ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26apr1663

A Self-Governing AI-Driven Agent for Advanced Patient Data Structuring

Dr. M. A. Kumar[1]; Pradeep Naidu Penta[2]; Grushnesh Rajagiri[3]; VTSNV Surya Teja[4]; Ramireddy Srihitha Reddy[5]

[1]Professor
[1,2,3,4,5]Department of AIT-CSE, Apex Institute of Technology, Chandigarh University, Punjab, India

Publication Date: 2026/05/09

Abstract: Abstract modern healthcare environments produce large amounts of various data, including structured EHR fields, narratives in clinical documentation, transcripts of telemedicine sessions, diagnostic imaging reports, and wearables and remote monitoring devices. A large fraction of this data is unstructured or semi-structured or inconsistently encoded, further contributing to documentation burden, impeding interoperability, and limiting the provision of timely and data-driven decision support.

This paper presents a self-directed, AI-based agent that is capable of autonomously organizing, harmonizing and integrating patient information that has multiple and heterogeneous sources. The suggested agent is a combination of transformer-based clinical natural language processing (NLP) and a reinforcement learning (RL)-controlled control layer that allows policy-directed choices regarding entity identification, semantic normalization, interrecord conflict, and general redundancy reduction.

The system is conceptually inspired by agentic AI architectures designed to operate in smart-city governance and clinical workflow orchestration; it consists of a few coordinated layers: (i) a data ingestion and harmonization layer, which connects with disparate clinical data sources; (ii) an NLP-oriented text processing module; (ii) a structured data normalization and The agent is tested by means of the digital-twin simulations and shadow-mode deployments built into the real clinical processes. Empirical findings show that it can identify entities with very high accuracy and resolve them, cut data curation and correction costs by over 80 percent, and produce quantifiable benefits on downstream tasks, like clinical decision support and predictive analytics. It is important to note that the system exhibits the ability to autonomously evolve in response to changing data schema and documentation patterns and to execute within a clearly stipulated set of safety, equity and regulatory limits.

This work bridges an important methodological gap between powerful clinical NLP methods and autonomous management of healthcare data by making the structuring of data a dynamic, self-organizing decision-making problem instead of a static, pipeline-based process. The suggested agent prepares the groundwork of scalable, high-fidelity representations of patient data that can be used to enable personalized medicine, continuous learning health systems, and powerful clinical and operational analytics.

Keywords: *Self-Governing AI, Intelligent Patient Data Structuring, Agentic AI, Clinical Natural Language Processing, Transformer Architectures, Reinforcement Learning, Healthcare Interoperability, Governance and Safety.*

How to Cite: Dr. M. A. Kumar; Pradeep Naidu Penta; Grushnesh Rajagiri; VTSNV Surya Teja; Ramireddy Srihitha Reddy (2026) A Self-Governing AI-Driven Agent for Advanced Patient Data Structuring. *International Journal of Innovative Science and Research Technology*, 11(4), 3939-3947. https://doi.org/10.38124/ijisrt/26apr1663

I. INTRODUCTION

Modern healthcare ecosystem relies on a sophisticated and fast growing set of data flows. These comprise extremely structured data, like EHR-based demographic data, vital signs, lab results, and order entries, as well as semi-structured forms and a huge amount of narrative data: clinician progress notes, discharge notes, radiology and pathology reports, transcripts of telemedicine visits, and wearable and remote monitoring wearable time-series signals.

Although rich, many of these data sources are in free text form or heterogeneous and inconsistent formats. Subsequently, this makes vital information not easily calculable, hinders cross-system and cross-institutional interoperability, secondary use in research and quality enhancement, and real-time decision support in the point of care. Clinicians are continually faced by both face-to-face and remote experiences with intensive documentation requirements. They need to spend many hours writing notes, populating structured fields, and reconciling discrepant records at the cost of face to face, high value interaction with patients [1, 2].

NLP and machine learning (ML) solutions have shown significant potential in identifying clinically relevant concepts, relationships and time information in unstructured text. The bases of these techniques are automated documentation support, clinical decision support systems, billing and coding processes, and registry reporting pipelines [3]. Concurrently, the development of agentic AI and multi-agent reinforcement learning (MARL) has also brought forth a potent paradigm of end-to-end autonomous workflow management in healthcare processes as well as smart-city environments. Triage, bed assignment, ambulance routing, inter-hospital coordination and resource allocation are some examples of tasks organized by agentic networks that are regulated by explicit policies, oversight and by digital-twin simulation environments [4, 5].

