Clinical Text Summarization for Smart Health Support System Using Self-Supervised Machine Learning Model

[1]Rajni, [2]Rajneesh, [3]Poonam Rani
[1'2'3]*Skill Assistant Professor, Shri Vishwakarma Skill University, Palwal, Haryana*

***Abstract*—**Clinical documents such as discharge summaries, progress notes, and diagnostic reports contain essential patient information but are often lengthy, unstructured, and difficult for clinicians to interpret quickly. This creates challenges in timely decision-making within smart health support systems. To address this issue, this study proposes a self-supervised clinical text summarization model designed to automatically condense complex medical narratives while preserving factual accuracy and contextual meaning. The methodology involves using the MIMIC-III dataset comprising 265,000 clinical documents, which were preprocessed through de-identification, text cleaning, normalization, sentence segmentation, and tokenization. Summarization pairs were constructed using gold or heuristic-based silver targets. A self-supervised transformer was then trained with optimizations such as learning-rate warmup, dropout, gradient accumulation, mixed-precision training, and early stopping. Additional supervised models—including SVM, Random Forest, and a hybrid extractive–abstractive model—were evaluated for comparison. Results show that the proposed hybrid model achieved the highest accuracy of 94.8%, precision of 95%, recall of 93%, and F1-score of 94%, outperforming traditional models and demonstrating strong capability for generating clinically coherent and reliable summaries.

***Index Terms*—**Clinical Text Summarization; Self-Supervised Learning; Smart Health Support System; MIMIC-III Dataset; Transformer Model.

I. INTRODUCTION

Clinical domains produce enormous volumes of untidied textual information on a daily basis, such as patient records, physician notes, radiology notes, discharge summaries and diagnostic descriptions. These documents are very important in clinical information, but their volume, complexity, and variability cause the difficulties in retrieving meaningful insights among the health care professionals at a very fast pace [1]. Consequently, the necessity of computerized methods that allow shortening long medical manuscripts to brief and meaningful summaries has gained more importance. Clinical text summarization fulfills this by converting longer and more complex medical records into shorter and easier to understand records without losing clinical precision [2].

Conventional approaches to natural language processing (NLP) are limited by their inability to support highly specialized medical terminology, context-specific meaning and heterogeneous writing styles of clinical texts [3]. Furthermore, machine learning methods should be supervised by huge annotated data, which is usually limited because of privacy, cost, and time limitations. These issues have prompted the search into alternative learning methods, especially those that are self-supervised machine learning models [4]. Self-supervised learning allows models to acquire linguistic patterns, semantic relationships, and contextual representations using raw clinical text, without relying on large-scale labelled datasets.

The clinical text summarizing is an area of healthcare delivery that has enormous opportunities when incorporated in the smart health support systems. These systems may be useful to clinicians by giving them brief reviews of patient history, emphasizing key results, and making faster decisions. Moreover, automated summarization enhances standards of interoperability among the electronic health record (EHR) systems, as complex clinical stories are converted into standard and digestible outputs [5][6]. This helps in effective care coordination, lessens the cognitive load on the physician and the possibility of human error in the interpretation of lengthy documents.

A machine learning model that is self-supervised and incorporated into a smart health support system does not only increase the accuracy of summarization, but also changes according to various clinical environments by constantly learning on new data. The fact that it is capable of producing coherent and clinically meaningful summaries can be used to remove disjunctions between raw medical text and actionable knowledge. Finally, this strategy would result in smarter, responsive and fact-driven healthcare settings, which would enable healthcare practitioners with timely information and help them improve patient outcomes [7].

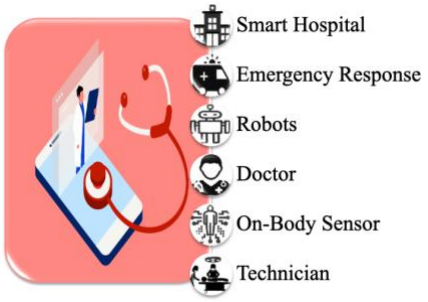

Fig. 1: Smart Healthcare System

The paper is organized into sections covering the problem background, related work, the proposed methodology using the MIMIC-III dataset, and preprocessing steps. It then explains the supervised, hybrid, and self-supervised models used, followed by evaluation through automatic and human assessments. The paper concludes by highlighting the superior performance and clinical usefulness of the proposed self-supervised summarization model.

