendation and feedback records for audit and system improvement.

The rule engine is the controlling stage for mandatory eligibility. The ranking algorithm operates only after this stage. This sequence is essential: the system must not allow a high similarity score to reintroduce a scheme that failed a mandatory condition. The system should also retain the rule version and scheme-data update information used for each recommendation so that results can be reviewed later.

IX. EXPECTED RESULTS AND DISCUSSION

The proposed system is expected to provide a consolidated and more understandable way to discover government welfare schemes. Instead of requiring citizens to search through multiple sources independently, the application uses a profile to identify relevant scheme records and organizes the results in a ranked list. The rule-based stage is expected to reduce recommendations that violate known mandatory conditions, while content-based ranking is intended to make the order of qualifying schemes more relevant to the citizen.

The explanation component is expected to improve transparency by showing which conditions matched and which require attention. It may also help citizens prepare the required documents and understand why a scheme was not included in the main recommendation list. The modular design supports future extensions such as multilingual interfaces, natural-language queries, richer behavioural information and more extensive scheme datasets.

No measured accuracy, precision, recall, response-time or user-satisfaction results are claimed in this paper because the described framework requires implementation and controlled evaluation. A future evaluation should use verified scheme rules and representative test profiles. Rule-engine evaluation should measure whether each mandatory condition is classified correctly. Ranking evaluation may use expert relevance judgments and suitable ranking metrics.

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 9.935

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 11, November 2026

|DOI: 10.15680/IJIRSET.2026.1511040|

Usability testing should assess whether citizens understand the recommendation status and explanations. Dataset freshness and the frequency of rule updates should also be monitored.

A. Suggested Evaluation Measures

Eligibility correctness can be measured by comparing system decisions with manually verified outcomes for a labelled set of citizen-scheme pairs. Ranking quality can be assessed using precision at K, recall at K or normalized discounted cumulative gain where relevance labels are available. Explanation quality can be assessed through expert review of whether each displayed reason is traceable to a stored rule. System performance can be evaluated through response time, database query time and behaviour under increasing profile and scheme volumes. These are proposed evaluation measures, not reported experimental findings.

X. LIMITATIONS AND FUTURE SCOPE

The quality of the recommendations depends on the completeness and correctness of citizen profiles, the accuracy of eligibility rules and the coverage of the scheme database. Government schemes may change their income limits, age requirements, documents or application procedures. Outdated information can lead to misleading recommendations even when the software logic is correct. A process for verifying official sources and updating records is therefore essential.

Incomplete profiles create uncertainty. The system must not infer sensitive or missing values and should clearly request additional information when necessary. The similarity score reflects relevance under the selected feature representation; it does not establish legal eligibility or guarantee approval. Behaviour-based recommendation may be difficult when interaction data is limited, and the system should avoid allowing popularity or historical usage to override mandatory conditions.

Future work may expand the scheme dataset, support multiple Indian languages, introduce natural-language processing for citizen queries, improve the explanation interface and evaluate ranking algorithms such as Logistic Regression, Decision Tree, Random Forest and XGBoost where appropriate training labels are available. Further work may also include secure data handling, consent-based profile management, accessibility testing, administrative review workflows and integration with official information sources. Any such extension should preserve the distinction between recommendation support and official benefit approval.

XI. IMPLEMENTATION CONSIDERATIONS AND SYSTEM ANALYSIS

A. Functional Requirements

The functional requirements describe the operations that the application is expected to perform. A citizen should be able to register, sign in, create or update a profile, and provide the attributes required for scheme matching. The system should validate entered values and indicate which required information is missing. It should retrieve the current scheme records, evaluate eligibility conditions, rank qualifying schemes, and display the reasons for each recommendation. The citizen should be able to open a scheme detail page, review its benefits and documents, and follow the available application guidance. Where a scheme has not passed the screening, the interface should communicate the reason in a clear and respectful manner.

The administrative functions include adding and updating scheme records, defining eligibility conditions, maintaining document lists, activating or deactivating schemes, and reviewing feedback or grievance information. Changes to a scheme should be traceable. An administrator should be able to identify the source and effective date of a rule, because an eligibility condition may be revised by the issuing department. A recommendation should therefore be associated with the version of scheme data used during its generation.

