

A Longitudinal Patient Preference-Aware Treatment Recommendation System Using Multi-Source Sentiment Fusion and Adaptive Utility Modelling

Jitendra Balaji Mendadhala¹, Jacob Victor Ganta²

¹M.Tech Graduate, Department of Computer science KI University, Vijayawada, AP, Email: jitendrabalaji1705@gmail.com

²Professor Department of Computer science KI University, Vijayawada, AP, Email: jacob.victor@gmail.com

Abstract

Treatment recommendation systems using artificial intelligence have shown significant gains in predictive accuracy and clinical decision support, but an inherent design challenge exists in virtually all existing frameworks, namely that patient preferences are either considered as fixed binary inputs, they are viewed as auxiliary features, or are not part of the optimization problem at all. There are severe clinical implications of such a gap, with guidelines that do not incorporate changing patient values being always linked with declining therapeutic adherence, diminishing patient contentment, and structural devaluation of shared decision-making, which is progressively implemented as a fundamental principle of ethical clinical care by regulatory authorities such as the FDA, the European AI Act, and the American Society of Clinical Oncology. To overcome this daunting constraint, this paper presents PreferenceNet, a longitudinal, preference-sensitive, treatment recommendation system, which can model, track, and optimize against patient preferences during repeated clinical visits over time. The system operates through four tightly integrated modules: a Multi-Source Preference Signal Extraction pipeline fusing patient-reported outcome surveys, clinical consultation transcripts processed via fine-tuned BioBERT, and wearable behavioral time-series encoded through a bidirectional LSTM into a unified preference representation; a Temporal Preference Drift Model employing a Milestone-Conditioned Hidden Markov Model that infers latent preference states and detects clinically meaningful transitions across longitudinal visit sequences from MIMIC-IV; an Adaptive Utility Optimization Layer with patient-specific, time-varying weights across clinical efficacy, preference alignment, side-effect tolerance, and treatment burden, updated via Bayesian inference at each clinical encounter; and a Preference-Concordant Recommendation module quantifying outputs through two originally defined metrics — the Preference Concordance Score and Utility Alignment Index — delivered via a clinician-facing interpretability dashboard. Assessed on 12,847 patients in comparison to four fixed baselines, PreferenceNet shows an AUC of 0.87, Preference Concordance Score of 0.79, Utility Alignment Index of 0.81, and Preference Drift Detection Accuracy of 83.2 percent, and an 88 percent clinical agreement rate under an additional vignette study, which together indicate.

Keywords: Hidden Markov Model, multi-source sentiment fusion, adaptive utility, BioBERT, Clinical natural language processing, Multi-source data fusion, Electronic health records, MIMIC-IV, Personalized medicine, Chronic disease management.

1. Introduction

Chronic diseases represent one of the major challenges faced by healthcare systems today at the global level. According to estimates, diabetes mellitus affects 537 million adults globally, a number predicted to increase up to 783 million by 2045. Additionally, breast cancer has become the most commonly detected form of cancer in women worldwide, surpassing even lung cancer in prevalence. However, treatment of chronic diseases requires significant efforts since the course of therapy is usually complicated and includes many options. Most importantly, patients themselves the primary decision makers in any case cannot be reduced to any biomarkers and survival curves as they have specific values and preferences which must be respected and considered. The practice of shared decision making has been increasingly adopted in modern medicine by organizations like WHO, ADA, and ASCO as both an ethical requirement and a way to ensure better results. When included in the process of treatment decision-making, patients improve adherence and experience higher levels of satisfaction and better results. On the other hand, ignoring their opinions is known to yield a predictable set of negative outcomes. Nonetheless, recent years have seen an impressive development of various artificial intelligence-based recommendation systems for treatment plans. From deep learning models like RETAIN and DeepCare to systems based on reinforcement learning and deep Q-learning and graph neural networks, these architectures have shown promise regarding the effectiveness of

recommendations, validation on clinical data, and performance. Such approaches have indeed brought much progress and deserve credit for what they accomplished regarding outcome predictions and clinical workflows. Still, there is one common disadvantage found across virtually all such systems, which makes them significantly less effective than they could be. Specifically, these algorithms attempt to maximize some metric of medical success survival probabilities, reduction in readmission rates, normalization of biomarkers, and others but neglect the values of patients themselves, representing them either as static inputs or binary variables which fail to reflect actual preferences. In a study reviewing 67 papers analyzing AI recommendation systems based on patient values published in 2025, the following shortcomings were noted lack of real-world testing and failure to actively include preference information. As confirmed by a 2024 systematic review in *npj Digital Medicine*, AI-driven decision tools can indeed add value when assisting in diagnostic screening and prognostication communication. Yet, the literature on the topic does not mention any system where patient preferences would be considered in the development of an algorithm for recommending treatment courses. Importantly, preferences are dynamic constructs that should be taken into account when considering patients' clinical trajectories. For instance, a patient who received a diagnosis of early-stage breast cancer could prefer radical treatment to achieve better chances of survival and endure severe side effects from treatment. Three years later, after cancer recurrence and another course of therapy, the patient's priorities could change drastically, with an increased focus on maintaining a high quality of life, avoiding hospitalization, and preserving functional independence instead of improving survival rates. This kind of longitudinal shift in patients' preferences is well-known to clinicians in the field of oncology and chronic diseases, yet no AI solution is currently capable of detecting and addressing such changes. PreferenceNet aims to directly fill this void by leveraging the system's objectives derived from the clinical insights provided earlier. First, PreferenceNet focuses on continuous, multi-source extraction of preference signal while refraining from relying solely on a single instrument or survey. PreferenceNet concurrently extracts preference signal from structured measurements based on patient-reported outcome measures (PROs), unstructured data based on clinical consultation conversations by applying fine-tuned BioBERT to the MedDialog-EN dataset, and passive data from wearable device activity measured through a long short-term memory neural network. All three sources provide a unified, cross-modal preference signal per each clinical visit. Second, PreferenceNet addresses the longitudinal nature of preference evolution by modeling preference drifts through utilizing a Milestone-conditioned Hidden Markov Model for modeling patient preference state transitions at each clinical visit based on the longitudinal visit sequences from MIMIC-IV electronic health records. Specifically, PreferenceNet will track and predict patient preference drift triggered by certain milestone events, such as relapse, remission, and treatment switching. Third, unlike using static patient-level parameters, PreferenceNet will learn patient-specific, dynamic utility weights based on multiple dimensions (efficacy, preference match, side effect tolerability, and treatment burden) at each visit via Bayesian inference to capture the patient's latest preference state. Finally, the preference-based recommendation should align with clinicians' perspective, therefore, PreferenceNet generates its recommendations together with two newly designed metrics (Preference Concordance Score and Utility Alignment Index) through the structured, yet interpretable dashboard interface that is meant to supplement, not replace, clinical decision-making.

2. Related Work

Treatment recommendation systems have evolved significantly over the past decade, moving from simple rule-based clinical guidelines toward sophisticated machine learning frameworks. This section reviews four interconnected streams of literature: AI-based treatment recommendation, patient-reported outcomes in clinical AI, shared decision-making tools, and temporal modelling of patient states in healthcare.

2.1 AI-Based Treatment Recommendation Systems

Initially, treatment recommenders were developed using collaborative filtering techniques and statistical models [1]. With the advent of deep learning, researches have shown that recurrent neural networks can effectively capture clinical events and provide recommendations by analyzing electronic health records (EHRs). These approaches include RETAIN [2] and DeepCare [3]. In recent times, reinforcement learning algorithms have proven their efficacy in providing effective treatment recommendations for sepsis patients [4] and diabetes patients by fine-tuning medication dosage [5]. Another milestone in the development of treatment recommenders was the introduction of

graph neural networks, which are capable of capturing the intricate relationship between disease diagnosis, medicines, and patients' demographic information [6], [7]. However, there exists an inherent deficiency in the current state-of-the-art treatment recommenders, wherein the recommendation process is considered to be solely based on the clinical outcomes, irrespective of the patient's preference for treatment [8].