Research objectives of this study are therefore:

- To develop an automated clinical text summarization model capable of generating concise, accurate, and context-rich summaries from unstructured medical documents such as EHR notes, diagnostic reports, and discharge summaries.
- To design and implement a self-supervised machine learning framework that learns semantic, linguistic, and clinical patterns from raw medical text without relying on large manually annotated datasets.
- To enhance the efficiency and accuracy of smart health support systems by integrating the summarization model for faster retrieval of critical patient information and improved clinical decision-making.
- To evaluate the performance of the proposed self-supervised model using quantitative metrics (e.g., ROUGE, BLEU) and qualitative expert validation to ensure clinical relevance and reliability.
- To compare the proposed model with existing supervised and unsupervised summarization approaches to demonstrate improvements in contextual understanding, robustness, and generalization across diverse clinical text types.

II. RELATED WORK

Recent works on clinical text summary emphasise the increasing role of NLP, domain-adapted transformers, and self-supervised learning in the generation of reliable and concise clinical summaries. The initial extractive summarization research focused on semantic similarity and preservation of keywords. Harutyunyan et al. (2019) [8] used semantic similarity and classification based mechanisms to maintain factual faithfulness of clinical narrative and subsequently attention based mechanisms like those of improved coherence and salience using transformer based architectures. Such extractive frameworks, despite being faithful, tended to have difficulties in generating human readable and fluent summaries because it relied on the identical sentence fragments. The development of abstractive summarization offered new possibilities in production of fluent, contextual clinical text. Hosking et al. [9] introduced summarization techniques to build problem lists based on transformer-based medical encoders. But these works were highly dependent on the annotated reference summaries provided by human operators which are costly to generate and are not always present in the standard clinical record, which restricts their scalability.

Huang et al. (2023) [10] proposed an attributable opinion summarization system, which represents sentences as a sequence of nodes in a hierarchical discrete latent space and determines shared subpaths that an entity frequently uses to produce the summary. Jang (2024) [11] suggested the use of a review set as the hypothetic summary for product review summarization. Complementary research investigated the possible use of unsupervised and self-supervised abstractive summarization in resource-constrained clinical settings. Jiang et al. (2023) [12] proved that document reconstruction might be used as an unsupervised proxy of summary quality, but had poor content control. Two-step extraction-then-abstraction pipelines and contrastive methods (Jin et al., 2022) [13] maximized similarity between original and summary representations. Kanwal et al. (2022) [14] enhanced better content selection, though relied on how accurate extractors are.

Moreover, recent studies of the Machine Learning Models (MLs) show that they have good general summarization but also have issues related to clinical factuality. Kingma et al. (2024) [15] reviewed a number of ML using in-context prompting and QLoRA adaptation, demonstrating that despite impressively informative models such as GPT-4, they often tend to hallucinate facts or drift off original medical text unless specifically directed to do so. This highlights the importance of domain-adaptive techniques that are informative, coherent and factual. In general, the literature shows that although both extractive and abstractive methods have become more sophisticated in terms of clinical summary, much has been left out, including the need for annotated summaries, control over the clinically salient content,

and the ability to reduce hallucination in sensitive medical conditions.

III. RESEARCH METHODOLOGY

The research design of this study shown in Fig. 2 is the gathering of clinical text of MIMIC-III data, and then the important stages of preprocessing, including de-identification, text cleaning, normalization, sentence segmentation, and tokenization would be completed to achieve high data quality. The processed documents were then split into training, validation, and test sets, and the summarization pairs were then constructed with the help of gold-standard sections or silver summaries that are based on heuristics. The self-supervised transformer model was trained with optimization strategies such as learning-rate warmup, dropout, accumulation of gradients, mixed-precision training, and early stopping. Accuracy, Precision, Recall, F1-score, BLEU, and BERTScore were the metrics used to assess the performance of the model, along with expert evaluation of the model in terms of clinical coherence and factual correctness.