B. Non-Functional Requirements

Usability is important because the intended users may have different levels of digital literacy. Form labels should use plain language, avoid unnecessary technical terms and provide examples where an attribute is difficult to interpret. The application should support responsive layouts so that the same functions remain accessible on desktop and mobile devices. Accessibility considerations include readable contrast, keyboard navigation, meaningful labels, error messages and a logical sequence through profile forms.

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 9.935

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 11, November 2026

|DOI: 10.15680/IJIRSET.2026.1511040|

Reliability requires consistent rule evaluation for the same profile and scheme-data version. The system should handle invalid inputs and missing information without silently generating a positive eligibility result. Performance requires that profile validation, scheme retrieval and ranking complete within a practical response time as the number of schemes and users increases. Maintainability requires separation of the user interface, business logic, recommendation functions and data access layer so that updates to one component do not unnecessarily affect the others.

Security and privacy are essential because the system processes personal and socioeconomic information. The application should collect only information necessary for the stated recommendation purpose, explain why it is requested and protect it against unauthorized access. Authentication, role-based access, secure communication, input validation, database access controls, backups and audit logs should be considered during implementation. Sensitive attributes should not be exposed in administrative screens unless access is required for an authorized task. Data retention and deletion procedures should be defined before deployment.

C. Use-Case Description

The principal actor is the citizen. A typical use case begins when a citizen creates an account and enters profile details. After validation, the citizen requests scheme recommendations. The application invokes the eligibility engine and then the ranking component. The result page displays qualifying schemes, their similarity-based order and the explanations. The citizen may select a scheme to read its full description, required documents and application instructions. The citizen may also update a profile value and request a fresh evaluation when circumstances change.

The administrator is responsible for maintaining the scheme catalogue. An administrator adds a scheme, records its source, enters the eligibility conditions, associates required documents and verifies the information before publishing it. When a government department revises a scheme, the administrator updates the record and retains the effective date or version information. The administrator may also review user feedback to identify unclear descriptions or incorrect records. These responsibilities are separate from official application processing and benefit approval.

D. Input Validation and Data Quality

The quality of the recommendation depends on the quality of input data. Profile fields should have validation rules appropriate to their data types. Age should be numeric and within a reasonable range; income should use a clearly defined period and currency; location fields should use a consistent administrative structure; and categorical attributes should use controlled values where possible. The system should explain validation errors and allow users to correct them before the recommendation process begins.

Scheme records require comparable care. Eligibility conditions should be stored in a structured form rather than only as free-text descriptions. Each condition should identify the profile attribute, comparison operator, required value, whether the condition is mandatory and the source from which it was obtained. Conditions that cannot be represented reliably as simple rules should be marked for manual review or additional guidance. The system should not convert ambiguous text into a definitive eligibility rule without validation by an authorized reviewer.

Data normalization is necessary when citizen and scheme records use different formats. For example, occupation names, district names and social-category labels may be written in different ways. A controlled vocabulary or mapping table can reduce mismatches. Numeric values may require unit consistency, and location matching may require state and district codes. The preprocessing stage should preserve the original user entry where appropriate while storing a normalized value for matching. The mapping itself should be documented so that administrators can correct errors.

E. Handling Missing and Conflicting Information

A citizen may not know or may choose not to provide every profile attribute. The system should distinguish between a value that is absent, a value that is not applicable and a value that has been verified. If a mandatory condition depends on an absent value, the outcome should be unverified. If two profile fields conflict, the system should request clarification rather than choosing one value without explanation. This is particularly important when a scheme uses household income, family composition or location as a decisive condition.

The user interface can group follow-up questions according to the schemes being considered. It should request only the additional information needed to evaluate relevant conditions and explain why the information is required. Where a user cannot provide a value, the system should retain the uncertainty and direct the user to the official scheme guidance

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 9.935

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 11, November 2026

|DOI: 10.15680/IJIRSET.2026.1511040|

or responsible department. This process reduces the risk of false certainty and supports informed decisions.

F. Recommendation Traceability

Each recommendation record should preserve enough information to reproduce or review the result. Useful fields include the citizen identifier, scheme identifier, rule-data version, evaluation date, condition outcomes, similarity score, rank and explanation text. Storing the rule outcomes is important because the scheme criteria may change after the recommendation is generated. A traceable record allows an administrator or authorized reviewer to determine whether the output reflected the information available at that time.