But there is still a large methodological gap. On the one hand, clinical NLP and ML systems can be frequently executed as a linear pipeline: they receive data, extract and normalize it, and send it down. Such systems are not usually given the autonomy to make decisions about when, how, and on which scale to organize or harmonize patient data across sources and do not explicitly optimize workloads or data quality over time. Conversely, agentic AI frameworks of clinical operations do not generally consider structuring and normalizing data to be one of their chief optimization aims, instead assuming that structured data is accessible and trustworthy.

To counter this, this paper suggests a self-governing, AI-based agent that is focused on smart patient data structuring. The agent combines clinical NLP models that are based on transformers with agentic control loops, and RL-based policy optimization, allowing the agent to autonomously learn strategies that trade off data completeness, consistency, redundancy reduction, and documentation effort. Gradually, the agent is able to learn to fit in to new documentation patterns, institutional coding practices, and regulatory limitations, and send uncertain or high-impact decisions to human specialists.

- *The Main Works of this Work are:*

A layer-based system architecture of a self-governed data-structuring agent that combines transformer-based clinical NLP, RL-based policy learning, and an end-to-end governance and safety shell in a synergistic way.

The elements of the formal problem framing are: re-framing patient data structuring as an autonomous sequential decision-making problem, which is subject to explicit safety, fairness, privacy, and regulatory restrictions.

The methodology that takes advantage of the digital-twin simulation and shadow-mode deployment, based on agentic healthcare and smart-city models, to train, validate and refine the policies of the agent in near-real-world scenarios in a safe manner.

An assessment plan that includes both fine-grained task-based measures (e.g., named entity recognition F1 scores, entity normalization accuracy, conflict resolution accuracy), and system-wide measures (e.g., rate of manual corrections, decrease in clinician documentation time, effect on downstream clinical and operational performance metrics).

II. RELATED WORK

- *Clinical Workflow Orchestration Agentic AI*

A multi-agent reinforcement learning framework by Warrier et al. [5] outlines a comprehensive and autonomous agentic AI to coordinate clinical operations, as well as triage, bed assignments, lab and imaging processes, patient transportation, and discharge plans. Their platform is based on such standards as HL7 FHIR and DICOM and can be used in a virtual hospital setting to facilitate real-time and self-adjusting orchestration of care processes under governance frameworks that are consistent with the NIST AI Risk Management Framework and the EU AI Act. Benefits reported are much faster response times by ambulances (by an average of 60%), shorter door-to-clinician times (by an average of 38%), and higher operating room throughput (by an average of 22%), during simulated and shadow-mode experiments.

This article highlights the revolutionary nature of agentic AI in changing reactive, human-controlled workflows into proactive, self-directed working systems. However, it is focused on macro level clinical operations and logistics and not the granular issues of clinical data extraction, normalization and reconciliation in the EHR. The structuring of patient data is an implicit assumption and not a clear object of optimization.

- *Smart Urban Infrastructure Governance Fuelled by AI*

Rathod et al. [4] present AutoGov-AI, a self-optimizing artificial intelligence platform to govern smart-cities, coordinating traffic, energy distribution, and response to emergency events with a network of decentralized edge RL agents that are supplemented by federated learning and simulations with digital twins. AutoGov-AI performs well (high decision accuracy (93.9%), high system resilience (up to 90.5%), and better latency performance) than the traditional centralized or purely edge-based systems, during normal and emergency operating environments.

Even though its use case is in urban infrastructure instead of healthcare, AutoGov-AI offers an interesting architectural prototype of modular, layered, self-managed agent systems which integrate distributed learning,

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26apr1663

simulation-based validation, and refinement of global policies. Conceptual design, with its focus on governance, scalability, and robustness, can be constructively applied to patient data management domain, to guide the design of an agent to autonomously organize and reconcile clinical data at scale.

➢ *NLP and ML of Healthcare Data*

Payal et al. [3] review the application of NLP and ML in the healthcare and pharmaceutical industries, and focus on text analytics in extracting value in the narrative medical record, laboratory narratives, and imaging reports. The applications of AI in the medical field include automatic speech recognition (ASR) to clinical dictation, AI-based documentation systems, automated registries and reporting systems, decision support systems, medical coding, and conversational virtual assistants. One of the key themes is how to convert free-text sources, such as handwritten notes, printed reports, and narrative EHR entries into structured, machine-readable formats that can be used to perform analytics, quality improvement, and research.