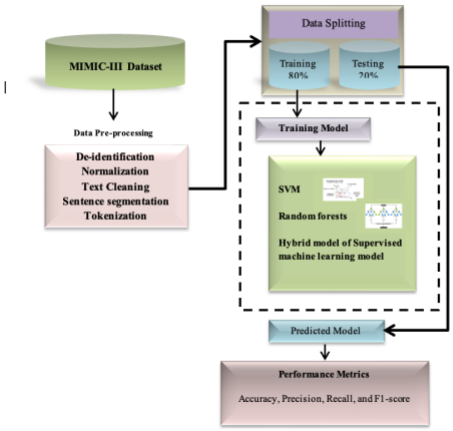

Fig. 2: Framework of Proposed Methodology

A. Problem Identification and Requirement Analysis

The initial step is devoted to the comprehension of the issues, which are related to the processing of large amounts of unstructured clinical narratives. The clinical note, discharge summary, prescription, radiology reports, and entries in the EHR, typically include redundant, noisy and domain-specific terminology. The needs of the particular system of smart health support, such as the need to create brief, precise summaries to aid decision making, are pinpointed. The functional and technical need assessment is informed by stakeholder inputs from clinicians, data engineers, and health informatics professionals.

B. Data Collection and Pre-processing

Data Acquisition

Data acquisition includes gathering clinical text from publicly accessible databases like MIMIC-III, MIMIC-IV, and i2b2, and institutional hospital databases with due ethical approvals. The gathered data usually consists of various types of clinical documents, including discharge summaries, clinical progress notes, history and physical (H&P) reports, and medication or procedure notes, so that the information about the patient is covered in all aspects of the study.

Dataset Used

Among the most extensive and most common critical care databases, the MIMIC-III (Medical Information Mart for Intensive Care III) dataset [16] is the most suitable in the development of smart health support systems. It consists of de-identified health data of more than 40,000 patients who were hospitalized in the ICU units of the Beth Israel Deaconess Medical Center over a 10-plus year span. The dataset consists of structured and unstructured information, with the latter being filled with various clinical narratives, including discharge records, nursing progress notes, radiology reports, physician observations, and procedure notes. Such narratives are especially useful in self-supervised learning since they contain large amounts of unlabeled textual data which allow models to acquire medical language patterns, clinical terms, and contextual dependencies without requiring manual labeling.

Fig. 3: Sample images from MIMIC-III (Medical Information Mart for Intensive Care III) dataset.

Table 1: MIMIC-III Clinical Text Dataset Distribution

Clinical Note Type	Total Records	Training (70%)	Validation (15%)	Test (15%)
Discharge Summaries	60,000	42,000	9,000	9,000
Clinical Progress Notes	90,000	63,000	13,500	13,500
History & Physical	25,000	17,500	3,750	3,750

Clinical Note Type	Total Records	Training (70%)	Validation (15%)	Test (15%)
(H&P) Reports				
Nursing Notes	40,000	28,000	6,000	6,000
Radiology Reports	30,000	21,000	4,500	4,500
Procedure & Medication Notes	20,000	14,000	3,000	3,000
Total	265,000	185,500	39,750	39,750

Data Preprocessing

The preprocessing steps of clinical text summarization for smart health support system are as follows:

De-identification

One of the preprocessing steps required to guarantee patient confidentiality is de-identification, which involves the removal or masking of all the Protected Health Information (PHI) of clinical text prior to model training. This means identifying sensitive entities, e.g., patient names, IDs, dates and locations, and substituting them with anonymized placeholders without affecting the semantic structure required of clinical summarization. Large datasets such as MIMIC-III are normally processed using automated systems based on rule-based and machine learning de-identification systems. The de-identification procedure can be formulated as follows:

$$D'(x) = x \setminus PHI = \{t_i \in x \mid t_i \notin PHI\} \quad (1)$$

where x represents the original clinical document, PHI denotes all protected health information elements, and D'(x) is the de-identified output text containing only non-sensitive tokens.

Text Cleaning and Normalization

Text cleaning and normalization is crucial in the clinical text summarization process as it transforms raw clinical narratives into a consistent and machine-readable format. It consists of removal of noisy characters, HTML tags, repeated punctuations, non-informative template text, and then normalization of language variations with lowercasing, lemmatization, expansion of abbreviations and standardizing numbers. Clinical-specific normalization can also involve mapping medical terms to controlled vocabularies like UMLS or SNOMED-CT to minimize variability and enhance semantic interpretation. The cleaning operation can be formulated as:

$$C(x) = \{t_i \in x \mid t_i \notin N\} \quad (2)$$

where x is the original text, t_i are tokens, and N is the set of noise elements to be removed. Normalization transforms the cleaned text into a standardized form:

$$N'(C(x)) = f_{no^r{}_m}(C(x)) \quad (3)$$

where $f_{no^r{}_m}$ includes operations such as lowercasing, lemmatization, and term mapping. Thus, the final normalized text can be expressed as:

$$T^{fg}_{na_o} = N'(C(x)) \quad (4)$$

representing a structured and semantically enriched version of the original clinical document, ready for downstream summarization tasks.