2.2 Patient-Reported Outcomes in Clinical AI

There are numerous PRO measures available to clinicians today, including PROMIS-29 [9], PRO-CTCAE [10], and EQ-5D [11]. Some research has used PRO data in clinical AI models as predictors of outcomes, specifically readmission [12] and treatment response prediction [13]. The problem is that in all these cases, the PRO measure is used as just another predictor without consideration of its dynamic characteristic as signalling patient preferences. As pointed out by Jensen et al. [14] in a review from 2023, although the inclusion of PRO measures in clinical AI pipelines leads to better predictions, their active feature as signals of patient preferences is not used yet. At the same time, the process of PRO collection is not consistent in everyday practice [15].

2.3 Shared Decision-Making and AI Decision Aids

Decision-making support (DMS) has become the cornerstone of contemporary medicine, being supported by major guidelines issued by the WHO [16], ASCO [17], and the ADA [18]. DMS applications have been built with the purpose of informing patients about their treatment options, potential outcomes, and tradeoffs through decision aids powered by AI [19]. In their 2024 review article published in *npj Digital Medicine*, Lee et al. [20] considered 43 decision aids that employ AI and concluded that, although the reviewed tools successfully enable the provision of information in consultations, no study described an integration of patient preferences signals into algorithms driving medical recommendations. Similarly, in their 2025 scoping review examining 67 AI-powered systems used for assessing patient values, Campos et al. [21] reported a critically low rate of real-world validation of such systems. Taken together, these results suggest that there has been no bridge built between SDM applications and recommender AI tools.

2.4 Temporal Modeling of Patient State in Healthcare AI

Modeling of patients' clinical states as they change over time has been studied extensively. HMMs have been successfully used for modeling of disease progression for illnesses like Alzheimer's [22], chronic kidney diseases [23], and others. LSTM neural network architectures have proven themselves effective in modeling longitudinal patterns in EHR for predicting sepsis [24] and ICU deterioration risk [25]. More recently, Transformer-based models such as BEHRT [20] and MedBERT [24] have outperformed other deep learning algorithms for the same purposes and achieved superior results in modeling long-term temporal dependencies in patient data. Still, regardless of their type, all algorithms for modeling longitudinal patient states presented in the existing healthcare AI literature have only focused on modeling clinical states, not patients' preference states [23]. It is crucial to understand that while patients' clinical states and patients' treatment preferences are related concepts, they cannot be considered one and the same thing. In other words, the process of modeling patients' clinical states and modeling their preferences should use completely different approaches. As far as we are aware, there is currently no research done on this subject [24], [25].

3. SYSTEM METHODOLOGY

The proposed PreferenceNet can be considered as an end-to-end treatment recommendation architecture that considers the dynamic preference of patients as a longitudinal feature. PreferenceNet uses data from two types of databases: the first one is MIMIC-IV, which provides information on structured longitudinal clinical features such as diseases, medications, laboratory tests, procedure history and other parameters that can be recorded during several visits to the clinic. Another data type comes from the MedDialog-EN corpus, which is a massive dataset containing real clinical consultation dialogues in English. The purpose of incorporating this data into our model is to extract unstructured preferences from conversations between physicians and patients. Hence, the two types of data can be viewed as two sides of the coin: the former represents a trajectory of clinical features of the disease of the patient, while the latter represents a trajectory of his/her preferences over the same period of time. PreferenceNet consists of four sequentially linked modules. In Module 1, multi-source preference extraction involves the merging of structured patient reported outcome encodings with extracted natural language preference signals and time-series

behavior data into a single preference representation vector at each clinical visit. Module 2 takes this series of vectors per visit and models its temporal dynamics using a Milestone-Conditioned Hidden Markov Model, obtaining latent preference states along with transition points at the longitudinal level of patients' visits. Module 3 integrates inferred preference states into the Adaptive Utility Optimization Layer and computes time-varying utility scores for each possible treatment by finding a balance between clinical effectiveness, preference matching, side effects tolerance, and burden of the treatment in accordance with Bayesian-updated preference weights. Finally, Module 4 gives the treatment recommendation in terms of the original metrics and an interpretable clinician dashboard. Figure 1, which demonstrates the end-to-end PreferenceNet pipeline from raw multi-source inputs through preference fusion, temporal drift modeling, utility optimization, and final recommendation output.




Fig 3.1: System Methodology

3.1 Module 1: Multi-Source Preference Signal Extraction

3.1.1 Patient-Reported Outcome Survey Encoding

For each clinical visit t at time v_t , patients fill out standardized self-reports of health measures that are part of the PROMIS-29 profile, providing quantitative self-assessments across domains including physical function, anxiety, depression, fatigue, sleep disturbance, social participation, and pain interference. The patient's score for each item is collected on a five-point Likert scale and then converted into a continuous variable through the equation:

$$\tilde{x}_i = \frac{x_i - x_{\min}}{x_{\max} - x_{\min}}, \quad \tilde{x}_i \in [0,1]$$

where x_i refers to the actual response for item i , while $x_{\min}$ and $x_{\max}$ stand for the hypothetical minima and maxima for the scale. The normalized item vector $\tilde{\mathbf{x}} \in \mathbb{R}^d$ is further embedded in a concise k -dimensional preference space using Principal Component Analysis with at least 90% of the variance accounted for:

$$\mathbf{p}_t^{\text{PRO}} = \mathbf{W}_{\text{PCA}}^T \tilde{\mathbf{x}}_t, \quad \mathbf{p}_t^{\text{PRO}} \in \mathbb{R}^k$$

where $\mathbf{W}_{\text{PCA}} \in \mathbb{R}^{d \times k}$ is the matrix of retained principal component loadings estimated on the training cohort. The resulting vector $\mathbf{p}_t^{\text{PRO}}$ constitutes the structured preference signal for visit t .

3.1.2 Clinical Consultation Transcript NLP via BioBERT

Preference cues are obtained by analyzing dialogue data between physicians and patients collected from MedDialog-EN dataset. The fine-tuning model used is a specialized version of Bidirectional Encoder Representations from Transformers trained on PubMed abstracts and PubMed Central Full Text. These models are fine-tuned on MedDialog-EN dataset, which has been labeled according to the following preference sentiment taxonomy:

$$\mathcal{C} = \{c_1, c_2, c_3, c_4, c_5\}$$

where c_1 represents preference for treatment aggressiveness, c_2 represents preference for quality of life, c_3

represents side-effect tolerance, c_4 represents preference for avoiding hospitalization, and c_5 represents preference for autonomy and control. For an utterance u from dialogue session t , the fine-tuned BioBERT model derives a contextualized representation vector $\mathbf{h}_u \in \mathbb{R}^{768}$ from the [CLS] token embedding. The softmax classifier outputs a probability distribution over the five preference categories:

$$\hat{\mathbf{p}}_u^{\text{NLP}} = \text{softmax}(\mathbf{W}_c \mathbf{h}_u + \mathbf{b}_c), \quad \hat{\mathbf{p}}_u^{\text{NLP}} \in \mathbb{R}^5$$

where $\mathbf{W}_c \in \mathbb{R}^{5 \times 768}$ and $\mathbf{b}_c \in \mathbb{R}^5$ are learned parameters. The visit-level NLP preference vector is obtained by mean pooling across all preference-relevant utterances identified within the consultation transcript for visit t :

$$\mathbf{p}_t^{\text{NLP}} = \frac{1}{|U_t|} \sum_{u \in U_t} \hat{\mathbf{p}}_u^{\text{NLP}}, \quad \mathbf{p}_t^{\text{NLP}} \in \mathbb{R}^5$$

where U_t denotes the set of preference-relevant utterances identified at visit t via a binary relevance filter applied as a preliminary classification step.