Both the clinical significance and technical feasibility of intelligent patient data structuring are confirmed in this literature. But the surveyed systems are primarily a pipeline paradigm: they are driven by incoming data, and execute a predetermined series of extraction and structuring processes without adaptive, policy-driven control or self-optimization.

➢ *AI and NLP in Telemedicine Documentation*

Reis [2] examines how AI and NLP can be integrated throughout the entire spectrum of telemedicine: pre-consultation, intra-consultation and post-consultation. The article explains the use of AI in collecting and summarizing patient history automatically before a remote visit, in real-time during the consultation, updating on transcription and clinical concept extraction, and creating formal documentation, discharge instructions, and education after the visit. These workflows emphasize how they might ease the burden on documentation, enhance the completeness of notes, and consistency, as well as minimize transcription errors and increase patient-centered care in virtual care.

Although this work is a strong case of how AI-based solutions can enhance and simplify telehealth records, it is not put in an agentic framework of self-governance and policy. The structures outlined are more of tool-like aids to clinicians as opposed to self-directed agents that learn, evolve, and optimize documentation and structuring plans in the long run.

➢ *Clinical NLP and Transformers*

Nerella et al. [6] give an in-depth overview of transformer architectures in healthcare, both in applications to medical imaging and structured and unstructured EHR data, social media content, physiological signal processing, and biomolecular data. The survey outlines the performance of transformer models like BERT, BioBERT, ClinicalBERT, and Vision Transformer (ViT) in tasks that are important to patient data structuring, such as clinical named entity recognition, relation extraction, and entity normalization, semantic similarity assessment and automatic ICD coding.

The authors highlight that domain adapted transformer models, trained on large biomedical and clinical datasets, have a significantly better performance on clinically oriented NLP benchmarks than generic models. Simultaneously, they highlight the unresolved issues concerning interpretability of models, fairness of algorithms, data privacy, robustness and cost of computation. Overcoming these difficulties is essential to making transformer-based NLP systems functional in production-scale healthcare settings where safety, responsibility and compliance are the most important.

➢ *Identified Research Gap*

Collectively, these research strands indicate a unique and little-discussed gap:

Healthcare and smart-city Agentic AI systems effectively coordinate the complex, distributed processes but typically assume the presence of accurate, structured information. Neither the structuring of patient data nor a central optimization objective or a locus of autonomous decision-making are typically modeled.

Healthcare NLP/ML and transformer-based models have demonstrated high results in tasks like clinical information extraction and normalization, but are typically used as fixed components. They do not have systems to automatically change their actions in response to changing data quality, clinician workload, governance constraints or downstream system feedback.

This drives the creation of a self-regulating, AI-powered agent with the main purpose of organizing, normalizing, and aligning the patient information across a variety of different sources. The goal of such an agent is to minimize manual intervention and documentation efforts, to reduce redundancies or contradictory entries, and to be sure that data curation is compliant with specific safety, fairness, and regulatory limitations.

III. FORMULATION OF PROBLEMS AND SYSTEM OBJECTIVES

➢ *Problem Definition*

We assume the issue of dealing with the heterogeneous constantly incoming streams of patient data based on:

- EHR elements, which are structured (demographic profiles, vital signs, laboratory test results, medication orders, procedure codes).
- Narrative content, such as admission and discharge summaries, progress and consultation notes, radiology and pathology reports, operative reports, and telemedicine or virtual care transcripts are unstructured.

Semi-structured documentation (e.g., templated forms, checklists and partially standardized clinical questionnaires) and wearable and monitoring device outputs (time-series measurements, alerts and derived features).

The overall aim is to create an autonomous, self-governing agent that is able to consume these heterogeneous inputs, decide what parts of the data may need to be structured

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26apr1663

or normalized, choose the most suitable extraction and coding strategies, resolve the conflicts between sources, and continue to have a coherent, high-quality, longitudinal picture of every patient. This should be done without breaching limits regarding safety, fairness, privacy, institutional policies, and clinician oversight, as well as without on-going efforts to lessen the overall documentation and data-curation burden on healthcare professionals.

➢ *Design Objectives*

The suggested agent is aimed at achieving the following general goals:

- *Autonomy:*

Allow context-sensitive, decentralized decision-making in data structuring, normalization as well as integration with minimal human oversight. Informed policies, clear limitations, and operating environment drive decisions, thus minimizing the need to use manual intervention.

- *Interoperability:*

Use healthcare data standards, especially HL7 FHIR and DICOM, to guarantee interoperability with any current electronic health record (EHR), imaging repositories and new agent-based infrastructures throughout the healthcare ecosystem.