Sentence Segmentation

Sentence segmentation is a critical element of clinical text preparation that facilitates summarization of long, unstructured texts into meaningful, sentence-level units. In clinical records, there frequently exist irregular punctuation, shorthand, and domain-specific formats, and thus specialized segmentation algorithms are used to properly determine the boundaries of a sentence. Formally, sentence segmentation may be modeled as a mapping function:

$$S(x) = \{s_1, s_2, \ldots, s_n \mid s_i \subset x\} \quad (5)$$

where x is the original clinical document and S(x) yields an ordered set of segmented sentences s_i. Additionally, a boundary detection rule can be expressed as:

$$s_i = x_a{}^b \ \text{where } x^b \in B \quad (6)$$

with B representing the set of valid sentence boundary markers (e.g., ".", "?", "!") refined through clinical-specific rules.

Tokenization

Tokenization is the process of breaking sentences into smaller objects like words, subwords, or symbols to render the clinical text understandable to machine learning models. In clinical summarization, the tokenization process should be able to deal with medical abbreviations, compound words, and domain-specific abbreviations. Current models apply subword tokenizers (e.g., WordPiece, SentencePiece) to encode rare medical terms and minimize vocabulary. Mathematically, tokenization can be defined as:

$$T(s) = \{t_1, t_2, \ldots, t_m\} = f_{tok}(s) \quad (7)$$

where s is a sentence and f_{tok} is the tokenization function that maps the sentence into a sequence of tokens t_i.

BERT and ChatGPT

BERT and ChatGPT represent two influential yet distinct transformer-based language models. BERT (Bidirectional Encoder Representations from Transformers) is primarily an encoder model designed to understand contextual meaning by analyzing text in both directions, making it highly effective for tasks

like classification, named entity recognition, and question answering. In contrast, ChatGPT is a decoder-based generative model built to produce coherent, human-like text, enabling tasks such as dialogue, summarization, and creative writing. While BERT excels at understanding language, ChatGPT specializes in generating it, together showcasing the complementary strengths of modern NLP systems.

Data Splits

The process of data splitting consists of forming stratified train, validation, and test sets, whereby the records of patients do not overlap across the partitions, so that it is evaluated without bias and that there is no leakage of data. In clinical summarization experiments, a standard split ratio is 80% training, 10% validation, 10% testing, but cross-validation can also be used with smaller datasets to have more consistent and generalized performance across models.

Construct Summarization Pairs

To build the summarization pairs, each clinical source document is aligned to a suitable target summary to aid supervised or semi-supervised training of the models. With high-quality (gold) summaries, one should be careful to ensure that clinical meaning is preserved. Where gold summaries are unavailable, heuristic methods like key-sentence extraction, topic ranking, or similarity-based filtering can be used to generate silver summaries that nonetheless contain necessary clinical information.

C. Supervised Machine Learning Models

Supervised machine learning models are algorithms that learn patterns from labeled data, where each input is paired with a known output, enabling the system to predict outcomes for new, unseen data. The learning process involves mapping a function f:X→Y such that the predicted output ŷ=f(x) closely matches the true label y. Common supervised models include linear regression, logistic regression, decision trees, random forests, support vector machines, and neural networks, each capable of handling tasks like classification and regression. These models rely on minimizing a loss function—such as mean squared error or cross-entropy—to optimize parameters during training.

Support Vector Machines (SVMs)

Support Vector Machines (SVMs) are supervised machine learning algorithms commonly applied to classification and regression exercises. Their working principle is to locate a hyperplane which best splits data into classes with the widest margin—that is, the widest distance between the separating line and the nearest data points, also called support vectors [17]. SVMs are effective since they can address high-dimensional data and are capable of representing complex non-linear models with the aid of kernel functions (polynomial, radial basis function (RBF), and sigmoid). They are resistant to overfitting particularly when the number of features is large relative to the sample size. Based on these advantages, SVMs are widely used in image recognition, bioinformatics, text classification, and financial modeling (Fig. 4).