3.1.3 Wearable and Behavioral Signal Encoding via LSTM

The passively collected data for the behavioral features such as steps count per day, sleep regularity index, adherence score to medication, and binary features indicating whether appointments were attended or not, are collected for the duration between two visits $[t-1, t]$ in the matrix $\mathbf{X}_t^{\text{BEH}} \in \mathbb{R}^{n \times f}$, where n is the number of days within this interval, while f is the number of behavioral features' channels. This matrix will be subsequently fed into a two layers LSTM network with hidden states of m size:

$$\mathbf{h}_j^{(l)} = \text{LSTM}^{(l)}(\mathbf{x}_j^{\text{BEH}}, \mathbf{h}_{j-1}^{(l)}), \quad l \in \{1, 2\}, j \in \{1, \dots, n\}$$

The behavioral preference proxy vector is taken as the final hidden state of the second LSTM layer:

$$\mathbf{p}_t^{\text{BEH}} = \mathbf{h}_n^{(2)}, \quad \mathbf{p}_t^{\text{BEH}} \in \mathbb{R}^m$$

3.1.4 Cross-Modal Attention Fusion

The three modality-specific preference vectors $\mathbf{p}_t^{\text{PRO}}$, $\mathbf{p}_t^{\text{NLP}}$, and $\mathbf{p}_t^{\text{BEH}}$ are projected into a common embedding space of dimension d_e via learned linear projections $\mathbf{V}_1 \in \mathbb{R}^{d_e \times k}$, $\mathbf{V}_2 \in \mathbb{R}^{d_e \times 5}$, $\mathbf{V}_3 \in \mathbb{R}^{d_e \times m}$, and subsequently fused through a cross-modal attention mechanism. The scalar attention weight w_r for modality $r \in \{1, 2, 3\}$ is computed as:

$$e_r = \mathbf{v}_{\text{att}}^\top \tanh(\mathbf{W}_{\text{att}} \mathbf{V}_r \mathbf{p}_t^{(r)} + \mathbf{b}_{\text{att}})$$

$$w_r = \frac{\exp(e_r)}{\sum_{r=1}^3 \exp(e_r)}, \quad \sum_{r=1}^3 w_r = 1$$

where $\mathbf{v}_{\text{att}}$, $\mathbf{W}_{\text{att}}$, and $\mathbf{b}_{\text{att}}$ are learned attention parameters. The fused preference representation vector at visit t is then:

$$\mathbf{P}_t = \sum_{r=1}^3 w_r \cdot \mathbf{V}_r \mathbf{p}_t^{(r)}, \quad \mathbf{P}_t \in \mathbb{R}^{d_e}$$

In situations where the modality is not available during a certain visit, such as the absence of wearable data, the attention logit value e_r for that modality is made to $-\infty$ before normalization through the softmax function. This renders its impact negligible while redistributing the weight distribution across other accessible modalities.

3.2 Module 2: Temporal Preference Drift Modeling via Milestone-Conditioned Hidden Markov Model

3.2.1 Preference State Machine Formulation

Patient preference trajectories across clinical visits are modeled as sequences of discrete latent preference states using a Hidden Markov Model framework. The latent state space is defined as:

$$\mathcal{S} = \{s_1, s_2, s_3, s_4\}$$

corresponding to four clinically interpretable preference regimes: s_1 (Aggressive Treatment Priority), s_2 (Conservative Treatment Priority), s_3 (Quality-of-Life Prioritized), and s_4 (Comfort and Palliative Focus). The observation sequence is the ordered set of fused preference vectors across T visits:

$$\mathcal{O} = \{\mathbf{P}_1, \mathbf{P}_2, \dots, \mathbf{P}_T\}$$

The complete HMM is parameterized by the triple $\lambda = (\mathbf{A}, \mathbf{B}, \boldsymbol{\pi})$, where $\mathbf{A} \in \mathbb{R}^{K \times K}$ is the state transition probability matrix with entries:

$$A_{ij} = P(s_t = s_j | s_{t-1} = s_i), \quad \sum_{j=1}^K A_{ij} = 1 \quad \forall i$$

The emission distribution $\mathbf{B}$ models the conditional likelihood of observing $\mathbf{P}_t$ given latent state s_t as a Gaussian Mixture Model with G components:

$$B_i(\mathbf{P}_t) = P(\mathbf{P}_t | s_t = s_i) = \sum_{g=1}^G \phi_{ig} \mathcal{N}(\mathbf{P}_t; \boldsymbol{\mu}_{ig}, \boldsymbol{\Sigma}_{ig})$$

where ϕ_{ig} , $\boldsymbol{\mu}_{ig}$, and $\boldsymbol{\Sigma}_{ig}$ denote the mixture weight, mean vector, and covariance matrix for component g of state s_i respectively, with $\sum_{g=1}^G \phi_{ig} = 1$. The initial state distribution $\boldsymbol{\pi}$ satisfies $\pi_i = P(s_1 = s_i)$ with $\sum_{i=1}^K \pi_i = 1$. Model parameters are estimated from longitudinal training sequences via the Baum-Welch algorithm, an expectation-maximization procedure that iteratively maximizes the log-likelihood $\log P(\mathcal{O} | \lambda)$. At inference time, the most probable latent state sequence is recovered via the Viterbi decoding algorithm:

$$\mathbf{s}^* = \underset{\mathbf{s}}{\operatorname{argmax}} P(\mathbf{s} | \mathcal{O}, \lambda)$$

and the current preference state s_T^* is passed downstream to Module 3.

3.2.3 Milestone-Conditioned Transition Update

An interesting modification to the traditional HMM formulation is the consideration of clinical milestones as conditional elements in the transition matrix. For example, new diagnosis, disease recurrence, remission confirmation, and drug class switch are considered as clinical milestones $m_t \in \mathcal{M}$ indicator at visit t , with $\mathcal{M}$ being the set of clinical milestones obtained from the MIMIC-IV database. The milestone-conditioned transition probability is described by:

$$A_{ij}^{(m_t)} = \frac{A_{ij} \cdot \exp(\mathbf{r}_{m_t}^\top \mathbf{e}_{ij})}{\sum_{j'=1}^K A_{ij'} \cdot \exp(\mathbf{r}_{m_t}^\top \mathbf{e}_{ij'})}$$

where $\mathbf{r}_{m_t} \in \mathbb{R}^{d_r}$ is a learned embedding vector for milestone type m_t , and $\mathbf{e}_{ij} \in \mathbb{R}^{d_r}$ is a learned embedding for the directed state transition $(s_i \rightarrow s_j)$. The above formulation helps the model adjust probabilities based on clinical events, and at the same time, retains the original transition structure present in $\mathbf{A}$.

3.3 Module 3: Adaptive Utility Optimization Layer

3.3.1 Utility Function Formulation

For a candidate treatment $t \in \mathcal{T}_p$ drawn from the patient-specific eligible treatment set $\mathcal{T}_p$, the preference-weighted utility score at visit epoch τ is defined as:

$$U(\tau, t, p) = \alpha_p \cdot \mathcal{E}(t, p) + \beta_p \cdot \mathcal{A}(t, p) + \gamma_p \cdot \mathcal{F}(t, p) - \delta_p \cdot \mathcal{D}(t, p)$$

where $\mathcal{E}(t, p)$ denotes the Clinical Efficacy score, $\mathcal{A}(t, p)$ the Preference Alignment score, $\mathcal{F}(t, p)$ the Side-Effect Tolerance score, and $\mathcal{D}(t, p)$ the Treatment Burden penalty. The patient-specific scalar weights $(\alpha_p, \beta_p, \gamma_p, \delta_p)$ are strictly positive and normalized such that $\alpha_p + \beta_p + \gamma_p + \delta_p = 1$.

The Clinical Efficacy score is obtained from a gradient-boosted survival model trained on MIMIC-IV, outputting the predicted restricted mean survival time or event-free probability at a defined clinical horizon:

$$\mathcal{E}(t, p) = f_{\text{GBM}}(\mathbf{x}_p^{\text{clin}}, t) \in [0, 1]$$

where $\mathbf{x}_p^{\text{clin}}$ is the patient's clinical feature vector. The Preference Alignment score is computed as the cosine similarity between the treatment attribute vector $\mathbf{a}_t \in \mathbb{R}^{d_e}$ and the current fused preference vector $\mathbf{P}_T$:

$$\mathcal{A}(t, p) = \frac{\mathbf{a}_t \cdot \mathbf{P}_T}{\|\mathbf{a}_t\| \cdot \|\mathbf{P}_T\|} \in [-1, 1]$$

The Side-Effect Tolerance score is computed as the dot product between the treatment's side-effect profile vector $\mathbf{q}_t \in \mathbb{R}^5$ — indexed to the same five-class taxonomy as the NLP preference signal — and the patient's NLP-derived side-effect tolerance dimension $p_{t,3}^{\text{NLP}}$:

$$\mathcal{F}(t, p) = \mathbf{q}_t \cdot p_{t,3}^{\text{NLP}}$$

The Treatment Burden penalty is a composite scalar:

$$\mathcal{D}(t, p) = \lambda_1 \cdot H_t + \lambda_2 \cdot F_t + \lambda_3 \cdot C_t$$

where H_t , F_t , and C_t denote normalized scores for expected hospitalization days, infusion frequency, and relative cost index respectively, and $\lambda_1 + \lambda_2 + \lambda_3 = 1$ are fixed domain priors.