The explanation should be generated from the same rule outcomes used in the screening stage. It should not state that a citizen is officially approved, guaranteed to receive a payment or definitively eligible unless the responsible authority has confirmed that status. Preferred wording includes "matches the information provided," "appears relevant," or "requires verification." The distinction between a system recommendation and an official decision should remain visible throughout the interface.

XII. DETAILED MODULE DESCRIPTION

A. Registration and Profile Management

The registration module provides account creation and authentication. After signing in, a citizen can enter profile information in a structured form. Profile management should allow the user to correct or update details such as age, income, occupation, education, family size and location. The application should provide short descriptions of fields whose meaning may not be obvious. Profile updates should trigger a new evaluation when the citizen requests recommendations, because a change in income, occupation or residence may alter the set of matching schemes.

B. Scheme Catalogue and Search

The scheme catalogue is the central source of scheme descriptions used by the recommendation engine. Each entry should include a scheme name, administering department, purpose, benefit description, target group, eligibility summary, detailed rules, required documents and official application reference where available. Search and filter functions may allow users to browse by sector, state, beneficiary group or keyword. The recommendation engine uses the structured fields, while the detail page presents the full explanatory information to the citizen.

C. Eligibility Service

The eligibility service receives a citizen profile and a set of scheme conditions. It evaluates each mandatory rule and returns a structured result. A rule outcome should include the condition identifier, the profile attribute used, the comparison performed and the result status. This structured response can be consumed by the ranking and explanation components. Keeping the service independent makes it easier to test rule logic and update scheme conditions without rewriting the user interface.

A. Ranking Service

The ranking service accepts only the scheme records permitted by the eligibility policy. It prepares feature representations for the citizen and each candidate scheme, applies the selected similarity measure and sorts the candidates. If two schemes have similar scores, the system may use additional transparent tie-breaking rules, such as a defined priority or the completeness of information, provided that these rules are documented. The system should retain the score for audit and should avoid displaying a score as a probability of approval.

D. Explanation and Scheme Comparison

The explanation service converts the rule outcomes into user-readable statements. It can group conditions into "matched," "not matched" and "needs more information." A comparison view may show the benefits, target group, key eligibility conditions and required documents for two or more schemes. The purpose is to help citizens understand alternatives without hiding important differences. The text should remain consistent with the official scheme description and should link to the relevant source where possible.

E. Benefit Wallet and Family Views

The proposed benefit wallet is intended to help a citizen track schemes that have been availed, are pending or may have been missed. Such tracking requires user-entered status or reliable application information; the system should not

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 9.935

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 11, November 2026

|DOI: 10.15680/IJIRSET.2026.1511040|

assume that a recommendation means an application was submitted or approved. A family dashboard can organize potential schemes for different family members where the user has supplied the necessary profile details and has permission to manage that information. Family-level aggregation should preserve privacy and access controls.

F. Age-Based Benefit Timeline

An age-based timeline can present schemes that may become relevant at different life stages, such as education support for students or social-security programmes for older citizens. The timeline is a discovery aid and should not be interpreted as an automatic determination of future eligibility. The system must evaluate the scheme conditions using the current profile and the applicable scheme version whenever a citizen requests a recommendation.

G. Feedback and Administrative Review

Feedback can help identify confusing explanations, missing scheme information or errors in the catalogue. A grievance-related feature may record a citizen's concern and route it to an appropriate administrative process if such integration is available. Feedback should not directly alter eligibility rules or recommendation rankings without review. Administrators should verify reported issues against official sources before changing published scheme data.

XIII. TESTING AND EVALUATION PLAN

A. Unit and Integration Testing

Unit tests should cover individual rule operators, boundary conditions, missing values, invalid data types and normalization functions. For example, an age rule should be tested with values below, equal to and above the threshold. An income ceiling should be tested at the boundary and with values on both sides of it. Tests should also verify that an absent value returns an unverified status rather than a satisfied status. Integration tests should confirm that the profile service, scheme database, rule engine, ranking service and explanation component exchange consistent data.