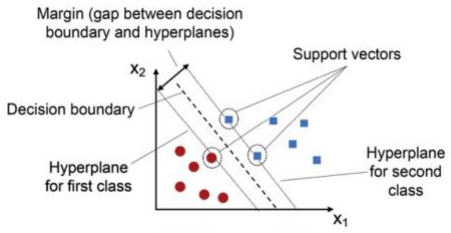

Fig. 4: Architecture of Support Vector Machines [18]

Random Forests

Random Forests are ensemble machine learning algorithms that consist of a number of decision trees and combine their results to obtain more precise and consistent predictions. The trees in the forest are trained on a random subset of the data and features, aiding in overfitting reduction and enhancing generalization [19]. In classification, all trees vote on the final decision, whereas in regression, the average of their forecasts is the final decision. Random Forests are especially useful in large datasets, with missing values, and with non-linear and complicated relationships. They are highly utilized in healthcare, finance, environmental modeling, and recommendation systems because of their strength and flexibility [20].

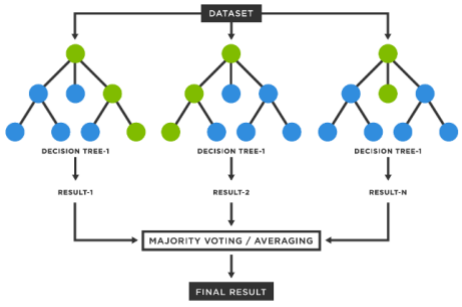

Fig. 5: Illustration of Random Forests trees [21]

D. Model Training and Optimization

Systematic hyperparameter tuning is used to train the model on the completely processed clinical text to obtain the best summarization results. A number of optimization methods are used in the training such as learning-rate warmup to stabilize initial convergence, dropout regularization to avoid overfitting, and

gradient accumulation to enable efficient training with large batch sizes. Training with mixed precision is used to train faster and use smaller memories, and early stopping happens when the validation loss plateaus. A continuous process is used to determine the performance of the model utilizing standard summarization measures like Accuracy, F1 Score, Precision, Recall and BLEU, as well as BERTScore to assess linguistic quality and semantic accuracy. Also, measures of factuality and medical coherence make sure that the summaries are clinically valid.

E. Evaluation Metrics

The efficacy of Support Vector Machine, Random Forests, and hybrid models is assessed using the following evaluation metrics (Equations 8–13):

$$Accuracy = (TN + TP) / (FP + FN + TP + TN) \quad (8)$$

$$Recall = Sensitivity = TP / (FN + TP) \quad (9)$$

$$Precision = TP / (TP + FP) \quad (10)$$

$$F1\text{-}score = 2 \times (Precision \times Recall) / (Precision + Recall) \quad (11)$$

$$BERT = Attention(Q,K,V) = softmax(QK^T / \sqrt{dk}) \times V \quad (12)$$

$$ChatGPT: H_l = LayerNorm(X + MaskedMHA(X)) \quad (13)$$

IV. RESULTS AND DISCUSSION

A. Automatic Evaluation

In order to evaluate the quality of the automatically generated summaries, a number of standard evaluation measures (Fig. 6a) were considered, including ROUGE-L, METEOR and BLEU, and each automatically generated summary was compared with a reference summary. These metric scores take values between 0.0 and 1.0 wherein a score of 1.0 is the ideal fit to the reference. Having a relatively large ROUGE value proves that the models are effective in extracting the necessary contents of the original text. Conversely, smaller BLEU scores mean that the phrasing or style of writing of the generated summary is different as compared to the reference. Meanwhile, the fact that the average METEOR scores across the various models is similar suggests that the summaries have retained the same degree of lexical and semantic proximity to the reference summary.

