

UMass-BioNLP at MEDIQA-M3G 2024: DermPrompt - A Systematic Exploration of Prompt Engineering with GPT-4V for Dermatological Diagnosis

Parth Vashisht*, Abhilasha Lodha*, Mukta Maddipatla*,
Zonghai Yao, Avijit Mitra, Zhichao Yang, Junda Wang, Sunjae Kwon, Hong Yu

CICS, University of Massachusetts, Amherst, MA, USA
{pvashisht, alodha, mmaddipatla, zonghaiyao}@umass.edu

Abstract

This paper presents our team’s participation in the MEDIQA-ClinicalNLP 2024 shared task B. We present a novel approach to diagnosing clinical dermatology cases by integrating large multimodal models, specifically leveraging the capabilities of GPT-4V under a retriever and a re-ranker framework. Our investigation reveals that GPT-4V, when used as a retrieval agent, can accurately retrieve the correct skin condition 85% of the time using dermatological images and brief patient histories. Additionally, we empirically show that Naive Chain-of-Thought (CoT) works well for retrieval while Medical Guidelines Grounded CoT is required for accurate dermatological diagnosis. Further, we introduce a Multi-Agent Conversation (MAC) framework and show its superior performance and potential over the best CoT strategy. The experiments suggest that using naive CoT for retrieval and multi-agent conversation for critique-based diagnosis, GPT-4V can lead to an early and accurate diagnosis of dermatological conditions. The implications of this work extend to improving diagnostic workflows, supporting dermatological education, and enhancing patient care by providing a scalable, accessible, and accurate diagnostic tool.¹

1 Introduction

Diagnosing skin conditions demands a complex blend of visual inspection, patient history examination, and deep clinical acumen, a skill set that dermatologists spend extensive years acquiring (Mangion et al., 2023). Despite the critical nature of these skills, many regions worldwide face a stark scarcity of dermatological expertise (Benner et al., 2009). Even in areas with adequate services, the demand for such specialized knowledge frequently surpasses its availability. The recent global

health crisis has also expedited the shift towards remote clinical diagnostics and treatments, further highlighting the challenges in diagnosing skin diseases (Behar et al., 2020). These challenges include the scarcity of dermatological expertise and the need to accommodate asynchronous patient interactions, including e-visits, emails, and messaging platforms, to ensure continuity and quality of care.

In response to these challenges, recent advancements in Artificial Intelligence (AI), particularly through the development of large language models (LLMs), offer promising solutions to significantly support dermatologists by enhancing clinical diagnosis and treatment processes (McDuff et al., 2023; Singhal et al., 2023b; Tu et al., 2024). Moreover, AI facilitates asynchronous patient services, offering a cost-effective and convenient alternative to traditional methods. Previous works have primarily utilized deep learning for tasks such as skin lesion classification (Udriștoiu et al., 2020; Esteva et al., 2017; Brinker et al., 2019), and dermatopathology (Hekler et al., 2019; Jiang et al., 2020) focusing predominantly on dermoscopic images (Cruz-Roa et al., 2013). These efforts, however, have relied on image-only models, indicating a need for broader applications.

Our research aims to extend the capabilities of AI in dermatology by diagnosing skin diseases and devising appropriate treatment plans based on patients’ dermatological images, queries, and medical histories. This approach mirrors the diagnostic process of dermatologists, who rely on high-quality images and comprehensive patient histories to make informed decisions. Although previous studies have explored fine-tuning models on multimodal data (e.g., SkinGPT (Zhou et al., 2023), and MedBLIP (Chen et al., 2023a)), our task is particularly challenging due to data availability and image quality limitations, reflecting real-world constraints where high-quality data is either scarce or expensive to obtain.