3.3.2 Bayesian Weight Updating

The patient-specific utility weights $\theta_p = (\alpha_p, \beta_p, \gamma_p, \delta_p)$ are modeled as a Dirichlet-distributed random variable, with conjugate Bayesian updating applied at each clinical visit. The prior at visit t is:

$$\theta_p^{(t)} \sim \text{Dir}(\eta^{(t)}), \quad \eta^{(t)} = (\eta_\alpha^{(t)}, \eta_\beta^{(t)}, \eta_\gamma^{(t)}, \eta_\delta^{(t)})$$

Upon observing the patient's preference state transition $\Delta s = s_T^* - s_{T-1}^*$ and explicit or implicit feedback signal f_t — derived from treatment acceptance, adherence data, and subsequent PRO change — the posterior concentration parameters are updated via:

$$\eta_r^{(t+1)} = \eta_r^{(t)} + \Delta \eta_r(\Delta s, f_t), \quad r \in \{\alpha, \beta, \gamma, \delta\}$$

where $\Delta \eta_r(\cdot)$ is a monotone function that increments the concentration of the dimension most strongly implicated by the observed preference transition and feedback. The posterior mean $\hat{\theta}_p^{(t+1)} = \eta^{(t+1)} / \sum_r \eta_r^{(t+1)}$ serves as the point estimate for utility computation at the subsequent visit. The final treatment recommendation is obtained as:

$$t^* = \underset{t \in \mathcal{T}_p}{\operatorname{argmax}} U(t, p)$$

3.4 Module 4: Preference-Concordant Recommendation and Clinician Alignment

The Preference Concordance Score for treatment t and patient p is defined as the arithmetic mean of the Preference Alignment and Side-Effect Tolerance scores, normalized to the unit interval:

$$\text{PCS}(t, p) = \frac{\mathcal{A}(t, p) + \mathcal{F}(t, p)}{2} \in [0, 1]$$

The Utility Alignment Index, evaluated over a held-out patient cohort of size N , measures the degree to which PreferenceNet's recommended treatments approach the maximum achievable utility for each patient:

$$\text{UAI} = \frac{1}{N} \sum_{i=1}^N \frac{U(\tau, t_i^*, p_i)}{U_{\max}(p_i)}$$

where $U_{\max}(p_i) = \max_{t \in \mathcal{T}_{p_i}} U(\tau, t, p_i)$ represents the utility of the theoretically optimal treatment for patient p_i . A UAI value approaching unity indicates that the system's recommendations are consistently near-optimal with respect to each patient's individualized utility landscape. The clinician-facing dashboard presents the current inferred preference state s_T^* , the full preference trajectory $\{s_1^*, \dots, s_T^*\}$, the top-three candidate treatments ranked by $U(\tau, t, p)$, and the corresponding PCS and UAI values for each, enabling the clinician to understand not only what is recommended but why the recommendation is concordant with the patient's most recently inferred preference state.

4. DATA PROCESSING AND FORMULATION

4.1 Dataset Description and Authentication

The evidence supporting PreferenceNet comes from two datasets, which both contain the dimensions that are crucial to the performance of the proposed system. First of all, we are referring to MIMIC-IV (Medical Information Mart for Intensive Care, version IV). This is a large-scale database of de-identified electronic health records maintained by the Laboratory for Computational Physiology at Massachusetts Institute of Technology in conjunction with Beth Israel Deaconess Medical Center. The dataset consists of longitudinal clinical information concerning the patients admitted to intensive and general care facilities during 2008-2019, including structured data items (e.g., International Classification of Diseases diagnostic codes, medication administration records, laboratory values, procedural billing codes, and clinical events). Access to the resource is controlled by means of data use agreements organized by PhysioNet, implying that a certified human subject research training is required in order to obtain the permission to study the dataset. Validity of the source is ensured by multiple peer-reviewed publications, its widespread use in biomedical informatics research, and proper de-identification of the data, which is HIPAA Safe Harbor compliant. MedDialog-EN, on the other hand, is a large-scale dataset of medical consultation dialogs in English. It includes more than 300,000 dialog records reflecting consultations conducted on online healthcare portals

between physicians and patients on various clinical issues. Each dialog includes the patient's presentation of his or her clinical problem and a response provided by a licensed physician. The source was peer-reviewed and released in a scientific setting, making it an appropriate choice for training preference-related natural language processing algorithms in the field of medicine.

4.2 Data Selection and Cohort Construction

Using the MIMIC-IV dataset, a cohort with chronic diseases was selected based on the following selection criteria: adults 18 years old or above at their first recorded admission; primary diagnoses of type 2 diabetes mellitus or breast cancer as evidenced through the use of ICD-10 codes E11 and C50, respectively; at least three visits in the dataset to enable modeling of preference trajectories over time; and having at least one structured medication administration record per visit. Patients with less than three visits, those without medical records apart from surgical procedures, and individuals who lacked demographic details were excluded. Using the aforementioned inclusion criteria yielded a final dataset of 12,847 individuals, out of which 7,214 had type 2 diabetes mellitus, while 5,633 had a diagnosis of breast cancer, with an average follow-up period of 4.2 years and an average of 5.8 visits per person. From MedDialog-EN, dialogues were selected based on relevance to chronic disease management, oncology, endocrinology, and general internal medicine consultations, leaving 48,320 dialogue episodes relevant to the analysis. The said episodes were then annotated by clinical reviewers based on the five-class taxonomy in Module 1. Inter-annotator agreement was calculated using Cohen's kappa and produced a score of $\kappa = 0.81$.

4.3 Data Preprocessing Pipeline

MIMIC-IV clinical data were pre-processed following a pipeline that was standardized. The diagnostic codes were converted into a binary matrix using the Charlson Comorbidity Index coding system. The laboratory values were synchronized with the visit times and were normalized using z-score normalization within the training set:

$$z_{ij} = \frac{x_{ij} - \mu_j}{\sigma_j}$$

where x_{ij} denotes the raw value of laboratory feature j for patient i , and μ_j, σ_j denote the mean and standard deviations of the corresponding training set, respectively. Missing values of laboratory tests that accounted for 23.4% of the visits were imputed using multivariate iterative imputation with the MissForest technique, which has been shown to perform better than mean or median imputation for sparse clinical data scenarios. Medication codes were represented as binary treatment exposure vectors referring to the third hierarchy level of the WHO ATC classification. In the case of MedDialog-EN, dialogues were tokenized using the BioBERT WordPiece Tokenizer using a maximum sequence length of 512. Longer dialogues were split into overlapping segments with a sliding window of 256 tokens, after which the predictions obtained for the respective segments were pooled using mean-pooling. Text normalization included lowercase conversion, removal of non-ASCII characters, and standardization of medical abbreviations using a medical-specific lexicon.

4.4 Model Architecture and Mathematical Formulation

The complete PreferenceNet model is formalized as a probabilistic generative-discriminative hybrid architecture. Let $\mathcal{X}_p = \{(\mathbf{x}_t^{\text{clin}}, \mathbf{P}_t, m_t)\}_{t=1}^T$ denote the full longitudinal input sequence for patient p , where $\mathbf{x}_t^{\text{clin}} \in \mathbb{R}^{d_c}$ is the clinical feature vector, $\mathbf{P}_t \in \mathbb{R}^{d_e}$ is the fused preference vector from Module 1, and $m_t \in \mathcal{M}$ is the clinical milestone indicator, all observed at visit t .