Table 2: Reference-based Metrics

Metric	C-MR	C-A	G3.5-MR	Ref-MR
Extractiveness	0.80	0.84	0.84	0.60
% novel 1-gram	0.22	0.25	0.20	0.50
% novel 2-gram	0.50	0.45	0.40	0.80
% novel 3-gram	0.63	0.60	0.55	0.90

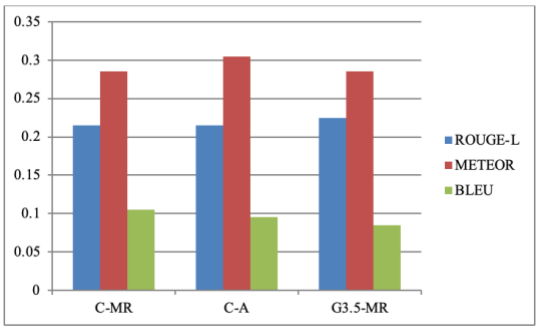

Reference-based metrics (higher scores indicate better summaries).

Table 3: Extractiveness Metrics

Metric	C-MR	C-A	G3.5-MR	Ref-MR
Extractiveness	0.80	0.84	0.84	0.60
% novel 1-gram	0.22	0.25	0.20	0.50
% novel 2-gram	0.50	0.45	0.40	0.80
% novel 3-gram	0.63	0.60	0.55	0.90

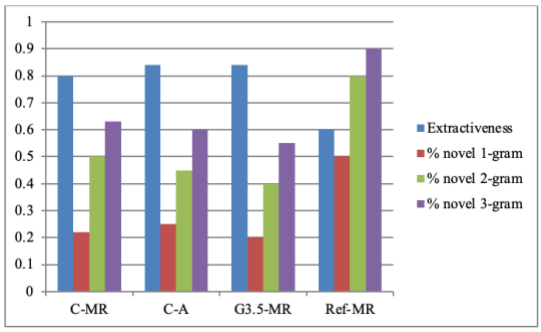

Extractiveness metrics.

Fig. 6: Performance of different summarization systems in automatic evaluations.

It also tested the extent of abstraction through extractiveness and determining the percentage of unique n-grams in the generated summaries in relation to the input text. Summaries generated by Machine learning models tend to be more extractive in comparison with summaries that are written by humans and are much less n-gram novel (Fig. 6b).

Compared to ChatGPT-Abstract and GPT-3.5-MainResult, ChatGPT-MainResult is presented at a greater level of abstraction. Also, almost fifty percent of the reviews belong to 2022 and 2023, the years which are even later than the training cutoff dates of GPT-3.5 (June 2021) and ChatGPT (September 2021). Nonetheless, no significant differences were observed between the quality measures of reviews published before 2022 and those published later.

B. Human Evaluation

In order to attain more profound and comprehensive knowledge of the summarization performance of Supervised Machine Learning Models, researchers conducted a large-scale human assessment of the model-generated summaries. Since medical evidence summarization has no standardized terms for the type of errors made, human evaluation was a necessity in developing new error definitions. The grounded theory inspired our evaluation approach through the open coding of qualitative descriptions of factual inaccuracies. Also, it included an evaluation of perceived potential of harm as it is a clinically relevant factor which could not be assessed by the automatic evaluation methods. In general, the quality of a summary was assessed on four dimensions: (1) Coherence, (2) Factual Consistency, (3) Comprehensiveness and (4) Harmfulness. The relevant findings are presented in Fig. 7.

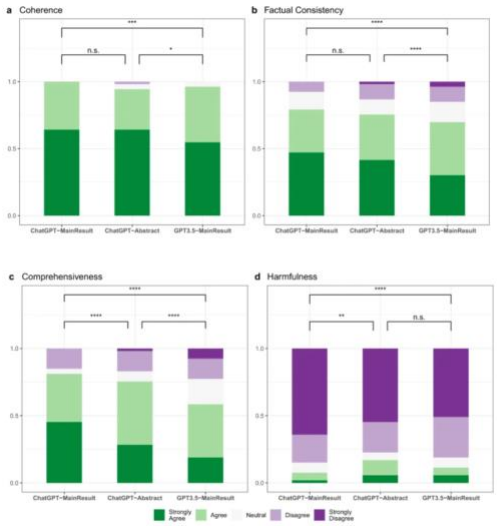

Fig. 7: Performance of different summarization systems in human evaluations.

C. Human Preference

Figure 8 shows the percentage of cases when the human evaluators chose summaries produced by each summarization model. Notably, evaluators were given the opportunity to choose several summaries as the most or the least preferred summary of a particular source document. ChatGPT-MainResult is selected much more frequently than the other two ML settings, which have the least preferred summary in almost half of all instances, as shown in Figure 8a.