* Equal Contribution

¹The code is released at [Github](#)

Studies by OpenAI (Nori et al., 2023a) and Microsoft (Nori et al., 2023b) have demonstrated that generalist foundation models, such as GPT-4, can surpass specifically fine-tuned medical models on various medical benchmarks by employing specialized prompting strategies. Building on these insights, our research leverages both the textual and visual capabilities of GPT-4, targeting the specific task of dermatology diagnosis and treatment.

The diagnostic process for skin lesions or conditions requires meticulous evaluation and is informed by methodologies such as dermatoscopy (Panagoulas et al., 2024), which enables dermatologists to observe skin abnormalities in greater detail. Dermatologists usually follow a common guideline for assessing skin lesions, emphasizing the importance of visual descriptors like shape, size, color, texture, and pattern in differential diagnosis. Inspired by these practices, we have integrated advanced Chain-of-Thought (CoT) techniques with visual features to create medical guidelines tailored for GPT-4V, enhancing its diagnostic precision. This enables the model to emulate dermatologists' diagnostic process.

Furthermore, our research integrated a Multi-Agent Conversation (MAC) Framework (Tao et al., 2024a; Wu et al., 2023; Li et al., 2023b), which involves multiple AI agents that generate additional context and critiques for various candidate skin conditions. These agents collaborate, debate, and consolidate their findings to determine the most accurate skin disease diagnosis from the candidates identified from the retrieval step. This, therefore, introduces a level of dynamic interaction and comprehensive analysis that mirrors the complex decision-making process in clinical dermatology.

Hence, our contributions are twofold:

- We deploy GPT-4V within a novel retrieval and re-ranking framework, critically evaluating the effectiveness of various prompting strategies. These include both naïve prompts and those meticulously crafted based on detailed medical guidelines (CoT), across different stages of our setup. This exploration aims to highlight the adaptability and precision of GPT-4V in simulating the diagnostic reasoning of dermatologists.
- We explore the Multi-Agent Conversation (MAC) Framework in the context of clinical dermatology, examining its potential to en-

rich the diagnostic process. Through this discussion, we identify and delineate the framework's strengths and limitations, offering insights into its applicability and performance in accurately diagnosing skin diseases.

2 Related Work

The interdisciplinary fusion of Artificial Intelligence (AI) and dermatology has spawned a myriad of approaches to enhance the diagnosis of skin conditions. Historically, these approaches have often treated diagnosis as a classification task, with literature extensively documenting the use of convolutional neural networks (CNNs) and other deep learning architectures like ResNets for lesion classification from dermoscopic images, which are typically limited to dermatological clinics due to image acquisition constraints (Ba et al., 2022), (Gouda and Amudha, 2020).

Recent advancements have moved beyond traditional clinic-bound methods, exploring the utility of clinical images for broader classifications, such as skin cancer and onychomycosis (Sharma et al., 2022). While these efforts have made significant contributions to disease diagnosis, they have not fully addressed the generative and comprehensive nature of clinical diagnosis, which encompasses treatment planning and patient interaction beyond mere classification.

The evolution of large language models (LLMs) has significantly widened the scope of AI applications in healthcare. LLMs like PubMedBERT (Gu et al., 2020) and BioGPT (Luo et al., 2022) have been fine-tuned on extensive corpora of medical literature, achieving state-of-the-art performance in tasks ranging from biomedical reasoning to question-answering. In the realm of domain-adapted LLMs, models like Meditron and Med-PALM have demonstrated remarkable capabilities in language understanding and generation, setting new benchmarks across biomedical datasets (Chen et al., 2023b; Singhal et al., 2022).