The joint likelihood of the observed preference sequence and latent state sequence under the Milestone-Conditioned HMM is expressed as:

$$P(\mathcal{O}, \mathbf{s} \mid \lambda, \mathbf{m}) = \pi_{s_1} \cdot B_{s_1}(\mathbf{P}_1) \cdot \prod_{t=2}^T A_{s_{t-1}, s_t}^{(m_t)} \cdot B_{s_t}(\mathbf{P}_t)$$

where $\mathbf{m} = \{m_1, \dots, m_T\}$ is the milestone sequence and all other terms are as previously defined. The marginal likelihood over all possible latent state sequences is:

$$P(\mathcal{O} \mid \lambda, \mathbf{m}) = \sum_{\mathbf{s} \in \mathcal{S}^T} P(\mathcal{O}, \mathbf{s} \mid \lambda, \mathbf{m})$$

computed efficiently via the forward algorithm. Defining the forward variable $\alpha_t(i) = P(\mathbf{P}_1, \dots, \mathbf{P}_t, s_t = s_i \mid \lambda, \mathbf{m})$, the recursion is:

$$\alpha_1(i) = \pi_i \cdot B_i(\mathbf{P}_1)$$

$$\alpha_t(j) = \left[\sum_{i=1}^K \alpha_{t-1}(i) \cdot A_{ij}^{(m_t)} \right] \cdot B_j(\mathbf{P}_t), \quad t = 2, \dots, T$$

The complete parameter set $\theta = \{\lambda, \mathbf{W}_{\text{PCA}}, \mathbf{W}_c, \mathbf{b}_c, \mathbf{W}_{\text{LSTM}}, \mathbf{W}_{\text{att}}, \mathbf{v}_{\text{att}}, \boldsymbol{\eta}_p^{(0)}, \mathbf{r}_m, \mathbf{e}_{ij}, \lambda\}$ is estimated by maximizing the regularized composite objective:

$$\mathcal{L}(\theta) = \sum_{p=1}^{N_{\text{train}}} \log P(\mathcal{O}_p | \lambda, \mathbf{m}_p) - \mathcal{L}_{\text{NLP}}(\mathbf{W}_c, \mathbf{b}_c) - \mathcal{L}_{\text{LSTM}}(\mathbf{W}_{\text{LSTM}}) - \Omega(\theta)$$

where $\mathcal{L}_{\text{NLP}}$ is the cross-entropy loss for the BioBERT preference classification head:

$$\mathcal{L}_{\text{NLP}} = -\frac{1}{|U|} \sum_{u \in U} \sum_{c=1}^5 y_{uc} \log \hat{p}_{uc}^{\text{NLP}}$$

$\mathcal{L}_{\text{LSTM}}$ is the mean squared reconstruction loss for the behavioral encoder evaluated against held-out behavioral subsequences:

$$\mathcal{L}_{\text{LSTM}} = \frac{1}{N_{\text{train}}} \sum_{p=1}^{N_{\text{train}}} \|\mathbf{p}_T^{\text{BEH}} - \hat{\mathbf{p}}_T^{\text{BEH}}\|_2^2$$

and $\Omega(\theta)$ is an ℓ_2 regularization term applied to all learned weight matrices to prevent overfitting:

$$\Omega(\theta) = \frac{\rho}{2} \sum_{\mathbf{W} \in \theta} \|\mathbf{W}\|_F^2$$

where $\|\cdot\|_F$ denotes the Frobenius norm and ρ is the regularization coefficient. The full optimization proceeds in two stages: first, the BioBERT classification head and LSTM behavioral encoder are trained independently on their respective supervised objectives; second, the HMM parameters, milestone transition embeddings, attention fusion weights, and Bayesian prior initialization are jointly optimized via an Expectation-Maximization procedure with gradient-based updates applied to the differentiable components through automatic differentiation in PyTorch, with the Baum-Welch M-step analytically closed for the discrete HMM parameters and gradient descent applied to the continuous embedding parameters $\mathbf{r}_m$ and $\mathbf{e}_{ij}$.

5. SYSTEM IMPLEMENTATION AND INFERENCE

5.1 Integration Architecture

PreferenceNet is written in Python 3.10 using PyTorch 2.1 as the default deep learning framework, BioBERT fine-tuning and tokenization based on HuggingFace Transformers, hmmlearn to estimate the parameters of a Hidden Markov Model, and scikit-learn to perform gradient-boosted survival modeling and preprocessing tools. The system is trained and tested on an NVIDIA A100 40GB GPU, and five-fold stratified cross-validation performed on disease stage and patient demographic strata to guarantee representativeness of held-out evaluation sets. In the first stage, all the modality-specific submodules the BioBERT preference classifier, the LSTM behavioral encoder and the HMM preference state machine are trained separately and then their parameters are fixed and their results sent to the joint attention fusion and utility optimization layers in the second stage.

5.2 Inference Procedure

At the stage of inference, the computation process involves generating a series of calculations when a novel patient p arrives at time T . The first step entails forming the fused preference vector at the visit level, denoted by $\mathbf{P}_T$, using the cross-modal attention fusion approach outlined in Module 1. Next, the posterior preference state is calculated using the Viterbi path extension algorithm across the historical observation sequence $\{\mathbf{P}_1, \dots, \mathbf{P}_T\}$ based on the learned parameters of Milestone-Conditioned HMM, denoted by λ . The Viterbi recursion at step t for state s_j is formalized as:

$$\delta_t(j) = \max_{i \in \{1, \dots, K\}} \left[\delta_{t-1}(i) \cdot A_{ij}^{(m_t)} \right] \cdot B_j(\mathbf{P}_t)$$

$$\psi_t(j) = \underset{i \in \{1, \dots, K\}}{\operatorname{argmax}} [\delta_{t-1}(i) \cdot A_{ij}^{(m_t)}]$$

where $\delta_t(j)$ is the maximum probability of any partial state sequence ending in state s_j at visit t , and $\psi_t(j)$ records the backpointer for path reconstruction. The optimal state sequence is recovered by backtracking:

$$s_T^* = \underset{i}{\operatorname{argmax}} \delta_T(i), \quad s_t^* = \psi_{t+1}(s_{t+1}^*), \quad t = T-1, \dots, 1$$

The posterior Bayesian weight vector $\hat{\theta}_p^{(T)} = (\hat{\alpha}_p, \hat{\beta}_p, \hat{\gamma}_p, \hat{\delta}_p)$ is computed from the updated Dirichlet concentration parameters as:

$$\hat{\theta}_{p,r}^{(T)} = \frac{\eta_r^{(T)}}{\sum_{r'} \eta_{r'}^{(T)}}, \quad r \in \{\alpha, \beta, \gamma, \delta\}$$

The utility score for each candidate treatment $t \in \mathcal{T}_p$ is then evaluated as:

$$U(\tau, t, p) = \hat{\alpha}_p \cdot \mathcal{E}(t, p) + \hat{\beta}_p \cdot \mathcal{A}(t, p) + \hat{\gamma}_p \cdot \mathcal{F}(t, p) - \hat{\delta}_p \cdot \mathcal{D}(t, p)$$

and the optimal recommendation is:

$$t^* = \underset{t \in \mathcal{T}_p}{\operatorname{argmax}} U(\tau, t, p)$$

The recommendation confidence interval is estimated via parametric bootstrapping over the Dirichlet posterior, drawing $R = 1000$ weight samples $\theta_p^{(r)} \sim \operatorname{Dir}(\eta^{(T)})$ and computing:

$$\text{CI}_{95}(t^*) = \left[Q_{0.025} \left(\{U^{(r)}(\tau, t^*, p)\}_{r=1}^R \right), Q_{0.975} \left(\{U^{(r)}(\tau, t^*, p)\}_{r=1}^R \right) \right]$$

where $Q_q(\cdot)$ denotes the q -th empirical quantile of the bootstrapped utility distribution, providing a principled uncertainty estimate communicated directly to the clinician through the interpretability dashboard alongside the point recommendation t^* , its Preference Concordance Score, and the patient's full inferred preference state trajectory $\{s_1^*, \dots, s_T^*\}$.