B. Functional Scenario Testing

Representative scenarios should include a student profile, a farmer profile, a salaried worker profile and a household profile with multiple dependants. These are test personas, not claims about real beneficiaries. For each scenario, expected rule outcomes should be prepared by reviewing the scheme conditions. The application output can then be compared with the expected result. Additional scenarios should cover incomplete profiles, conflicting location information, changed income and an updated scheme rule.

C. Ranking Evaluation

Ranking evaluation requires a relevance reference set. Domain reviewers can assess whether each scheme is relevant to a profile after the mandatory conditions have been considered. Precision at K measures the proportion of relevant schemes among the top K results, while recall at K measures how many of the relevant schemes appear in those results. Normalized discounted cumulative gain can account for graded relevance and the position of each scheme in the ranked list. The evaluation should report the size and composition of the test set and the method used to obtain relevance judgments.

D. Explainability and Usability Evaluation

Explanation evaluation should verify that every displayed statement corresponds to an actual stored condition and that no unsupported claim is introduced. Reviewers can check whether explanations correctly distinguish satisfied, unsatisfied and unverified criteria. Usability evaluation may ask participants to complete tasks such as entering a profile, locating a scheme, understanding why it was recommended and identifying the documents required. Feedback can reveal unclear field labels, confusing status messages or barriers for users with limited digital experience.

E. Performance and Maintenance Evaluation

Performance tests should measure the time required for profile validation, scheme retrieval, eligibility evaluation and ranking as the number of schemes and users grows. Database tests should examine indexing and query efficiency for scheme attributes and eligibility conditions. Maintenance evaluation should simulate a scheme update and verify that the new version is used for subsequent recommendations while earlier recommendation records remain traceable. These tests provide evidence about operational readiness; they should be completed before making claims about real-world effectiveness.

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 9.935

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 11, November 2026

|DOI: 10.15680/IJIRSET.2026.1511040|

XIV. ETHICAL, PRIVACY AND GOVERNANCE CONSIDERATIONS

A welfare recommendation system can influence how citizens understand public benefits. Incorrect information may discourage an eligible person from applying or create an expectation that a benefit will be granted. For this reason, the system should communicate uncertainty, identify official sources and avoid presenting recommendations as government decisions. Citizens should be able to inspect the information they supplied and correct errors. Where a rule is unclear or a profile is incomplete, the system should recommend verification rather than making an unsupported conclusion.

Personal data should be processed for a clearly stated purpose and protected through appropriate technical and organizational safeguards. Access to citizen profiles should be limited to authorized functions. Administrative users should have only the permissions necessary for their duties, and changes to scheme rules should be logged. The system should avoid collecting attributes that are not needed for the eligibility or relevance functions. Any use of behavioural information should be transparent and should not disadvantage citizens who have little or no prior interaction history.
Fairness requires attention to data coverage and rule representation. If the scheme catalogue omits programmes for a particular region or beneficiary group, the recommendations will also be incomplete. If profile categories are inconsistent, citizens may be incorrectly screened. Regular review of the dataset, test cases across different profile types and a channel for reporting errors are therefore necessary. The system should also make clear that citizens can consult the responsible department or official portal for authoritative guidance.

XV. DISCUSSION

The principal contribution of the proposed design is the separation of eligibility verification from relevance ranking. In a conventional similarity-only recommender, a scheme may appear near the top because its descriptive features resemble a citizen profile even though a mandatory condition is not satisfied. In the proposed process, the rule engine controls which schemes can enter the eligible ranking. This is suitable for a domain in which recommendations must respect explicit conditions rather than merely maximize user preference similarity.

The approach also has limitations. Rule-based screening depends on the ability to translate official scheme conditions into structured and current rules. Some schemes may contain exceptions, discretionary provisions or documentation requirements that cannot be fully represented by a simple comparison. The system must therefore support unverified outcomes and human review. Content-based ranking also depends on the selected attributes, encoding method and weights. A high similarity score may reflect feature overlap without capturing every aspect of a citizen's actual needs.
A further design consideration is the balance between personalization and simplicity. Collecting many profile attributes may improve matching for some schemes, but it increases the burden on users and the amount of personal data processed. The system should collect only the information required for the schemes under consideration and request additional details progressively. Clear explanations and a transparent data-update process are as important as the recommendation algorithm itself.