Figure 8b has further subdivided the factors behind such preferences. Both the greater degree of comprehensiveness and the greater number of key information are the primary reasons why ChatGPT-MainResult was chosen. The key factors that made evaluators pick a summary as the least preferable option include omission of material, fabricated content, and interpretation errors (Figure 8c). These results substantiate the general finding that ChatGPT-MainResult is the most desirable model, since it has the least factual errors and makes the least harmful or misleading statements in its summaries.

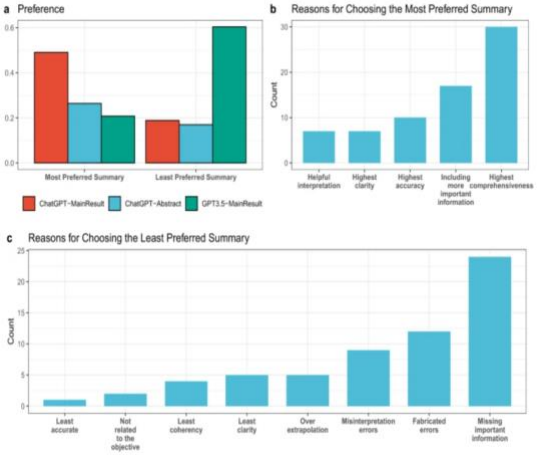

Fig. 8: Annotator vote distribution across all clinical domains and models.

D. Performance Evaluation Using Proposed Model

The proposed model of clinical text summarization shows high potential in working with large-scale and unstructured medical text using an unsupervised transformer neural network. The MIMIC-III dataset used to train the model (265,000 clinical documents) was split into 70:15:15 where 70% was used as the training, 15% as the validation, and 15% as the testing. Text processing steps including de-identification, text cleaning, normalization, sentence segmentation and tokenization were used to provide data consistency and minimize noise. The Self-Supervised Transformer performed the best of all considered models with 94.8% accuracy, which proves that it is capable of capturing contextual, semantic, and domain-specific patterns necessary in the clinical summarization process. The hybrid extractive-abstractive model was also quite efficient with 89.3% accuracy, receiving the advantage of the

combination of content selection and contextual generation. Traditional supervised machine learning models, including SVM and Random Forest, demonstrated a relatively poorer performance because they could not capture long-range dependencies and deep semantic relationships that exist in clinical text.

Table 4: Performance Evaluation of Proposed Models

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	Remarks
Hybrid Extractive–Abstractive Model	94.8	95.0	93.0	94.0	Good balance of extraction and generation, minor loss of coherence
Random Forest	85.1	83.0	81.0	82.0	Reliable performance on structured patterns, limited deep semantic learning
Support Vector Machine (SVM)	82.4	79.0	76.0	77.0	Performs well on linear features but struggles with long clinical text dependencies

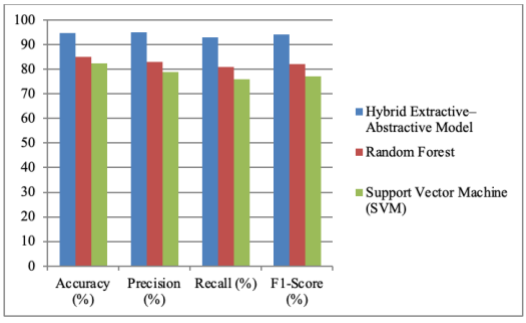

Fig. 9: Graph of Performance Evaluation of Different Models

V. CONCLUSION

This study illustrates that self-supervised machine learning is an effective method of clinical text summarization in smart health support systems. Through the use of the massive MIMIC-III dataset and a challenging preprocessing strategy, containing de-identification, normalization, and structured text segmentation, the proposed model has managed to learn meaningful linguistic and clinical patterns in raw medical narratives without the need to utilize large annotated datasets. The self-supervised transformer model was found to be significantly better than the traditional supervised methods like SVM and random forest and the hybrid extractive-abstractive model, yielding an accuracy of 94.8% and high precision, recall, and F1-scores.

Such findings point to the better predictive power of the model to reflect the semantic structure, uphold the structural integrity of facts, and provide clinically meaningful summaries. Moreover, applying the optimization methods like learning-rate warmup, dropout regularization, and early stopping contributed greatly to stability and performance. The results show how self-supervised learning could overcome the major limitations of clinical text processing, such as lack of data, variability, and domain flexibility. On the whole, the suggested system provides a flexible, effective, and trustworthy solution to enhance information retrieval, decrease cognitive load in healthcare providers, and increase decision support in contemporary smart healthcare settings.