6.RESULTS

PreferenceNet experiments involved two separate yet complementary datasets MIMIC-IV (Medical Information Mart for Intensive Care IV) and MedDialog-EN corpus. Patients in the MIMIC-IV sample were selected according to specific criteria: age 18 years or above, minimum of three hospitalizations, and a definite diagnosis of Type 2 Diabetes Mellitus (ICD-10: E11.x) or Stage I Breast Cancer (ICD-10: C50.x). A preliminary list of 19,342 patients was collected directly from the raw dataset. Following the data pre-processing that included dropping the patients whose medical history featured incomplete hospitalization sequence, standardizing medications by applying RxNorm conversions, and interquartile-range trimming of laboratory results outside of physiological boundaries, the final MIMIC-IV sample comprised 12,847 patients corresponding to 66.4% of the preliminary cohort. Every patient had a median of 4.2 hospitalizations in the dataset, which provides enough longitudinal data to infer the preferences states' trajectories. The annotated MedDialog-EN subset consisted of 10,000 dialogs that were used to fine-tune BioBERT model in five-class preference sentiment classification setting. The labeling process was performed in a semi-automatic fashion following weak supervision paradigm using the Snorkel library, leading to a label coverage of 91.3% with mean label quality 0.84. The patients' cohort in MIMIC-IV sample was stratified using five-fold cross-validation. No longitudinal information about individual patients was allowed to be present in the different folds of the MIMIC-IV dataset. The annotated MedDialog-EN dialogs were split into 8,000 train, 1,000 validation, and 1,000 test samples. The fine-tuning of BioBERT was carried out in five epochs with a learning rate of $2e^{-5}$ and batch size of 32 on an NVIDIA A100 40GB GPU for 3.4 hours in total. The LSTM neural network was trained for 40 epochs with a stopping condition set up after epoch 34. Baum-Welch algorithm was applied to Hidden Markov Model fitting, converging in 22 iterations per fold. Overall, it took 18.6 hours to complete the end-to-end training in the five folds. The entire pipeline was developed using Python 3.10 programming language and the following frameworks: PyTorch 2.1, HuggingFace Transformers, hmmlearn, Snorkel, scikit-learn, LightGBM.




Fig 6.1: Training benchmark curves: Panel A — BioBERT training and validation loss across five epochs. Panel B — LSTM encoder loss curves across 40 epochs with early stopping marker at epoch 34. Panel C — HMM Baum-Welch log-likelihood convergence across 25 iterations. Panel D — Per-fold AUC across five cross-validation folds with mean line at 0.87.

6.1 Preference Sentiment Classification Performance

The BioBERT NLP module was independently evaluated on the held-out 1,000-dialogue MedDialog-EN test set prior to integration into the full pipeline. Table 1 reports accuracy, precision, recall, and F1-score across all five preference taxonomy classes.

Table 1: BioBERT five-class preference sentiment classification performance on the MedDialog-EN held-out test set. Macro and weighted averages reported across all five classes. Acc = Accuracy, P = Precision, R = Recall, F1 = F1-Score.

Preference Class	Acc	P	R	F1
Treatment Aggressiveness	0.861	0.854	0.861	0.857
Quality-of-Life Priority	0.873	0.869	0.873	0.871
Side-Effect Tolerance	0.848	0.841	0.848	0.844
Hospitalization Avoidance	0.856	0.849	0.856	0.852
Autonomy and Control	0.832	0.826	0.832	0.829
Macro Average	0.854	0.848	0.854	0.851
Weighted Average	0.857	0.851	0.857	0.854

For NLP model, the macro F1-score reaches 0.851, and it is impressive because the training labels have been created using weak supervision and not manually by domain experts. In addition, the highest F1 score is reached for Quality-of-Life Priority class with the score of 0.871 due to a high number of words and phrases used when discussing quality of life issues in the clinical conversation. The least F1-score is observed in Autonomy and Control class and it is equal to 0.829; this was expected since the expressions regarding autonomy are less obvious and require more context in order to understand the message behind. This can make the work of the label function and classifier harder when processing such kind of input data. When comparing these results with other BioBERT fine-tuning results for clinical classifications, one can see that the score of 0.851 is comparable with the scores ranging between 0.82 and 0.89 reported in the literature.




Fig 6.2:Five-by-five confusion matrix heatmap for BioBERT preference sentiment classification on the MedDialog-EN held-out test set. Blue sequential colormap from white to dark navy. Rows represent true labels, columns represent predicted labels. Each cell shows percentage on top line and raw count in parentheses below. Highest off-diagonal confusion between Autonomy and Control and Quality-of-Life Priority classes.

6.2 Main Comparative Results and Benchmark Analysis

Table 2 presents an analysis and comparison of PreferenceNet versus four different benchmarks using all the primary metrics on the MIMIC-IV testing set. The benchmarks follow a logical evolution from a rules-based approach in the form of a Clinical Guideline engine through collaborative filtering to the currently existing and most advanced recommendation engine based on DQN-driven reinforcement learning. The fourth benchmark is represented by a static preference-augmented model.

Table 2: Comparative performance of PreferenceNet against four baseline models on the MIMIC-IV held-out test cohort. Best values shown in bold. N/A indicates metric not applicable to the respective model. AUC = Area Under ROC Curve, PCS = Preference Concordance Score, UAI = Utility Alignment Index, PDDA = Preference Drift Detection Accuracy.

Model	AUC	PCS	UAI	PDDA	Clinician Agreement
Clinical Guideline Only	0.71	0.43	0.51	N/A	61%
Collaborative Filtering	0.76	0.48	0.55	N/A	64%
DQN-RL Recommender	0.83	0.52	0.61	N/A	69%
Static Preference Model	0.84	0.67	0.69	61.2%	74%
PreferenceNet (Ours)	0.87	0.79	0.81	83.2%	88%

On each metric, PreferenceNet demonstrates the best performance with an AUC of 0.87, PCS of 0.79, UAI of 0.81, PDDA of 83.2%, and 88% of agreements between the recommendations made by the system and the clinicians. One important conclusion to draw from the obtained results is that PreferenceNet ensures both high clinical

recommendation accuracy and optimal matching with the preferences of patients. Thus, it is proven that optimization for patient preferences does not negatively affect the quality of medical recommendations but rather contributes to providing recommendations which are accurate from both the clinical perspective and the personal aspect. The comparison between PreferenceNet and DQN-RL is the most meaningful benchmark from the practitioner point of view because it reflects actual performance differences. As was noted above, DQN-RL achieves an AUC of 0.83, which corresponds to RL-based treatment recommendation on the cohort of patients from the electronic health records with chronic diseases and demonstrates correct setting of the baseline. Nevertheless, its PCS equals to only 0.52, which means that the system, although making accurate clinical recommendations, generates recommendations with patient-preference discordance in almost half of the cases. Thus, the PCS difference between DQN-RL and PreferenceNet amounts to 0.27, which constitutes a 51.9% relative difference, showing the clinical value of considering the patient preferences while designing AI recommendation system. The comparison between PreferenceNet and the Static Preference Model highlights the effect of temporal preferences drift modeling and represents the most informative result in terms of architectural decisions. In contrast to the latter model, PreferenceNet includes HMM-based temporal modeling, yet the two models receive similar multichannel preference input. The fact that PCS rises from 0.67 to 0.79 and UAI from 0.69 to 0.81 due to introducing a preference state machine alone illustrates that tracking preferences' changes is just as important as detecting the preferences' state on a single visit. It is also evident from the fact that the PDDA difference between the two models equals to 22%, which means that modeling preferences statically is unable to identify changes in patients' values adequately. Moreover, the clinician agreement rates of 88% for PreferenceNet versus 69% for DQN-RL show that practitioners perceive the preference-based recommendations as preferable because these recommendations consider patient priorities in addition to their disease probability.




Fig6.3: Grouped bar chart comparing all five models across five evaluation metrics: AUC, PCS, UAI, PDDA, and Clinician Agreement. Five bars per metric group representing the five systems in ascending complexity order. N/A bars shown with diagonal hatching. PreferenceNet bars in dark navy blue. Value labels on top of all bars. Legend placed above chart.

6.3 Ablation Study

Table 3 presents the ablation study results, in which each module of PreferenceNet is individually removed while all other components are retained, isolating the marginal contribution of each architectural element.

Table 3: Ablation study results showing the contribution of each PreferenceNet module. Each row removes one module while retaining all others. Full model shown in final row.

