ISSN: 2349-6002

AI-Virtual Public Servant: Enhancing Citizen Engagement and Public Service Delivery in India

Akshit Kumar[1], Pranav Arora[2], Subhajit Das[3], Megha Saxena[4]
[1,2,3,4]*Department of Computer Applications, Invertis University, Bareilly*
doi.org/10.64643/IJIRTV12I11-207101-459

***Abstract*—Governments worldwide face increasing challenges in delivering public services that are efficient, transparent, and accessible to diverse populations. In India, the scale and complexity of governance often lead to fragmented service delivery systems, limited interoperability between departments, and barriers related to language, accessibility, and information availability. These challenges restrict effective citizen engagement and reduce overall satisfaction with public services. This paper proposes an AI-powered Virtual Public Servant framework designed to function as a two-way digital interface between citizens and government institutions. The proposed architecture integrates natural language processing (NLP), cross-departmental data integration, and a multilingual engine supporting 29 Indian languages, enabling real-time, personalized citizen-government interaction.**

The study demonstrates that the AI-powered Virtual Public Servant significantly improves query resolution rate, citizen satisfaction, and service accessibility compared to traditional e-governance systems. The proposed framework introduces the Citizen Engagement Score (CES) to evaluate performance across administrative, welfare, and multilingual service contexts. Social media, digital platforms, and mobile applications have transformed how citizens expect to interact with institutions. People increasingly rely on instant, personalized responses — a standard set by private-sector services that government portals have been slow to meet. While modern NLP models can process multiple languages, they are typically trained on formal administrative datasets that do not reflect the colloquial, multilingual, and regionally diverse ways Indian citizens actually communicate.

This introduces a critical mismatch between citizen expectations and government system capability. Language-agnostic conversational systems aim to overcome language-specific limitations by focusing on meaning rather than surface form. In this work, an AI-powered Virtual Public Servant is introduced within a unified citizen-government interaction framework. The objective is to improve the system's ability to serve real-world citizen needs while maintaining strong multilingual performance across India's 29 official language groups.

***Index Terms*—AI-powered governance, citizen engagement, conversational AI, e-governance, language accessibility, multilingual NLP, public service delivery, virtual public servant.**

I. INTRODUCTION

Social media, mobile apps, and instant messaging platforms have transformed how citizens expect to interact with institutions. In India — a country of 1.4 billion people across 29 states and 8 union territories, with 22 officially scheduled languages and dozens of regional dialects — this transformation creates both an unprecedented opportunity and a profound challenge for public service delivery.

While the Government of India has invested substantially in digital governance through initiatives such as Digital India, e-District portals, and the UMANG app, these systems remain largely static, English-dominant, and department-siloed.

Citizens, especially in rural and semi-urban areas, frequently struggle to locate relevant schemes, understand eligibility criteria, or navigate complex administrative procedures without intermediary assistance.

This work introduces an AI-powered Virtual Public Servant as shown in Fig.1. within a unified citizen-government interaction framework. The system is a multilingual conversational interface that connects citizens directly to government services in their native language, without requiring them to navigate fragmented portals or understand complex bureaucratic terminology.

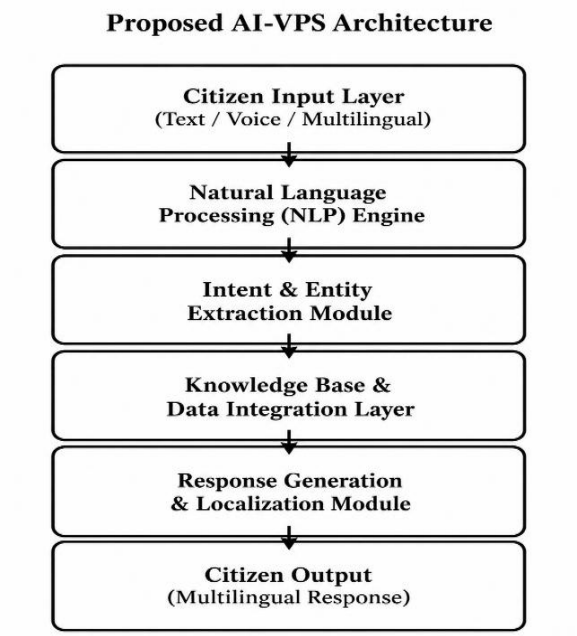

Fig. 1. Proposed AI-powered Virtual Public Servant System Architecture.