- *Safety and Governance:*

Ensure privacy, security, fairness, and clinical safety with seamlessly integrity into a runtime assurance mechanism and detailed audit trail. Policy checks are provided on all actions and system behavior is constantly checked to ensure regulatory and institutional compliance.

- *Adaptivity:*

Alternatively, to improve the process of structuring and normalization, constantly improve it based on the feedback received in the form of manual corrections, downstream decision errors, and performance indicators. To develop policies with better accuracy and enhanced retrospective correction requirements are minimized; reinforcement learning techniques are used to develop the policies as time progresses.

- *Scalability:*

Support high volume, multi-modal data processing between heterogeneous departments, institutions and networks. The system is to be efficient in distributed and federated learning systems that maintain data locality, regulatory jurisdiction and institutional integrity.

IV. PROPOSED ARCHITECTURE

Our architecture is based on ideas of AutoGov -AI [4] in the governance of smart cities, and agentic AI models of clinical workflow alignment [5] to create a multi-layered, modular architecture of a self-organizing patient data structuring agent. The architecture isolates ingestion, processing, decision-making and governance concerns, making the architecture easily extendable and with a strong lifecycle management.

➢ *Data Ingestion and Harmonization Layer:*

This is a data-processing layer that ingests and harmonizes data for processing.

This base layer has the duty of obtaining and initializing the data of various types of clinical and operational sources, such as:

- *EHR Systems:*

Available over FHIR-based APIs and event streaming capabilities, delivering structured clinical information like encounters, orders, medications, and problem lists.

- *Telemedicine Platforms:*

Providing consultation transcripts and text based on audio/video sessions and including synchronous and asynchronous interactions in virtual care.

- *Imaging Systems (PACS/RIS):*

Consuming DICOM metadata, imaging study description and related radiology or cardiology reports.

- *Wearables and other tracking devices:*

Including structured and semi-structured measures (e.g. vital signs, activity measures, continuous monitoring data).

In this layer, schema mapping, data cleaning, and first harmonization are done. The heterogeneous inputs are converted into the intermediate canonical form, e.g., standardized FHIR bundles, and supplemented with source identifiers, timestamps, and provenance. This offers a regular substrate to downstream NLP and normalization modules.

➢ *Unstructured Data Processor is Based on NLP*

Unstructured and semi-structured texts—such as clinical notes, operative and radiology reports, telehealth conversation transcripts, and call center logs—are processed by a transformer-based clinical NLP pipeline. Core components include:

- *Tokenization and Pre-processing:*

De-identification of health data to eliminate or obscure the presence of the data, segmentation of sentences, standardization of medical abbreviations, and other pertinent text cleaning, based on de-identification techniques presented in transformer-based clinical NLP studies [6].

- *Named Entity Recognition (NER):*

Clinically relevant entity detection (diagnosis and problems, medication, procedure, laboratory test, symptom, and social determinant of health, etc.) is automatically detected. It is executed with domain-specific transformer networks (e.g., BioBERT, ClinicalBERT, and their variations).

- *Relation Extraction (RE):*

The recognition and representation of semantic relations among entities, e.g., drugdose, testresult, conditiononset, and problemprocedure, leading to a formalized expression of clinical facts, which is usually thought of as a clinical knowledge or fact graph.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26apr1663

- *Entity Normalization:*

Mapping of extracted entity mentions to standard terminologies and code sets, including SNOMED CT, ICD and LOINC. Entity normalization methods based on transformers are employed to improve the accuracy and address lexical variation, synonyms, and contextual nuances.

- *Summary and Structure of Notes:*

Creation of structured, sectioned clinical notes (e.g., History, Assessment, Plan) based on free text, and consistent with existing EHR documentation templates and telemedicine workflows [2]. This aids human readability and sub downstream computation.

This module, together, converts unstructured textual information into a rich and structured format of clinical ideas, including standardized coding and prepared to be combined with other parts of the patient record.

➢ *Structured Data Normalization and Integration Module*

In cases where data are already available in form, like laboratory data, vital signs, medication lists and procedure codes, this module is concerned with consistency, completeness and longitudinal consistency:

- *Normalization of Units and Codes:*

Measurement unit, reference range and coding scheme validation and correction. Local or proprietary codes are coded to accepted standard and any apparent anomalies noted or rectified.