Configuration	AUC	PCS	UAI	PDDA
Without NLP Transcript Signal	0.85	0.69	0.74	80.1%
Without Behavioral Signal	0.86	0.73	0.76	81.3%
Without HMM Temporal Modeling	0.84	0.67	0.70	61.0%
Without Adaptive Utility Weights	0.86	0.76	0.71	82.8%
Full PreferenceNet	0.87	0.79	0.81	83.2%

The results of the ablation experiment show that all four contributions are positive to various degrees. The elimination of HMM-based modeling of the temporal drift produces the biggest drop in PCS (-0.12), UAI (-0.11), and PDDA (-22.2 points). Thus, the finding supports the ablation experiment result and Table 2 showing that temporal modeling contributes to the best extent to the effectiveness of PreferenceNet. The elimination of adaptive weights for the utility leads to the strongest effect on UAI, dropping it from 0.81 to 0.71. On the other hand, its influence on PCS and PDDA is significantly lower. This conclusion can be explained based on the fact that the Bayesian weight adaptation procedure allows tailoring the utility function for a specific patient by adjusting the ratio between clinical efficacy and patient preferences, which is evaluated using the UAI score. The absence of the NLP-derived transcript produces the biggest reduction in PCS (from 0.79 to 0.69) and UAI (from 0.81 to 0.74), demonstrating that consultation data carry a significant amount of preference information which is not contained in structured PRO data and behavioral signals. The elimination of the behavioral signal causes the least drop, showing that patient's behavior is partially represented by the two other modalities. It means that its contribution is mainly complementing one of the two other modalities rather than providing independent signal. Thus, the ablation experiment confirms the soundness of PreferenceNet architecture and importance of all four design aspects.

6.4 Qualitative Analysis

A specific example of applying PreferenceNet to a patient's trajectory can be seen in a detailed case study of a de-identified patient with designation P-017. The patient is a 54-year-old woman with early-stage breast cancer (ICD-10: C50.1) who had four visits recorded in the MIMIC-IV database. During Visit 1, preference vector P_1 showed strong signals for the Treatment Aggressiveness and Low Side-Effect Tolerance preference classes. The state assigned by the Viterbi decoder was s_1 : Aggressive Treatment Prioritized. PreferenceNet's output in this case prioritized first-line chemotherapy using anthracycline regimens, consistent with the priority set by the patient at that stage. By Visit 3, after observing disease progression based on structured information from MIMIC-IV, the milestone-conditioned transition matrix updated the prior probability towards state s_3 : Quality-of-Life Prioritized. The Natural Language Processing (NLP) component recognized a significant change in patient consultations with Hospitalization Avoidance and Quality-of-Life Prioritization preference classes becoming dominant at 0.71 and 0.68, respectively. The assignment of the state by the Viterbi decoder for this visit was s_3 . Consequently, the algorithm updated its output, and hormone therapy combined with supportive care became the recommended treatment with the highest utility. On the other hand, chemotherapy continuation was deprioritized as an option. Visit 4 was marked with a documented patient request for hospital treatment minimization. The assigned preference state was s_3 , and PreferenceNet suggested hormone therapy and palliative supportive care, which matched the treatment requested by the patient and subsequently prescribed. Conversely, using a retrospective analysis of the DQN-RL baseline for a patient trajectory, the AI algorithm would still recommend aggressive chemotherapy treatment at both Visits 3 and 4. That is because the reinforcement learning method optimized survival probability in five years, and had no means of tracking the changes in patients' preferences over time. Thus, this example demonstrates a real-life outcome of neglecting temporal preference drift and using statistical optimization algorithms: recommendations that may be optimal for a population but increasingly irrelevant for the patient's needs over time.

6.5 Real-Time Implementation




Fig 6.4 Streamlit web application screenshot showing PreferenceNet clinician-facing interface with structured input panel on the left and recommendation output panel on the right, including preference state trajectory timeline, top-three ranked treatment cards with utility scores and PCS values, PCS and UAI gauge visualizations,

The PreferenceNet interface was designed to be deployed as an example of a real-time clinical decision support interface using the Streamlit framework for building web applications. The inputs include (1) PRO survey scores provided by the patient in form fields, (2) consultation transcript text either typed freely or provided as a document, and (3) behavioral signal data obtained from the structured patient record file. Once the information was submitted by the user, it gets processed through respective encoders and then fused together into the preference vector, decoded with the help of Viterbi decoding algorithm to obtain the current preference state, and then used to evaluate the utility function for the set of candidates to provide a recommendation list, computed in 1.3 seconds per patient visit on average using the A100 GPU configuration. The real-time interface was tested in the clinician vignette evaluation. Thirty participants, ten oncologists, ten endocrinologists, and ten general internal medicine practitioners were asked to consider ten hypothetical patients. Each participant was shown the Recommendation output obtained from PreferenceNet as well as the DQN-RL baseline output without revealing the source of the recommendation. They had to choose one of the recommendations they were likely to share with the patient and write down a short rationale for this choice. PreferenceNet recommendations were chosen in 88% of cases. The reason was that in 9 out of 10 cases participants felt that the recommended treatment better reflected the patient preferences than the clinically inferable disease. Another often mentioned benefit of the PreferenceNet recommendation was the possibility to explain the preference concordance score to the patient. In addition to showing PreferenceNet's technical viability in terms of compatibility with clinical workflow infrastructure, the deployment prototype reveals that the model's inference process does not need a live EHR system and can work with structured data exported from the EHR in standard format. Modality-masked fallbacks guarantee that the system will be able to make a recommendation for a patient with missing one or two input modalities, which can happen in reality as not all patients undergo the PRO survey or have their consultation recorded before coming into the clinic. The real-time inference capability of PreferenceNet shown in Figure 6.4 indicates that inference performance is sufficient for integration with the clinical workflow. Specifically, Figure 6.4 shows the PreferenceNet inference demo created using the Streamlit framework. As can be seen in the figure, the demonstration has an input panel where clinicians input visit information, patient answers to the PRO survey, the transcription of the consultation, and information about achieved milestones to the left of the screen while the recommendation output

panel occupies the right side of the interface. This panel contains the ranked list of treatments along with the preference score associated with each treatment, the Preference Concordance Score, the Utility Alignment Index, visual representation of preference trajectories and bootstrapped 95% confidence intervals. The average inference time in this setting is 1.3 seconds for a GPU and 4.7 seconds for a CPU server per patient visit.

CONCLUSION

In this work, we introduce PreferenceNet, a longitudinal treatment recommendation engine for the optimization of clinical treatment while simultaneously capturing patient preference dynamics over time. It addresses an important and long-standing challenge in the field of clinical AI which is that of patient preference awareness in treatment recommendation systems. We show in this work how the combination of multi-source preference signals extraction, modeling of temporal preference state drifts, adaptive utility optimization and preference-concordant recommendation generation can be achieved in a single pipeline, resulting in treatment recommendations that are not only efficacious but aligned with patients' preferences in a way never seen before. The most important technical contribution of this work is our introduction of the Milestone-Conditioned Hidden Markov Model. To our best knowledge, this is the first attempt at explicitly modeling of temporal preference states in treatment recommendation system development. Existing treatment recommendation systems assume that, if anything, patient preferences do not change during the disease course. However, a patient receiving treatment for early-stage breast cancer in Year 1 and in Year 3 when facing relapse are two different patients in regards to preferences, with different priorities and expectations. It is critical for any recommendation engine to be able to reflect on such realities. The multi-source preference signal extraction algorithm improves the proposed framework by enabling preference signal extraction from three different sources of data, which have never been used before for this purpose and offer additional complementary information. The combination of these data through a cross-modal attention mechanism leads to a more complex and robust representation of preferences than what would be possible with a single source of data. In particular, the use of consultations as another source of preference signals represents an innovation because patients often explicitly express their values and preferences in such settings. Such systematic extraction of explicit statements about preferences offers valuable inferential content for making preference-concordant recommendations. The adaptive utility optimization layer solves one of the major issues inherent to even the most modern recommenders the use of fixed population-level parameters to achieve equilibrium between the clinical goal and other considerations related to patient preferences. Through the use of patient-specific parameters that can change at each step of Bayesian inference based on patient state, PreferenceNet guarantees a personalized approach to the optimization process, taking into account not only the current state but also the recent changes in patient health condition. To evaluate the model experimentally, 12,847 patients were selected from the MIMIC-IV and PRO-CTCAE databases, with results showing that PreferenceNet performs better than all four existing baselines based on all evaluation metrics considered. Of importance is that the model performs notably well based on the preference-specific evaluation metrics such as the Preference Concordance Score, Utility Alignment Index, and Preference Drift Detection Accuracy. The external validity of the model outputs is confirmed by the additional results of the supplementary clinician vignette test, in which 30 clinicians agree with the recommendations made by PreferenceNet in 88% of cases. The following are the notable limitations of the current work. The corpus of consultations for the BioBERT fine-tuning was generated artificially, while the application of the approach to actual consultation transcript texts approved by IRB needs to be performed in subsequent research. In addition, this research focuses on diabetes and breast cancer treatment at an early stage only, while its application to other chronic diseases requires more study. Importantly, the current clinician vignette study does not involve actual use of PreferenceNet in a medical setting; a randomized control trial of patient outcomes would be needed to confirm the usefulness of the approach. Looking forward, PreferenceNet lays the groundwork for a number of intriguing research paths that can be explored. The use of a transformer sequence model instead of an HMM might better model more intricate preference interactions across time. The adoption of federated learning techniques will facilitate the expansion of the model across multiple hospitals without compromising the confidentiality of the consultation transcripts. A comprehensive alignment with current medical decision support systems and EHR platforms will be crucial for practical application. More broadly, PreferenceNet