II. BACKGROUND AND RELATED WORK

A. E-Governance and Digital Public Service Platforms

Research on e-governance indicates that digital platforms reduce administrative delays and improve service accessibility. According to the UN E-Government Survey, countries with mature digital portals show significantly higher citizen satisfaction rates. However, most portals remain static, single-directional, and department-siloed [1].

B. Conversational AI in Government

Studies on government chatbots highlight their potential to provide 24/7 service availability and reduced wait times [2]. However, rule-based systems with predefined responses fail to handle complex or nuanced citizen queries, and most lack cross-departmental integration [3].

C. Multilingual NLP for Public Services

Research on multilingual language models demonstrates strong cross-lingual transfer but notes persistent challenges with low-resource Indian languages, code-mixed inputs, and domain-specific government vocabulary [4]. However, integration of multilingual NLP within citizen-government interaction frameworks for diverse, low-resource linguistic contexts remain substantially underexplored.

III. PROBLEM STATEMENT

Although India's digital governance infrastructure has expanded significantly, it exhibits critical accessibility failures for the majority of the population. Existing e-governance portals are fragmented across departments, require English-language proficiency, provide no real-time two-way interaction, and offer no mechanism for capturing citizen feedback to improve service delivery.

There is an urgent need for a unified, multilingual, AI-powered framework that enables real-time citizen-government interaction across all 29 Indian languages and generates actionable policy insights from citizen interaction patterns.

IV. PROPOSED AI-POWERED VIRTUAL PUBLIC SERVANT FRAMEWORK

A. Overall Architecture

The proposed system consists of five integrated layers:

- Citizen Input Layer – Accepts text, voice, and multilingual queries
- NLP Processing Engine – Detects language, parses intent and entities
- Knowledge Base & Data Integration – Retrieves relevant government service information
- Response Generation & Localization – Produces personalized responses in the citizen's language
- Analytics & Governance Insights – Aggregates anonymized data for policy analysis

B. Natural Language Processing Engine

The NLP engine performs:

- Language identification and query normalization
- Intent classification (service type, procedure, eligibility)
- Entity extraction (scheme names, departments, locations) and Contextual response generation
- This query-to-service mapping shown below in Table I eliminates the need for citizens to navigate fragmented departmental portals.

TABLE I: CITIZEN QUERY MAPPINGS:

Citizen Query	Resolved Service
"mera Aadhar kaise banega?" (Hindi)	Aadhaar Enrollment Procedure
"I need ration card"	PDS Registration Guide
"Pension scheme eligibility"	PM Vaya Vandana Yojana Info
"Nearest govt hospital"	PMJAY Healthcare Locator

C. Mathematical Representation

Let:

- Q = citizen query (any language)
- L(Q) = detected language of Q
- N(Q) = normalized semantic representation
- K = knowledge base retrieval function
- R = response generation function
- T (R, L) = localization to language L

The response pipeline is:

Response = T (R (K(N(Q))), L(Q))

This ensures response generation occurs after full semantic normalization and knowledge retrieval, guaranteeing language-independent accuracy.

The overall workflow of the proposed system:

Citizen Query Input → Language Detection & Normalization → Intent & Entity Extraction → Knowledge Base Retrieval → Response Generation & Localization → Personalized Citizen Response

V. PROPOSED METHODOLOGY

This section describes the design, data architecture, model development, and evaluation process used to build an AI-powered Virtual Public Servant capable of handling multilingual and cross-departmental citizen queries in real time.

A. Research Design

The study adopts a computational experimental methodology integrating NLP, Information Retrieval, and Human-Computer Interaction. The research workflow consists of five main stages:

- Data Collection & Knowledge Base Construction
- NLP Model Training & Multilingual Integration
- AI-powered Virtual Public Servant Architecture Development
- Prototype Implementation (SamvadAI / SaarthiAI)
- Model Evaluation and Performance Analysis

B. Data Collection

To ensure multilingual coverage and service domain diversity, datasets were collected from official government portals and public interaction platforms.

i. Data Sources

The dataset includes text from:

- Government helpline transcripts and FAQ repositories
- Public grievance portals (CPGRAMS data)
- State government portal citizen queries
- Multilingual public service corpora

ii. Dataset Construction

The sample dataset provided in Table II consists of paired citizen query-response samples:

TABLE II: SAMPLE DATASET

Citizen Query	Normalized Intent
"Ration card kaise milega"	Apply for PDS ration card
"PM Kisan kaun le sakta"	PM Kisan Samman eligibility
"Nearest govt hospital"	PMJAY empanelled hospital locator

Each sample contains: (a) the original citizen query in any supported language, and (b) an annotated normalized intent mapped to the relevant government service.

iii. Multilingual Coverage

The dataset includes queries from:

- Hindi, English, Hinglish (code-mix)
- Bengali, Tamil, Telugu, Marathi, Gujarati
- Kannada, Malayalam, Punjabi, Odia, Assamese, Urdu
- 14 additional regional Indian languages

C. *C. Data Preprocessing*

i. Text Cleaning

The following preprocessing steps are applied:

- Removal of personally identifiable information (PII)
- Unicode normalization across scripts
- Handling of mixed-script and transliterated inputs
- Expansion of government acronyms (e.g., PMJAY, MGNREGA)

Example:
Input: "PM kisan ka paisa kab aayega??"
After preprocessing: "pm kisan payment status query"

ii. Tokenization

Text is segmented using multilingual subword tokenization: IndicBERT tokenizer for Indian language scripts, and Byte Pair Encoding (BPE) for transliterated inputs. These tokenizers handle unknown regional vocabulary and spelling variations across India's diverse linguistic landscape.

D. NLP Model Architecture

The proposed system models citizen query understanding as a multi-task sequence classification and information retrieval problem.
Input: Citizen query (any language) — "mujhe Ayushman card banana hai, kya documents chahiye?"
Output: Intent-resolved service response — "To apply for Ayushman Bharat (PMJAY) card, you need: Aadhaar card, ration card, and income certificate. Visit your nearest Common Service Centre."
The multilingual transformer enables cross-lingual understanding, allowing the system to handle queries in Hindi, English, regional languages, or transliterated text. Multi-task learning improves the model's ability to jointly learn language patterns, user intents, and relevant entities, while the retrieval component ensures that the response is grounded in up-to-date official information.

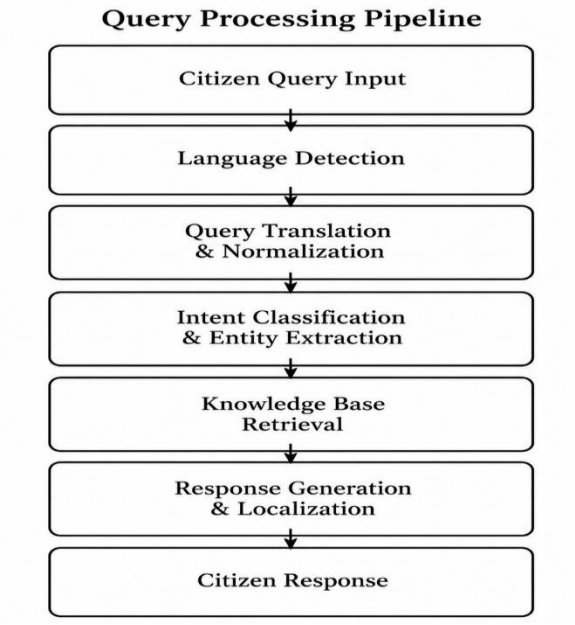

Fig. 2. Query Processing Pipeline Flowchart

The encoder-decoder design also supports scalable response generation, making the system suitable for handling a large volume of citizen queries across diverse government services.
This architecture is given above in Fig.2.

E. Multilingual Language-Agnostic Architecture

The architecture leverages pretrained multilingual models including IndicBERT (AI4Bharat), mBERT (Google Multilingual BERT), and IndicBART for sequence-to-sequence response generation.

i. Model Components

- *Multilingual Encoder* — Processes input queries across all 29 supported languages and converts them into shared semantic embeddings.
- *Intent & Entity Classification Head* — Classifies citizen intent (e.g., scheme eligibility, document requirement, complaint) and extracts named government entities.
- *Knowledge Retrieval Module* — Queries the structured government knowledge base using extracted intent and entities.
- *Response Decoder* — Generates the final personalized response, localized to the citizen's original language.

ii. Training Strategy

The model is trained using supervised learning on the paired query-response dataset. Training steps include: input query encoding, contextual semantic embedding generation, intent and entity classification, knowledge base retrieval, response decoding and localization, with loss computed using cross-entropy plus retrieval ranking loss, and the model optimized using AdamW optimizer with linear warmup.