- *Record Linkage & Record Deduplication:*

Find and unify overlapping or duplicate records between different encounters, providers and systems. Similarity metrics and learned embeddings which are based on machine learning are used to facilitate powerful entity resolution and merging strategies.

- *Longitudinal State Maintenance of Patients:*

Building and upkeep of a temporally sequenced, longitudinal perspective of clinical condition of the patient. This representation can be directly used to input to transformer-based EHR sequence models (e.g., BEHRT, Med-BERT-like architectures) to support predictive analytics and decision support as optional downstream uses.

➢ *Self-Governing Agent Layer*

The system is centred on agentic decision-making layer which coordinates the application of structuring, normalization and integration steps over time. This layer is based on adaptive, policy-driven control inspired by the multi-agent rule learning (MARL) and hierarchical RL paradigms used by AutoGov-AI [4] and agentic clinical workflow orchestrators [5].

- *Key Functions Include:*

✓ *Policy-Driven Structuring Decisions:*

Decision on when to use more aggressive normalization (e.g. to auto-merge two or more entries with very similar values, to normalize heterogeneous fields) or when to take a more conservative stance and retain the raw details. Policies that have been learnt, confidence estimates and constraints of governance influence decisions.

✓ *Conflict Resolution:*

In case of conflicting information between disparate data sources (e.g., divergent lists of problems, different medication dosage) the agent chooses and implements a resolution strategy. Some of the alternatives can be majority voting, source reliability weighting, recency heuristics or escalation to human adjudication.

✓ *Task Prioritization:*

Prioritization of structuring activities according to their anticipated contribution to the active clinical processes, quality improvement initiatives, registries, and downstream analytics. This priority level is similar to the resource-allocation rationality used in agentic clinical operations models [5].

✓ *Reinforcement Learning (RL):*

On-going policy optimization based on the feedbacks like the rate of manual corrections, the occurrence of downstream clinical decision errors, audit reports, and improvement in the overall data quality. The use of RL aims at minimizing the future interventions, and optimizing the performance on several objectives.

The layer can be implemented as a single global agent covering a whole institution or a federation of specialized agents (e.g. medication-specific agents, problem-list specific agents, clinical-documentation specific agents). Such agents are able to synchronize with each other based upon multi-agent coordination schemes that have been established in clinical workflow coordination [5].

➢ *Governance, Safety, and Human-in-the-Loop*

By combining the ideas of governance and runtime assurance shells in healthcare agent models [4,5], all agent choices are followed by explicit, machine-readable policies that state legal, institutional and clinical limitations.

- *Key Components Include:*

✓ *Runtime Assurance Shell:*

A security coating that intercepts and scrutinizes suggested actions, averting unprotected and non-compliant activities, including removal of important diagnoses, creation of verified records, or massive volume changes. It is also able to impose rate constraints and structural protections about large-scale modifications.

✓ *Confidence-Based Escalation:*

Low model confidence or large ethical/clinical ambiguity in decision lead to human action instead of automatic action. This prevents edge cases and high-risk situations to be processed by unqualified clinicians or data stewards.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26apr1663

- ✓ *Provenance and Auditability:*

Detailed logging of all transformations, such as the input data, model versions, policy rules used, rationale behind decisions, and time-stamps. This facilitates similar traceability and accountability as provenance graphs suggested in agentic AI governance literature [5].

- ✓ *Non-prejudice and Equity Surveillance:*

Continuous monitoring of the quality of structuring and error distributions in demographic subgroups to find possible algorithmic bias. Based on perceived fairness in the deployment of transformers and LLMs [6], corrective actions may be triggered (e.g., re-training a model, policy modifications) in the case of detected disparities.

V. METHODOLOGY

- ➢ *Model Selection and Pretraining*

The selection of models and pretraining occurs during this step.

Based on survey data about transformer-based clinical NLP [6], pretrained models with domain optimization, including BioBERT, ClinicalBERT, and PubMedBERT are used to initialize the NLP components. These architectures have proved to perform well with biomedical NER, relation extraction, question answering and semantic similarity tasks.

These models are then optimized on institution specific, de-identified corpora of clinical notes and telemedicine transcripts to aid:

- *Entity Recognition and Concept Extraction:*

High fidelity diagnosis, medication, procedure, laboratory tests, and social determinants of health.

- *Relation Extraction:*

Building of clinically meaningful tuples and relations (e.g., medication–indication, test-finding, condition-onset) that construct the building blocks of structured clinical representations.

- *Entity Normalization:*

Proper mapping of entities extracted to standard terminologies and code sets (SNOMED CT, ICD, LOINC, etc.) among local documentation variations.