With the advent of multimodal models, integrating visual and textual data has further refined AI applications in medical domains. Vision-language models such as Med-CLIP, Med-BLIP, and Llava-Med have exhibited promising results in image-text retrieval, zero-shot classification, and even multimodal conversations, respectively (Chen et al., 2023a; Li et al., 2023a). Specifically, in the context of dermatology, the Skin-GPT4 model (Zhou

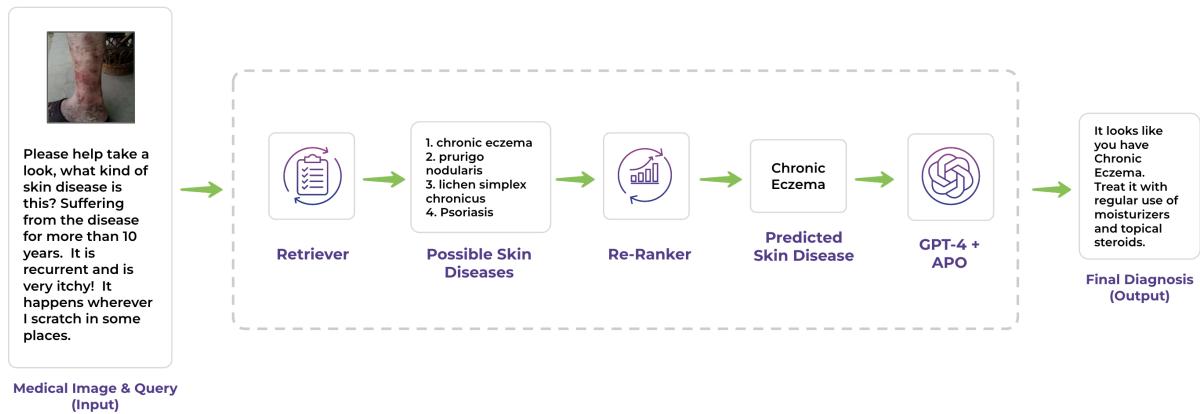


Figure 1: Overview of the AI-assisted dermatology diagnosis pipeline, from initial patient input through to the GPT-4V generated final diagnosis and treatment plan.

et al., 2023) represents a pioneering effort in creating a multimodal setup tailored for skin disease identification and patient interaction.

Studies have shown that generalist foundation models like GPT-4, with their expansive knowledge bases and specialized prompting techniques, outperform domain-specific models such as MedPALM on various medical benchmarks (Nori et al., 2023a). GPT-4’s application in dermatology, particularly in melanoma identification and medical exam question answering, underscores its potential as an assistive tool for educational and diagnostic purposes (Miao et al., 2024; Mishra et al., 2024; Yang et al., 2023).

Our research builds upon these foundations, employing GPT-4’s multimodal capabilities (GPT-4V) to enhance dermatological diagnostic processes. By integrating Chain-of-Thought (CoT) techniques and a Multi-Agent Conversation (MAC) Framework (Tao et al., 2024a; Wu et al., 2023; Li et al., 2023b), we aim to emulate the complex decision-making process of dermatologists, enriching the GPT-4V’s ability to generate diagnostic and treatment plans from multimodal data. This work not only taps into the multimodal analytical strength of GPT-4V but also seeks to optimize the model’s performance in a domain where the nuances of patient history and visual inspection are paramount.

Thus, our contribution to the field involves the innovative use of GPT-4V within a retrieval and re-ranking framework, leveraging both naïve and medically informed CoT prompting strategies.

3 Methodology

Our methodology delineates the comprehensive approach we adopted to address the task of mul-

timodal medical answer generation. This process involves two primary stages: retrieving potential diagnoses and the ranking of these to identify the most probable skin condition and treatment plan.

Task Description: The objective of our research is to develop a system capable of diagnosing a possible skin condition and recommending a corresponding treatment plan based on a patient’s medical query and associated image. To accomplish this, we propose a two-step pipeline consisting of a retrieval module followed by a ranker module. Specifically:

- **Retrieval Module:** This component extracts a list of possible skin conditions from the given image and medical query.
- **Ranker Module:** This module’s task is to select the most accurate skin condition diagnosis from the list generated by the retrieval module.

Our overall methodology is mentioned in Figure 1.