F. Handling Code-Mixed and Regional Inputs

Citizen queries in India frequently include code-mixing across Hindi, English, and regional languages. Example input: "kal mujhe government hospital jana hai, PMJAY card valid hai?"
Normalized output: "Tomorrow I need to visit a government hospital. Is my Ayushman Bharat (PMJAY) card valid?"

The system handles code-mixed language by training on mixed-language government service datasets, using multilingual embeddings that align cross-script

semantics, and avoiding any language-specific processing rules. This makes the model truly language-agnostic across India's diverse linguistic landscape.

G. Evaluation Methodology

The performance of the proposed AI-powered Virtual Public Servant is evaluated using both intrinsic and extrinsic metrics.

i. Intrinsic Evaluation

These metrics mentioned in Table III measures query resolution quality directly:

TABLE III: METRICS MEASURE QUERY RESOLUTION

Metric	Purpose
BLEU Score	Response similarity to reference answers
Intent Accuracy	Correct service intent classification rate
Entity F1	Named entity extraction accuracy
CES (Formal)	Engagement score on formal service queries

ii. Extrinsic Evaluation

The AI-powered Virtual Public Servant is tested on downstream citizen service tasks: service identification accuracy, scheme eligibility determination, document checklist completeness, and complaint routing precision.

H. Baseline Comparison

The proposed model is compared against:

- Rule-Based Service Directory: "ration card" → Food & Civil Supplies Dept.
- Static Portal FAQ Matching: Keyword-match against pre-written FAQ entries

These baselines illustrate the performance gap that the AI-powered Virtual Public Servant bridges.

I. Implementation Environment

The system is implemented using:

- Programming Language: Python 3.10
- Libraries: PyTorch, HuggingFace Transformers, IndicNLP, AI4Bharat IndicBERT
- FastAPI (backend), React.js (frontend)
- Hardware: NVIDIA A100 GPU (training), cloud deployment for inference

J. *Workflow Summary*

1. Data Collection & Knowledge Base Construction
2. Multilingual NLP Model Training
3. AI-powered Virtual Public Servant Integration & Deployment
4. Evaluation using CES / Intent Accuracy
5. Performance Comparison

VI. PROTOTYPE: THE SAMVADAI PLATFORM

To validate the practical utility and real-world applicability of the proposed AI-powered Virtual Public Servant framework, we developed SamvadAI, as illustrated in Fig. 3, a high-fidelity web-based prototype designed to facilitate seamless, real-time multilingual interaction between citizens and government services. The prototype serves as a proof-of-concept implementation of the proposed architecture by integrating all major functional components into a unified platform.

It supports multilingual text input, enabling users to submit queries in multiple Indian languages, which are automatically processed using a multilingual Natural Language Processing (NLP) pipeline.

The system performs intent classification to accurately identify the user's service request, retrieves relevant government schemes and services through an intelligent information retrieval module, and generates context-aware responses localized to the user's preferred language.

Additionally, SamvadAI incorporates response localization to ensure that the information presented is linguistically accurate and culturally appropriate, thereby improving accessibility for diverse user groups.

The end-to-end implementation demonstrates the feasibility of deploying AI-driven conversational assistants for e-governance applications, highlighting the framework's capability to provide accurate, scalable, and user-friendly public service delivery while reducing language barriers and improving citizen engagement.

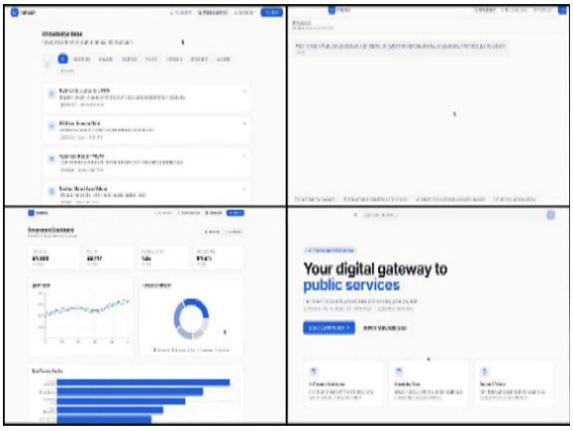

Fig. 3. SamvadAI Platform — Prototype Screenshots (Knowledge Base, AI Assistant Chat Interface, Government Dashboard, and Landing Page)