In the case of structured data, record-linkage and deduplication machine learning models are used together with rule-based integrity checks (e.g., consistency rules, plausible ranges) to guarantee reliability and minimize spurious merges.

- ➢ *Training Agents through reinforcement learning*

The self-directed agent policy is trained in the context of reinforcement learning that is similar to MARL-based solutions in smart city control [4] and planned clinical procedures [5].

- *State Representation:*

Captures the existing composition and wholeness of the patient record, data quality metrics (including missingness, redundancy and inconsistency), the operational workload (including unstructured note backlog awaiting structuring), and operational governance or compliance flags.

- *Action Space:*

Incorporates the choice of structuring strategies (e.g., strict vs. conservative normalization), merging or separating of entities, delaying actions or routing certain items to human reviewers.

- *Reward Function:*

Combines several goals: a decrease in the load of manual correction, a reduction in the time spent by clinicians on documentation, an increase in the completeness and internal consistency of data, and penalties based on governance breaches, unsafe transformations, or downstream clinical errors.

- *Training Environment:*

A simulation environment simulates historical EHR edit logs, telemedicine visits and previous structuring workflows in a digital-twin style. This enables the agent to experiment, study, and test strategies in a risk-free environment, similar to those of digital twins outlined in AutoGov-AI [4] and agentic health systems [5].

Off-policy RL algorithms (e.g., Soft Actor-Critic) and hierarchical RL schemes can be employed. The policies at the higher level regulate the overall structuring policies and conflict-resolution patterns, whereas the controllers at the lower-level implement particular transformation actions. This top-down breakdown is reflective of agent structures in crisis and operations management in healthcare [5].

- ➢ *Continuous learning and human feedback*

It is first instantiated in a so-called shadow-mode, with the agent creating suggested structuring and normalization actions, but does not directly edit live production records. These proposals may be accepted, rejected, or edited by human users, who are clinicians, coders, and data stewards.

Their answers form a stream of feedback with labels that are constantly used to do two things:

- *Transformer Model Refinement:*

Gradual refinement of NLP elements to enhance accuracy of extraction, disambiguation and normalization performance based on the institution documentation style and patient population.

- *Policy: RL Optimization:*

Continued optimization of the decision policies of the agent based on what actions were accepted or rejected and also based on the downstream results observed. This is based on the previous agentic healthcare work of the so-called shadow-mode + digital twin evaluation and learning paradigm [5].

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26apr1663

The system will acquire more autonomous, reliable structuring under this continuous feedback loop and retain the human control over complex or high-risk cases.

VI. EVALUATION PLAN

➢ *Task-Level Metrics*

In order to measure the performance of the NLP and structuring components on a granular level, we suggest the following measures:

- *NER, Relation Extraction and Normalization Performance:*

Precision, recall and F1-score compared to expertly annotated gold standard corpora.

- *Structuring Accuracy:*

Concordance with agent-generated structured fields and expert coder/clinician fields in the EHR, based on the use of relevant accuracy and concordance measures.

- *Redundancy and Conflict Reduction:*

Alterations in the proportion of duplicate, overlapping or conflicting entries per patient record compared to baseline workflows.

These measures are correlated with the known clinical NLP standards and assessment schemes applicable in the research of transformer-based information extraction [3,6].

➢ *System-Level and Workflow Metrics*

In order to measure the wider effect of the agent on clinical processes and the user experience, we use the metrics of the telemedicine documentation literature [2] and agentic orchestration literature [5]:

- *Documentation Time:*

Shortening of clinician time on manual data entry, note structuring and reconciliation per encounter, calculated through time-and-motion research, or EHR interaction logs.

- *Manual Correction Rate:*

Fraction of structured fields that need to be changed following the agent processing, which itself is a direct measure of agent accuracy and reliability in the real world application.

- *Turnaround Time:*

Time elapsed between receiving of data (e.g., after a consultation or a diagnostic report is available) and having a completely structured and integrated patient record.

- *Resilience and Latency:*

The capability of the system to maintain throughput and performance with the peak loads both in terms of average and tail structuring latency and the rates of successful completion of tasks during the stress conditions. These analyses are parallel to resilience and responsiveness analyses in previous agentic orchestration frameworks.

Taken as a whole, these task-level and system-level measures allow assessing both the technical and operational usefulness of the self-organizing patient data-structuring agent.