3.1 Retrieval Module

The Retrieval Module is the initial phase of our diagnostic approach, reflecting the dual aspects a dermatologist considers when evaluating a medical condition: visual inspection and patient history. Inspired by recent works such as MedGENIE (Frisoni et al., 2024), along with others (Yu et al., 2022; Zhang et al., 2023; Su et al., 2022), we leverage LLMs as strong context generators instead of traditional retrieval methods, such as keyword-based methods (e.g., BM25 Robertson et al. (2009)), vector-similarity-based methods (e.g., ColBERT Khattab and Zaharia (2020)), and

some internet tools (e.g., Google API). In this retrieval step, we treat LLMs as a knowledge base (Singhal et al., 2023a) to generate potentially valuable information for subsequent steps. Our module employs two distinct strategies:

3.1.1 Context-Independent Retrieval (Image-Only)

Recognizing scenarios where comprehensive medical context (patient’s medical history and medical queries) may not be readily available, we engage in context-independent retrieval. This approach leverages GPT-4V to identify possible skin conditions based solely on image data. We compared this model’s performance against a widely-used online AI tool, First Derm²

3.1.2 Context-Dependent Retrieval (Image + Context)

The inclusion of medical context is pivotal for accurate diagnosis. Particularly, incorporating details about systemic conditions and patient history can significantly influence differential diagnosis, a critical aspect of clinical dermatology. To this end, we utilize CoT prompting, a technique that simplifies complex problems into manageable objectives, enabling the model to address the larger task.

Within the Context-Dependent Retrieval, we experiment with two strategies:

- Naive CoT: Here, GPT-4V is instructed to methodically analyze all relevant information from the images and medical query before generating a list of potential skin conditions. This process mimics the step-by-step procedural thinking a dermatologist might employ.
- Expert Guidelines Grounded CoT: This approach involves crafting prompts based on the [Clinical Guidelines](#) that dermatologists follow, encapsulating a generic framework for skin disease diagnosis. Such frameworks typically comprise patient history, visual inspection, and differential diagnosis. Our Expert-CoT strategy emphasizes key visual characteristics like the lesion’s shape, color, size, location, and texture. By integrating this data, the module produces a detailed list of differential diagnoses needed to enhance the model’s diagnostic precision further.³

²<https://firstderm.com/>

³The specific prompts utilized for both the Naive CoT and

3.2 Re-Ranker Module

After the retrieval module identifies potential skin conditions, the re-ranker module is critical in our diagnostic pipeline. Its primary objective is to meticulously refine the preliminary list, pinpointing the diagnoses with the highest probability of accuracy. To achieve this, we experiment with four re-ranking strategies:

1. Naive Chain of Thought (CoT)
2. Expert Guidelines Grounded CoT with Context
3. Expert Guidelines Grounded CoT without Context
4. Multi-Agent Conversation Framework

These structured approaches enable a systematic evaluation of the candidate’s conditions, ensuring that the decision-making process mirrors the analytical and deductive reasoning of a dermatologist. The specific prompts utilized for these three CoT techniques are presented in Tables 14, 15, and 16. The MAC framework is explained with an example in the Appendix A.

3.2.1 Naive Chain of Thought (CoT)

In the Naive CoT approach, GPT-4V is initially instructed to analyze the patient’s medical query and the associated images to extract relevant information. Subsequently, each candidate skin condition retrieved from section 3.1 is assigned a score ranging from 1 to 10, where 1 signifies the least probable and 10 denotes the most probable condition. The model identifies the most probable disease based on the scores and analysis. The prompt is mentioned in Table 14.