A. Practical Adoption

For practical adoption, the scheme catalogue should be introduced in stages. A pilot can begin with a limited set of schemes from selected sectors and a clearly defined geographic scope. Each scheme record should be checked by a reviewer against the official notification or portal before it is made available to citizens. The pilot should record how often users encounter missing information, how many recommendations require verification and which explanations are difficult to understand. This evidence can guide improvements to the profile form, rule representation and administrative workflow before the catalogue is expanded.

The project should also define responsibility for maintaining the data. Software administrators can manage the technical platform, but scheme content should be verified by personnel familiar with the relevant official sources. A review schedule can identify records that have not been checked recently. Where an official source changes or becomes unavailable, the scheme should be flagged for review rather than silently retained as current. This operational process is necessary because recommendation logic is only as reliable as the information on which it depends.

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 9.935

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 11, November 2026

|DOI: 10.15680/IJIRSET.2026.1511040|

B. Expected Benefits for Citizens and Administrators

For citizens, a structured recommendation interface can reduce the effort required to discover schemes and compare their conditions. A clear explanation can help users understand why a programme appears relevant and which details they may need to confirm. The benefit wallet and family views may provide an organized way to keep track of opportunities, subject to accurate user-entered status information. For administrators, structured records and feedback may help identify incomplete descriptions, repeated questions and scheme information that requires updating. These benefits are expected outcomes of the design and should be verified through field evaluation rather than assumed in advance.

The system is not intended to replace government portals, application procedures or officials who make eligibility decisions. Instead, it acts as an information and discovery layer that can direct citizens toward relevant official scheme details. A successful implementation must make this role explicit, provide links to authoritative sources where available and avoid collecting more personal information than is needed. These safeguards support responsible use and help maintain trust in digital public-service tools.

XVI. CONCLUSION

This paper presents the design of an AI-based personalized government welfare benefits recommendation system intended to improve access to relevant public welfare information. The proposed framework combines a rule-based eligibility engine with content-based similarity ranking and an explanation component. The rule engine evaluates mandatory conditions using structured citizen profiles, while the ranking component orders schemes that pass the screening according to their relevance to the profile. Explanations identify satisfied, unsatisfied and unverified requirements so that users can understand the basis of the results.

The proposed architecture uses Java and Spring Boot for backend services, MySQL for data management, HTML, CSS, JavaScript and Bootstrap for the interface, and Python with Scikit-learn as a possible environment for recommendation development and testing. The design can support scheme comparison, a benefit wallet, family-level views and an age-based benefit timeline. Its practical effectiveness will depend on verified scheme data, responsible handling of citizen information, clear explanations and systematic evaluation. The framework offers a structured direction for improving scheme discovery, while final eligibility and benefit approval remain with the responsible government authorities.

REFERENCES

1) “A Comprehensive Review of Recommender Systems: Transitioning from Theory to Practice,” 2024.
2) “A Survey of Recommender System Techniques and the E-Commerce Domain,” 2022.
3) R. Mandalageri et al., “AI-Powered Hybrid Recommender for Adaptive Government Service Allocation Using ML and Behavioural Filtering,” International Journal of Engineering Sciences (IJES), 2024
4) “Bridging Government Schemes with Citizens for Inclusive Governance,” International Journal of Innovative Research in Science and Engineering Technology (IJIRSET), 2025
5) “AI Driven Government Scheme Recommendation System,” Journal of Advances and Applications in Fluid Research (JAAFR), 2025
6) “AI-Driven Welfare Scheme Assistance and Management System,” International Journal of Research Publication and Reviews (IJRPR), 2025.
7) “Smart Beneficiary Schemes Mapping and Grievance Redressed System,” International Journal of Creative Research Thoughts (IJCRT), 2025
8) “Integrating NLP, AI, and Citizen Data for Smart Welfare Recommendation (AI-GSRS),” International Journal of Scientific Research and Engineering Development (IJSRED), 2025
9) Government of India, myScheme and UMANG digital service platforms, official portals.



INTERNATIONAL STANDARD SERIAL NUMBER INDIA

INNO SPACE SJIF Scientific Journal Impact Factor

doi crossref

निस्केयर NISCAIR

INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN SCIENCE | ENGINEERING | TECHNOLOGY

9940 572 462 6381 907 438 ijirset@gmail.com

www.ijirset.com

Scan to save the contact details