➢ *Simulation-Based and Shadow-Mode Evaluation*

To evaluate the proposed agent rigorously, we suggest using the protocol of a two-stage evaluation, similar to the one used by Warrier et al. [5], which includes the following steps:

- *Digital-Twin Simulation:*

During the initial phase, a digital-twin setting would be simulated with the aid of retrospective electronic health record (EHR) and telemedicine data. This simulated environment would simulate realistic data ingestion and normalization and structuring workflows, allowing systematic experimentation with agent policies under controlled variations in workload intensity, data sparsity, data quality and error patterns. By emulating the various operational conditions, including peak clinical workloads, noisy or incomplete records and mixed data formats, researchers can run the agent through its stress-tests to ensure its stability, responsiveness and policy flexibility before any actual deployment.

- *Shadow-Mode Deployment:*

This phase is a shadow-mode deployment in real clinical settings. In this case, the agent is run in parallel with the presence of documentation and data management processes but it does not affect production systems directly. Rather, it constantly produces suggested structuring actions, say, coding suggestions, entity normalization or record linkage, as human clinicians and staff continue to practice their normal documentation. The agent-human differences and convergences in the agent suggestions and the human performance produce a rich evaluation signal, providing a fine-grained understanding of the performance, safety, and clinical acceptability. These comparison traces can in turn be high-value supervised and reinforcement learning examples to perfect the underlying NLP models and the decision policies of the agent.

A battery of statistical tests, ablation studies, and sensitivity analyses ought to be carried out to unravel the effects of individual architectural elements. By adhering to analytical methods applied to complex, multi-component agentic systems [4,5], researchers will be able to independently quantify the contributions of (a) transformer-based clinical NLP modules, (b) policy optimization via reinforcement learning and (c) embedded governance, safety and fairness constraints. Such a component-based analysis is used to make principled design decisions and to tune the entire system based on evidence.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26apr1663

VII. DISCUSSION AND FUTURE DIRECTIONS

The self-directed patient data structuring agent suggested in this paper is a combination of multiple intersecting lines of research: transformer-based clinical natural language processors [6], more general NLP/ML pipelines to curate healthcare data [3], and agentic reinforcement learning systems originally designed to be used in smart cities [4] and self-managing clinical processes [5]. The agent aims to significantly decrease the administrative load on clinicians, improve the completeness and accuracy of clinical data and realize interoperable, longitudinal, and analytics-ready patient records by prioritizing the structuring, normalization and integration of patient data to become a primary design objective, instead of it being a by-product of documentation. Individualized care tracks, sophisticated decision support solutions, analytics on a population scale, and effective cross-institutional coordination can be based on such records, respectively.

Despite this potential, this vision poses a number of substantive challenges which need to be addressed in further research in a systematic way:

- *Data Privacy and Security:*

The training of large transformer models and reinforcement learning policies on sensitive clinical data brings about high privacy, confidentiality, and cybersecurity risks. There is an urgent need to change and improve the governance structures, access control policies, de-identification and pseudonymization policy, and secure logging policy, based on the recommendations of current agentic AI models [5] and transformer surveys [6]. Differential privacy, secure enclaves, and high auditability are methods worth considering in order to make sure that the standards of regulatory actions (e.g., HIPAA, GDPR) and institutional policies are adhered to.

- *Interpretability, Transparency, and Bias:*

Structured diagnoses, coded procedures, medication lists and other normalized fields of data, which are the outputs of the agent, directly inform downstream clinical decision-making. Improperly structured or biased information can propagate minor, cascading errors which can be hard to spot ex post. Therefore, the system should have means of interpretable decision traces, rationale logs, and explanations about any key structuring decisions that can be viewed by the user. Algorithms should be monitored continuously to ensure their bias and fairness across patient subgroups, and routine audits and performance dashboards that are subject to clinician, data steward and governance board scrutiny.

- *Integration and Interoperability in Hospital IT ecosystems:*

To integrate an agentic system into complex and heterogeneous hospital IT infrastructures, significant technical integration and management of organizational change is necessary. The agent should be able to work with EHRs system, imaging, laboratory information systems and auxiliary applications without any issues. In this regard, it is important to comply with standards of interoperability, including HL7 FHIR on clinical data exchange and DICOM on imaging, as stressed in the previous study [5]. Simultaneously, organizations will have to coordinate workflow redesign, employee training, alignment of stakeholders, and governance procedures to guarantee safe, acceptable and sustainable adoption.