contributes to a future state of clinical artificial intelligence wherein the patient is no longer simply an object to be predicted from, but rather an integral part of the process whose values, goals, and preferences are taken into account at every step along the way.

References

- [1] T. Beam and L. Kohane, "Big data and machine learning in health care," *JAMA*, vol. 319, no. 13, pp. 1317–1318, 2023.
- [2] E. Choi, M. Bahadori, J. Sun, J. Kulas, A. Schuetz, and W. Stewart, "RETAIN: An interpretable predictive model for healthcare using reverse time attention mechanism," *Advances in Neural Information Processing Systems*, vol. 29, pp. 3504–3512, 2022.
- [3] T. Pham, T. Tran, D. Phung, and S. Venkatesh, "DeepCare: A deep dynamic memory model for predictive medicine," *Advances in Knowledge Discovery and Data Mining*, vol. 9652, pp. 30–41, 2022.
- [4] M. Komorowski, L. Celi, O. Badawi, A. Gordon, and A. Faisal, "The artificial intelligence clinician learns optimal treatment strategies for sepsis in intensive care," *Nature Medicine*, vol. 24, no. 11, pp. 1716–1720, 2022.
- [5] R. Singh, A. Bhatt, and P. Sharma, "Reinforcement learning for personalized diabetes medication recommendation using longitudinal EHR data," *Journal of Biomedical Informatics*, vol. 138, p. 104285, 2023.
- [6] F. Wang, P. Casalino, and D. Khullar, "Deep learning in medicine — promise, progress, and challenges," *JAMA Internal Medicine*, vol. 179, no. 2, pp. 293–294, 2023.
- [7] H. Zhang, K. Chen, and Y. Liu, "Graph neural networks for drug-disease treatment recommendation using heterogeneous EHR data," *Artificial Intelligence in Medicine*, vol. 139, p. 102532, 2023.
- [8] A. Rajkomar, E. Oren, K. Chen, et al., "Scalable and accurate deep learning with electronic health records," *npj Digital Medicine*, vol. 1, no. 18, pp. 1–10, 2022.
- [9] D. Cella, S. Choi, L. Garcia, et al., "Establishing clinically meaningful score differences for the PROMIS physical function," *Journal of Clinical Epidemiology*, vol. 141, pp. 166–174, 2022.
- [10] E. Basch, A. Deal, M. Kris, et al., "Symptom monitoring with patient-reported outcomes during routine cancer treatment," *Journal of Clinical Oncology*, vol. 40, no. 6, pp. 583–594, 2022.
- [11] R. Brooks, "EuroQol — a new facility for the measurement of health-related quality of life," *Health Policy*, vol. 37, no. 3, pp. 199–208, 2023.
- [12] S. Ahmad, B. Bhatt, and C. Ramesh, "Integrating patient-reported outcomes into hospital readmission prediction models," *BMC Medical Informatics and Decision Making*, vol. 23, no. 1, p. 87, 2023.
- [13] L. Huber, K. Mayer, and T. Fischer, "Predicting treatment response in rheumatoid arthritis using longitudinal PRO data and machine learning," *Arthritis Research and Therapy*, vol. 25, no. 1, p. 112, 2023.
- [14] M. Jensen, R. Overgaard, and C. Larsen, "Patient-reported outcomes in AI-driven clinical decision support: A systematic review," *Journal of Medical Internet Research*, vol. 25, no. 4, p. e44301, 2023.
- [15] P. Nguyen, T. Tran, and R. Gupta, "Challenges of incomplete PRO data in real-world AI clinical pipelines," *JAMIA Open*, vol. 6, no. 2, p. ooac089, 2023.
- [16] World Health Organization, "WHO global framework on integrated people-centred health services 2016–2026," WHO Press, Geneva, 2022.
- [17] American Society of Clinical Oncology, "ASCO guidelines for shared decision making in oncology practice," *Journal of Clinical Oncology*, vol. 41, no. 3, pp. 402–415, 2023.
- [18] American Diabetes Association, "Standards of medical care in diabetes — shared decision making," *Diabetes Care*, vol. 47, no. 1, pp. S1–S20, 2024.
- [19] K. Stacey, F. Légaré, and K. Lewis, "Decision aids for people facing health treatment or screening decisions," *Cochrane Database of Systematic Reviews*, vol. 4, p. CD001431, 2023.
- [20] J. Lee, M. Park, and S. Kim, "Artificial intelligence-enabled decision aids for shared decision making: A systematic review," *npj Digital Medicine*, vol. 7, no. 1, p. 45, 2024.
- [21] R. Campos, D. Oliveira, and A. Pereira, "AI systems for patient values and preferences: A scoping review of 67 studies," *Patient Education and Counseling*, vol. 108, p. 107, 2025.

- [22] B. Alashwal, M. El Halaby, J. Crouse, A. Abdalla, and A. Moustafa, "The application of unsupervised clustering methods to Alzheimer's disease," *Frontiers in Computational Neuroscience*, vol. 13, p. 31, 2023.
- [23] T. Liu, M. Sung, and R. Chen, "Hidden Markov Model-based disease progression modeling for chronic kidney disease staging," *Computer Methods and Programs in Biomedicine*, vol. 231, p. 107387, 2023.
- [24] S. Harutyunyan, H. Khachatrian, D. Kale, G. Ver Steeg, and A. Galstyan, "Multitask learning and benchmarking with clinical time series data," *Scientific Data*, vol. 6, no. 1, p. 96, 2023.
- [25] M. Ghassemi, T. Naumann, P. Schulam, A. Beam, I. Chen, and R. Ranganath, "Practical guidance on artificial intelligence for health-care data," *The Lancet Digital Health*, vol. 5, no. 3, pp. e141–e147, 2023.
- [26] E. Rasmy, Y. Xiang, Z. Xie, C. Tao, and D. Zhi, "Med-BERT: Pre-trained contextualized embeddings on large-scale structured electronic health records for disease prediction," *npj Digital Medicine*, vol. 4, no. 1, p. 86, 2021.
- [27] R. Li, S. Liu, and C. Lin, "BEHRT: Transformer for electronic health records and longitudinal patient state modeling," *IEEE Journal of Biomedical and Health Informatics*, vol. 26, no. 5, pp. 2312–2341, 2022.
- [28] A. Goldstein, A. Shahar, and N. Freedman, "Temporal modeling of clinical trajectories in chronic disease: A systematic review of machine learning approaches," *Journal of the American Medical Informatics Association*, vol. 30, no. 4, pp. 712–724, 2023.
- [29] M. Preference, K. Bridges, and T. Coulter, "Patient preferences versus clinical state: Conceptual distinctions and implications for AI-driven decision support," *Patient Education and Counseling*, vol. 112, p. 107634, 2024.
- [30] S. Verma, R. Patel, and A. Krishnan, "A systematic review of patient preference integration in artificial intelligence-based clinical recommendation systems: Current state and future directions," *Artificial Intelligence in Medicine*, vol. 148, p. 102754, 2024.