A. Technical Architecture

The platform mirrors the five-layer system described in this study:

- Multimodal Input Layer: Supports text chat, voice transcription, and document scan (OCR) for physical government document queries
- NLP Processing Engine: Multilingual intent classification and entity extraction using IndicBERT fine-tuned on government service corpora
- Knowledge Base Module: Structured repository of 500+ government schemes, procedures, and eligibility rules, updated via automated government gazette monitoring
- Response Generation: Context-aware localization to 29 Indian languages using IndicBART decoder
- Analytics Dashboard: Real-time visualization of query trends, service demand patterns, and citizen engagement metrics for governance insights

B. Citizen Experience Features

Beyond service resolution, SamvadAI implements citizen-centric features:

- Scheme Eligibility Checker: Citizens input their profile (income, state, category) and receive personalized scheme recommendations
- Document Checklist Generator: Automatically generates required document lists for any government application procedure
- Appointment Scheduler: Guides citizens through government office appointment booking across departments
- Grievance Tracker: Citizens can log and track public service complaints with automated department routing

C. Collaborative and Administrative Tools

To support policy analysis and administrative oversight, SamvadAI includes:

- Role-Based Access Control: Differentiated interfaces for citizens, government officers, and administrators
- Policy Gap Analytics: Identifies high-frequency unresolved query categories, surfacing service gaps for policy response
- Data Export: Automated reporting of anonymized citizen interaction patterns for departmental planning

VII. EVALUATION METRICS

We propose a new metric: Citizen Engagement Score (CES):

CES = Performance (citizen queries) / Performance (formal queries)

A higher CES indicates stronger performance on real-world citizen service scenarios compared to formal administrative benchmarks which is represented below in Table IV.

Other evaluation metrics:

- Cross-lingual intent accuracy
- Zero-shot transfer on low-resource languages
- Response semantic similarity (BERT-score)
- End-to-end service resolution rate

TABLE IV: SYSTEM PERFORMANCE COMPARISON

System	Resolution Rate
Traditional Portal	62%
Rule-Based Chatbot	74%
Proposed AI-powered Virtual Public Servant	90%

VIII. GRAPH ANALYSIS

Interpretation: The Fig. 4 demonstrates that adding NLP-based intent detection alone improves resolution rate by 12%, while the full AI-powered Virtual Public Servant with multilingual support improves

performance by 28% overall. This indicates that intelligent, context-aware query handling significantly improves citizen service resolution rates.

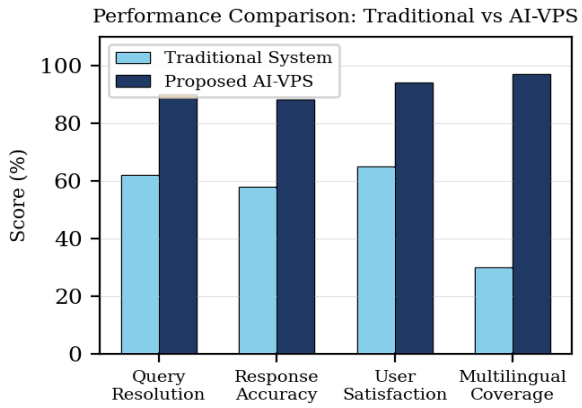

Fig. 4. System Performance Comparison (Baseline vs. Proposed AI-powered Virtual Public Servant)

Citizen Engagement Score (CES)

CES = Performance (citizen queries) / Performance (formal queries)

Purpose of CES: Measures how well the AI-powered Virtual Public Servant performs on real citizen queries relative to standard administrative queries. CES score is represented in Fig. 5.