3.2.2 Expert Guidelines Grounded CoT with Context

This method employs a sophisticated strategy by utilizing prompts meticulously designed around the [Clinical Guidelines](#) followed by dermatologists. These guidelines encapsulate a comprehensive visual assessment of the affected area, scrutinizing distinct characteristics such as shape, size, color, location, and texture (listed in Table 15). Within this framework, GPT-4V is initially directed to conduct an analysis of the patient’s condition, incorporating insights drawn from their medical query. Subsequently, the patient’s images undergo visual examination using the defined guidelines, using which

the Expert Guidelines Grounded CoT strategies are detailed in Table 13

relevant features are extracted. The final step involves considering a list of possible skin conditions (retrieved from section 3.1) and systematically ruling them out based on the gathered insights and visual inspections to identify the most probable skin condition from the set of candidates.

3.2.3 Expert Guidelines Grounded CoT without Context

This approach omits the user query, focusing exclusively on the visual examination of dermatological conditions as per established [guidelines](#). Utilizing GPT-4V, an initial step involves the generation of a detailed visual description, drawing upon ten specified visual features essential for dermatological assessment (as outlined in Table 16). Subsequently, each candidate's skin condition is described visually, emphasizing distinguishing features aligned with the visual guidelines. A comparative analysis is then conducted between the visual descriptions of the candidates and the initial image description, and a score ranging from 1 to 10 is assigned based on the level of match (1 being the lowest match and 10 the highest). The most probable candidate, determined by the highest score in the comparative analysis, is selected as the diagnosis.

3.2.4 Multi-Agent Conversation Setup

Inspired by the recent various applications of the Multi-Agent Conversation framework in the general and medical domains ([Wu et al., 2023](#); [Tao et al., 2024b](#)), we also implement a multi-agent conversation framework for our re-ranker module (Figure 2). This framework involves multiple AI agents, each specializing in a different aspect of dermatology diagnosis. These agents collaborate, debate, and consolidate their findings to identify the most accurate diagnosis, mirroring the collaborative approach often seen in medical panels. This multi-agent setup not only enriches the model's diagnostic capabilities but also introduces a level of dynamic interaction and comprehensive analysis that mirrors the complex decision-making process in clinical dermatology. Moreover, acknowledging performance gain and consistency improvement obtained using critique-based refinement in large language models, we incorporate feedback generation as an objective of multi-agent debate followed by refinement.

The main components of our multi-agent setup are defined in Table 1.

Process Flow

- Assignment and Analysis
 - The Coordinator assigns a distinct probable disease to each Diagnostic Specialist based on the case study and list of probable diseases.
 - Each specialist analyzes the case study, provides evidence supporting their assigned disease and critiques the applicability of other diseases.
- Compilation and Presentation of Findings
 - After receiving inputs from all specialists, the Coordinator compiles the evidence and critiques.
 - The compiled information is presented to the Admin for evaluation.
- Evaluation and Revision
 - The Admin reviews the evidence and critiques, identifying areas where additional clarity or strengthening is needed.
 - If necessary, the Admin requests revisions from specific specialists to enhance their evidence based on critiques.
- Final Diagnosis
 - With the revised evidence, the Admin conducts a final review to determine the most accurate diagnosis.
 - The process concludes once the Admin confirms the diagnosis.

3.3 Aligner

The Aligner Module represents the final step in our diagnostic process, focusing on optimizing the model's output to ensure it aligns with clinical standards and expectations. This involves adjusting the prompt to refine the model's language and structure, aiming to emulate the concise, informative style characteristic of professional medical advice. The optimization process is guided by analyzing real doctor responses in the dataset, identifying key elements such as terminology, format, and the inclusion of essential diagnostic and treatment information. The goal is to produce a diagnosis and treatment plan that not only accurately identifies the patient's condition but also provides actionable, understandable advice. This module highlights our commitment to bridging the gap between AI-generated content and the practical needs of clinical practice, ensuring that the output is not