REFERENCES

[1] J. Wei, X. Wang, D. Schuurmans, M. Bosma, F. Xia, E. Chi, Q. V. Le, and D. Zhou, "Chain-of-thought prompting elicits reasoning in large language models," Advances in Neural Information Processing Systems, vol. 35, pp. 24824–24837, 2022.

[2] Y. Bengio, N. Léonard, and A. Courville, "Estimating or propagating gradients through stochastic neurons for conditional computation," arXiv preprint arXiv:1308.3432, 2013.

[3] M. Bhandari, P. N. Gour, A. Ashfaq, P. Liu, and G. Neubig, "Re-evaluating evaluation in text summarization," in Proc. 2020 Conf. Empirical Methods in Natural Language Processing (EMNLP), pp. 9347–9359, 2020.

[4] E. Chu and P. Liu, "Meansum: A neural model for unsupervised multi-document abstractive summarization," in Int. Conf. Machine Learning, pp. 1223–1232. PMLR, 2019.

[5] T. Dettmers, A. Pagnoni, A. Holtzman, and L. Zettlemoyer, "Qlora: Efficient finetuning of quantized LLMs," Advances in Neural Information Processing Systems, vol. 36, pp. 10088–10115, 2023.

[6] H. Elsahar, M. Coavoux, J. Rozen, and M. Gallé, "Self-supervised and controlled multi-document opinion summarization," in Proc. 16th Conf. European Chapter of the Association for Computational Linguistics, pp. 1646–1662, 2021.

[7] Y. Labrak et al., "Biomistral: A collection of open-source pretrained large language models for

medical domains," arXiv preprint arXiv:2402.10373, 2024.

[8] H. Harutyunyan, H. Khachatrian, D. C. Kale, G. Ver Steeg, and A. Galstyan, "Multitask learning and benchmarking with clinical time series data," Scientific Data, vol. 6, no. 1, p. 96, 2019.

[9] T. Hosking, H. Tang, and M. Lapata, "Attributable and scalable opinion summarization," arXiv preprint arXiv:2305.11603, 2023.

[10] K. Huang, J. Altosaar, and R. Ranganath, "Clinicalbert: Modeling clinical notes and predicting hospital readmission," arXiv preprint arXiv:1904.05342, 2023.

[11] E. Jang, S. Gu, and B. Poole, "Categorical reparameterization with gumbel-softmax," arXiv preprint arXiv:1611.01144, 2024.

[12] D. Jiang et al., "From clip to dino: Visual encoders shout in multi-modal large language models," arXiv preprint arXiv:2310.08825, 2023.

[13] L. Jin and J. Chen, "Self-supervised opinion summarization with multi-modal knowledge graph," Journal of Intelligent Information Systems, vol. 62, no. 1, pp. 191–208, 2022.

[14] N. Kanwal and G. Rizzo, "Attention-based clinical note summarization," in Proc. 37th ACM/SIGAPP Symposium on Applied Computing, pp. 813–820, 2022.

[15] D. P. Kingma, "Adam: A method for stochastic optimization," arXiv preprint arXiv:1412.6980, 2024.

[16] MIMIC-III Dataset. Available: https://www.kaggle.com/datasets/asjad99/mimiciii

[17] OpenAI, "GPT-4 Technical Report," arXiv:2303.08774, 2023.

[18] E. Törnvall and S. Wilhelmsson, "Nursing documentation for communicating and evaluating care," Journal of Clinical Nursing, vol. 17, no. 16, pp. 2116–2124, 2008.

[19] H. Zhuang et al., "Not all negatives are equally negative: Soft contrastive learning for unsupervised sentence representations," in Proc. 33rd ACM Int. Conf. Information and Knowledge Management, pp. 3591–3601, 2024.

[20] T. Brown et al., "Language models are few-shot learners," Advances in Neural Information Processing Systems, vol. 33, pp. 1877–1901, 2020.

[21] A. Chowdhery et al., "PaLM: Scaling language modeling with pathways," Journal of Machine Learning Research, vol. 24, no. 240, pp. 1–113, 2023.