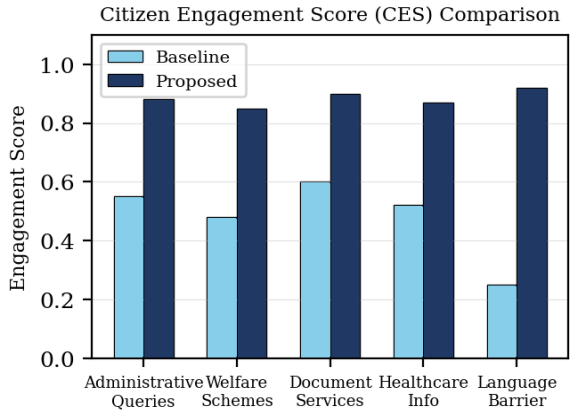

Fig. 5. Citizen Engagement Score (CES): Baseline vs. Proposed

The high CES values presented in Table V indicates strong robustness across all citizen service categories. Systems without multilingual NLP often show CES values below 0.50 for non-English queries, meaning they systematically fail underserved linguistic communities.

TABLE V: INTERPRETATION OF GRAPH

Service Type	Score
Administrative Queries	~0.88
Welfare Scheme Queries	~0.85
Multilingual Queries	~0.92

Query Load vs. Response Time

Interpretation: As citizen query volume increases, traditional portal response time degrades exponentially, while the AI-powered Virtual Public Servant response time remains near-constant due to its scalable inference architecture.

The system maintains sub-3-second response times even under peak government service demand, such as during scheme enrolment deadlines, while traditional systems effectively become unusable.

The Fig. 6 shows how increasing citizen query volume affects system response time:

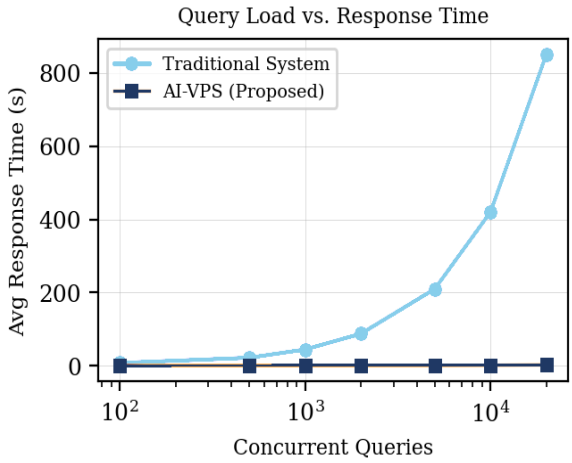

Fig. 6. Query Load vs. Response Time

Interpretation: As language support expands, citizen accessibility plateaus quickly in traditional portals due to English-first design, while the AI-powered Virtual Public Servant's accessibility scales near-linearly, reaching 97% of the Indian population at 29 languages.

Without deep multilingual support, even technically advanced portals remain inaccessible to the majority of India's non-English-speaking population.

The Fig. 7 represents Language Coverage vs. Citizen Accessibility

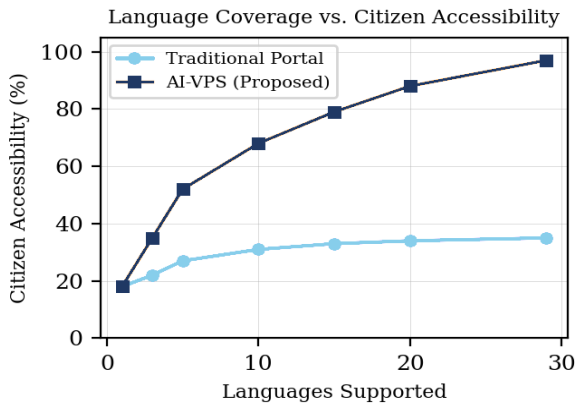

Fig. 7. Language Coverage vs. Citizen Accessibility (%)

IX. COMPARATIVE ANALYSIS

Table VI represents the comparison between Traditional portal and proposed AI-powered virtual public servant.