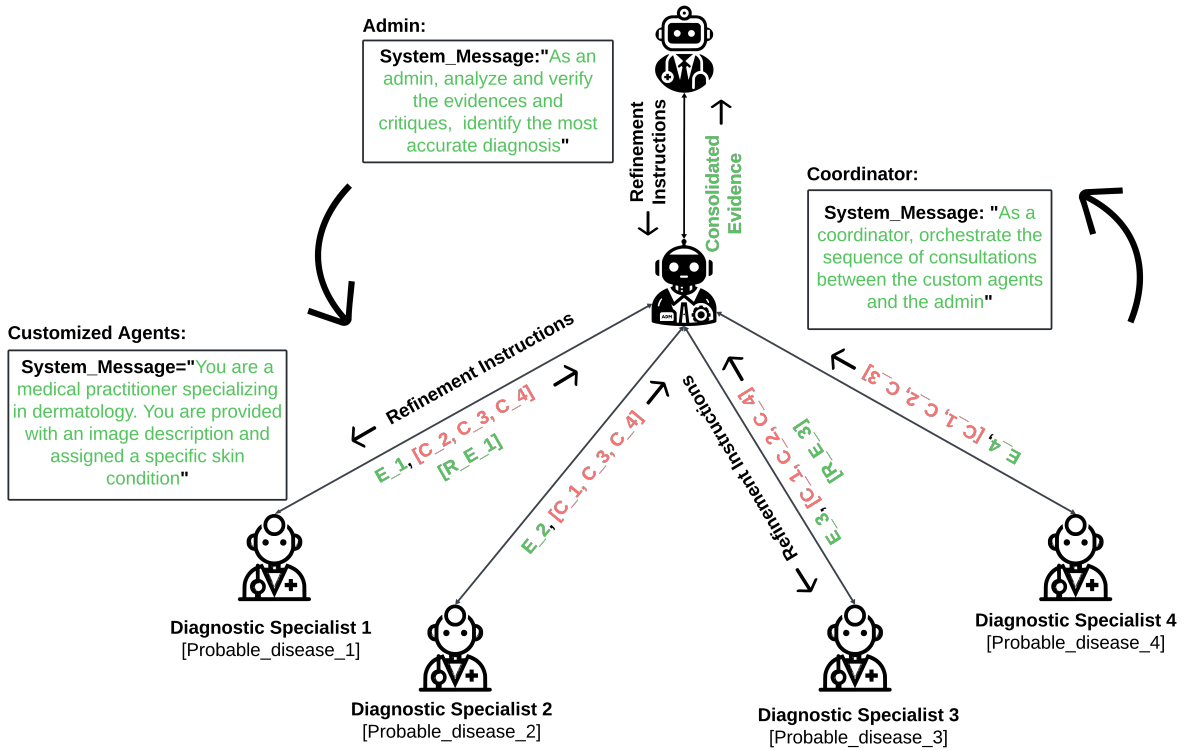


Figure 2: Multi-Agent Conversation (MAC) Setup

Note: E_1 corresponds to the evidence supporting *probable_disease_1* generated by *Diagnostic_Specialist_1*. C_2 , C_3 , and C_4 are the critiques for *probable_disease_2*, 3 & 4 generated by *Diagnostic_specialist_1*. Based on the Re-definment Instructions from the Admin, the diagnostic specialist returns Redefined evidence(R_E_1)

| Roles | Tasks |
|--------------------|---|
| <i>Coordinator</i> | Orchestrate sequence of consultations. Assign diseases to specialists and manage communications. |
| <i>Admin</i> | Evaluate evidence and critiques for accuracy. Request evidence enhancements and finalize the diagnosis. |
| <i>Agent</i> | Analyze case study, advocate for one disease. Provide evidence and critique alternative diagnoses. |

Table 1: Tasks for each role in a multi-agent setup

only technically accurate but also clinically relevant and usable in real-world medical contexts.