- *In the future, there are a number of potential avenues of research:*

Federated and Collaborative Learning: A natural extension of the proposed framework Federated agents can be thought of as collaboratively learning to structure policies in multiple hospitals or health systems, without the centralized aggregation of raw patient data. Based on the ideas of federated learning developed in smart city infrastructures [4] and new healthcare AI, every institution would be able to locally optimize and update agent components (e.g., policy networks, language models) and only exchange model parameters or gradients via secure aggregation protocols. This solution can encourage the cross-institutional generalization and strength and maintain the data sovereignty and privacy.

- *Integration with Specialized Large Language Models:*

One more way is to introduce large language models that are explicitly specialized and pre-trained on biomedical, clinical, and health policy knowledge, as reported in recent surveys of transformer-based large language models [6]. These large models can be fine-tuned using parameter-efficient fine-tuning algorithms, including adapters, LoRA or prompt-tuning, without violating computational resource limits or governance policies. This form of integration may improve the agent in terms of its capacity to read complex free-text notes, to summarize multi-modular information and to produce high quality structured representations.

- *Patient-Consumer Feedback and Human-Centered Design:*

Lastly, patient-in-the-loop system integration provides a way to more effectively align the system with the principles of ethical and human-centered AI in telemedicine literature [2] and agentic AI literature [5]. As an illustration, patients might have an access to organized summaries of their records safely and provided with the means to mark inconsistencies, lack of information, or a sense of misrepresentation. Their responses can be included in a refinement loop, which will enhance data fidelity and transparency, trust, and the ability to make decisions together between patients and healthcare providers.

VIII. CONCLUSION

This paper presented a self-regulated, AI-driven patient data structuring agent, aimed at filling the gap between non-intelligent, rule-based NLP pipelines and entirely autonomous healthcare data management systems. The proposed architecture through a strong connection between transformer-based clinical NLP modules and an agentic control layer based on reinforcement learning can extract, normalize, reconcile, and integrate patient data across various and heterogeneous sources autonomously. By so doing, it aims at lowering unnecessary manual data entry,

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26apr1663

inconsistencies, and liberating clinical personnel to concentrate more on patient care.

The agent is acting within given safety, equity and governance boundaries based on agentic models that were initially developed as smart urban infrastructures and coordination of clinical workflows. Its behavior can be optimized and tested in digital-twin simulation and shadow-mode deployments to enable robust, low-risk testing and improvement cycles before the adoption of the full operation.

The agent aims to address a fundamental need of modern healthcare systems, i.e., the transformation of fragmented, siloed and unstructured clinical data into high-quality, interoperable and semantically-enriched patient records. These records are critical pillars to accuracy and personalized medicine, real time decision support at the point of care, massive epidemiological and operational analytics, and coordinated care across institutional and geographical borders.

Future directions will focus on scaling out the framework along three main axes: (1) federated and distributed learning with a wide range of institutions to promote generalizability without sacrificing privacy; (2) more explicit integration with domain specific large language models, using parameter efficient adaptation methods; and (3) systematic integration of patient centric feedback loops and participatory design practices. Collectively, these initiatives are expected to establish the system to be more autonomous, reliable, as well as transparent, and thus enhance the trust of clinicians and their patients in AI-powered data infrastructure and the overall healthcare AI agenda of responsibility and human-centeredness.

REFERENCES

[1]. M. Payal, T. Ananth Kumar, K. Suresh Kumar, “Integrating Natural Language Processing (NLP) and Machine Learning Techniques for Healthcare Industries,” International Journal of Innovative Research in Computer and Communication Engineering, vol. 10, no. 5, 2022.
[2]. T. C. Reis, “Artificial Intelligence and Natural Language Processing for Improved Telemedicine: Before, During and After Remote Consultation,” Atencion Primaria, vol. 57, 103228, 2025.
[3]. S. Nerella et al., “Transformers in Healthcare: A Survey,” University of Florida, preprint, 2023.
[4]. D. B. Rathod et al., “AI-Driven Decision-Making Architectures for Future-Ready Smart Urban Infrastructures,” 2025 IEEE 5th International Conference on ICT in Business, Industry and Government (ICTBIG), 2025.
[5]. A. Warrier, A. K. S. et al., “Autonomous Agentic AI for Clinical Workflow Orchestration: Self-Managing Healthcare Operations,” Proceedings of the 6th International Conference on IoT Based Control Networks and Intelligent Systems (ICICNIS 2025), IEEE, 2025.
[6]. Additional foundational works on agentic AI and clinical NLP are cited within the text as per the original provided references.