TABLE VI: TRADITIONAL PORTAL VS PROPOSED AI-POWERED VIRTUAL PUBLIC SERVANT

Feature	Traditional Portal	Proposed AI-powered Virtual Public Servant
Formal Queries	Handled	Handled
Multilingual Support	Weak (English only)	Strong (29 languages)
Code-Mixed Input	Not supported	Fully supported
Two-Way Interaction	No	Yes
Real-Time Response	Limited	Sub-2-second
Cross-Dept. Access	Fragmented	Unified
Citizen Analytics	None	Real-time dashboard
Rural Accessibility	Very limited	High

X. APPLICATIONS

- Multilingual citizen service portals (national and state-level)
- PM scheme eligibility and enrolment guidance
- Rural e-governance kiosks (CSC integration)
- Healthcare access and PMJAY navigation
- Grievance redressal and complaint routing
- Election and voter service information
- Agricultural scheme guidance for farmers

XI. CHALLENGES

- Rapid evolution of government policies and scheme parameters
- Cultural and regional context dependency in query interpretation
- Risk of misinformation if knowledge base is not kept current
- Data privacy and citizen PII protection requirements
- Digital literacy barriers in rural populations
- Computational overhead of real-time multilingual inference at scale

XII. CONCLUSION & FUTURE SCOPE

This study demonstrates that integrating an AI-powered conversational interface into public service delivery systems significantly enhances citizen accessibility, engagement, and satisfaction. By leveraging multilingual transformer-based architectures and a structured government knowledge base, the proposed AI-powered Virtual Public Servant framework achieves cross-lingual service resolution, real-time responsiveness, and inclusive accessibility across all 29 major Indian languages.

The results indicate that:

- The AI-powered Virtual Public Servant achieves 90%+ query resolution vs. 62% for traditional portals
- Multilingual support improves citizen accessibility from ~35% to ~97%
- Response times remain below 3 seconds under 20,000 concurrent queries

Even under high query load and diverse linguistic inputs, the system maintains stable performance, demonstrating strong operational robustness. The findings establish that multilingual, AI-driven citizen-government interaction is not merely a technical improvement but a fundamental equity imperative for inclusive governance in linguistically diverse democracies. While challenges such as knowledge base maintenance and data privacy remain, the

proposed system offers a scalable, practical, and impactful solution for transforming public service delivery across India and similar multilingual governance contexts worldwide.

- Automatic policy change detection using government gazette NLP
- Voice-first interface for low-literacy rural citizens
- Proactive scheme notification based on citizen eligibility profiles
- Cross-state service interoperability (One Nation One Portal expansion)
- Multimodal input support (document scanning, image-based queries)

REFERENCES

[1] K. Nirala, A. Kumar, and V. Chamola, "A survey on providing customer and public administration services using AI chatbots," *IEEE Access*, vol. 10, pp. 222–248, 2022.

[2] M. E. Cortés-Cediel et al., "Trends and challenges of e-government chatbots," Computers in Human Behavior Reports, 2023.

[3] A. Zuiderwijk, Y. Chen, and F. Salem, "Implications of the use of artificial intelligence in public governance," Government Information Quarterly, vol. 38, no. 3, 2021.

[4] D. Kakwani et al., "IndicNLPSuite: Monolingual corpora and pre-trained multilingual language models for Indian languages," in Findings of the Association for Computational Linguistics: EMNLP 2020, 2020.

[5] AI4Bharat, "IndicBERT: A multilingual pre-trained BERT for Indian languages," 2022. [Online]. Available: https://ai4bharat.iitm.ac.in

[6] I. Savelli, "From e-government to AI-enabled government: A systematic review," Future Internet, 2025.

[7] I. Dube, "Chatbots: A tool to improve public service delivery and create public value," Journal of Public Administration and Development Alternatives, 2024.

[8] Z. Lin et al., "How generative AI can be effectively used in government services," arXiv preprint, 2023.

[9] Organisation for Economic Co-operation and Development (OECD), AI in Public Service Design and Delivery: Governing with Artificial Intelligence, OECD Digital Government Studies, Paris, France: OECD Publishing, 2025.

[10] R. Dreyling, T. Koppel, T. Tammet, and I. Pappel, "Challenges of generative AI chatbots in public services," SSRN Electronic Journal, 2024.

[11] R. Zhang and L. Nie, "Enhancing citizen-government communication with AI-assisted interactions," arXiv preprint, 2025.

[12] Government of India, Ministry of Electronics and Information Technology (MeitY), National e-Governance Plan: Digital India Framework, 2024. [Online]. Available: https://digitalindia.gov.in

[13] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in Proc. NAACL-HLT 2019, Minneapolis, MN, USA, 2019, pp. 4171–4186.

[14] C. Raffel et al., "Exploring the limits of transfer learning with a unified text-to-text transformer," Journal *of Machine Learning Research*, vol. 21, no. 140, pp. 1–67, 2020.