Recent work has introduced aligners to assist LLMs in generating harmless outputs (Ji et al., 2024), a concept previously applied in the style transfer domain to map model outputs to desired forms (e.g., formality style transfer (Rao and Tetreault, 2018; Yao and Yu, 2021)). When using third-party APIs like GPT, where updating the model’s weights is not an option, recent methods have explored the use of Automatic Prompt Optimization (APO) to improve prompts, assuming access to training data and an LLM API (Prasad et al., 2022; Pryzant et al., 2023). Recent studies

have also applied APO in the clinical domain to assist doctors in generating better note-generation prompts (Yao et al., 2023). Inspired by these efforts, we use human responses from training data as APO’s training input, allowing the LLM to derive appropriate aligner prompts to facilitate the final step of style transfer. The final prompt generated by APO can be found in Table 17.

3.4 Evaluation

The evaluation of our pipeline is dependent on the accuracy metric. Accuracy is defined individually

for each component.

$$\text{Accuracy} = \frac{\text{Number of retrieved GT}}{\text{Total number of data points}} \quad (1)$$

Number of retrieved GT: Total number of examples for which ground truth skin condition was present in the retrieved list of candidate skin conditions.

Total number of data points: Total number of examples for which ground truth skin condition is known. We skip all those examples for which ground truth is not known. In the validation we have 47 examples for which ground truth is known and a total of 56 examples.

Acknowledging the fact that a same skin condition can have multiple names, we implemented GPT-Eval as an evaluator to identify if two skin conditions are similar or not. Our evaluation strategy employs a rule-based approach to assess the similarity between two skin conditions, "A" and "B", according to four predefined rules (refer Table 18). These rules incorporate name identity, synonymity, common root condition, and shared effects and causes to determine similarity systematically. This method addresses the complex nature of dermatological conditions by providing a structured framework that considers linguistic, clinical, and etiological aspects of skin diseases.

4 Results

4.1 Retrieval Module

| Retrieval Strategies | Methods | Accuracy |
|----------------------------|------------|----------|
| <i>Context Independent</i> | First Derm | 0.468085 |
| | GPT-based | 0.595744 |
| <i>Context Dependent</i> | Naive CoT | 0.851063 |
| | Expert CoT | 0.744680 |

Table 2: Comparison of Retrieval Strategies and their Accuracy

The accuracy scores, as reported in Table 2, reveal significant insights into the efficacy of each strategy employed within our Retrieval Module. Firstly, we observed that Context-Independent Retrieval, which relies exclusively on image data, resulted in lower accuracy when compared to Context-Dependent strategies. This indicates that the absence of medical context limits GPT-4V’s ability to identify potential skin diseases accurately. Conversely, Context-Dependent Retrieval exhibited superior results. By incorporating medical

queries along with images, this method provides a richer context to GPT-4V, leading to more precise retrieval of potential skin conditions. It appears that the additional contextual data plays a pivotal role in enhancing the model’s diagnostic capabilities.

When comparing the two strategies within the Context-Dependent Retrieval, Naive CoT outperformed Expert CoT. This may initially seem counterintuitive, given that Expert CoT is grounded in medical guidelines, which one would expect to yield better results. However, our analysis suggests that the Naive CoT strategy’s ability to generate a broader range of potential candidates contributed to its higher accuracy. In contrast, the Expert CoT strategy, which employs differential diagnosis principles, likely eliminated some candidates during the retrieval phase, potentially leading to decreased accuracy.

From these observations, we hypothesize that differential diagnosis, while not as effective in the initial retrieval phase, may be better suited to the re-ranking phase of our diagnostic pipeline. The re-ranking phase requires a systematic evaluation to differentiate between closely related skin conditions, aligning with the differential diagnosis’s intrinsic nature. Therefore, the nuanced approach of systematically eliminating similar conditions could prove beneficial in the subsequent stage, where precision is paramount.

4.2 Re-Ranker Module

| Methods | Top-2 Accuracy | Top-1 Accuracy |
|------------------------------------|----------------|----------------|
| Naive CoT | 0.553191 | 0.425531 |
| Medical Guidelines (image+context) | 0.617021 | 0.531915 |
| Medical Guidelines (image only) | 0.553191 | 0.446808 |

Table 3: Comparison of Re-Ranker Strategies and their Accuracy

The re-ranker module is important in refining the initial list of potential diagnoses obtained from the retrieval module. The metrics used to evaluate the performance of our re-ranker module are:

- **Top-2 Accuracy:** This metric reflects the model’s ability to include the correct diagnosis within its top two predictions from the candidate conditions identified in the retrieval phase.

- **Top-1 Accuracy:** This is the precision with which the model identifies the correct diagnosis as its first and final choice from all possible conditions.

As illustrated in Table 3, the evaluation of our re-ranker strategies reveals several insights.

The Naive CoT and Medical Guidelines (image only) strategies exhibit comparable performance, with both Top-2 and Top-1 accuracies closely aligned. This suggests that even without the medical context, the model can leverage visual cues to a degree of effectiveness.

A notable increase in accuracy is observed with the use of Medical Guidelines alongside context (image + patient’s query). Incorporating the patient’s medical history and associated query, in conjunction with differential diagnosis techniques as outlined in Table 15, enhances the model’s discriminatory power. This aligns with our hypothesis that the systematic approach of differential diagnosis—filtering through similar skin conditions—proves more efficacious in the re-ranking phase.

4.3 MAC

| Methods | Accuracy |
|---------|----------|
| MG-GR | 0.53333 |
| MAC | 0.73333 |

Table 4: Re-Ranking - top 1 Accuracy using MAC. Here MG-GR is the (Medical Guidelines Grounded Re-Ranker).

The multi-agent conversation setup significantly outperforms the traditional top-1 re-ranking strategy, exhibiting a substantial improvement of nearly 20 percentage points. This enhancement was observed across 15 distinct examples where the number of potential solutions retrieved varied between three and five. We propose that the key mechanism driving this enhanced accuracy is the system’s critique-based conversational framework. Within this framework, each participating agent is subject to a rigorous process of critique and feedback from other agents. This collaborative interaction encourages continuous reassessment and refinement of each agent’s initial diagnoses and the evidence they present. Consequently, this iterative process likely contributes to more precise and reliable diagnostic outcomes, as each agent integrates insights gained

from the critiques to adjust and improve their reasoning and conclusions.

4.4 Aligner Module

| | DeltaBleu |
|------------|-----------|
| Before APO | 0.944723 |
| After APO | 2.737657 |

Table 5: DeltaBleu scores before and after Automatic Prompt Optimization (APO)

An important evaluation metric for the competition is the deltableu score. The "DeltaBLEU" score is a variation of the BLEU (Bilingual Evaluation Understudy) score, which is a widely used metric for evaluating the quality of text that has been machine-translated from one language to another. The BLEU score measures the correspondence between a machine’s output and that of a human, providing a quantitative assessment of translation accuracy. We leverage Automatic Prompt Optimization (APO) to learn a set of rules that bootstraps our prediction and align the responses. Table 5 shows the bleu score improvement by leveraging the rules learned by APO. The learned rules are mentioned in the Table 17.

- Before Alignment: "Based on the visual descriptions, it seems like the most probable condition is Chronic Eczema. I recommend applying topical steroids and moisturizers regularly for treatment."
- After Alignment: "Consider Eczema, which should manifest similarly on both sides. Treat it with regular use of moisturizers and topical steroids."
- Ground Truth: "Should be happening on both sides. Think of Eczema."

5 Discussion and Conclusion

In our study, we systematically explored the merits of various prompting strategies within an information retrieval-based dermatology diagnostic framework. By evaluating these strategies through the lens of accuracy metrics, we found that a naive Chain of Thought (CoT) strategy effectively simulates a retrieval module typical of information retrieval systems. This approach is adept at returning a sufficient number of candidate diagnoses, setting a foundational stage for further analysis. Our findings underscore the importance of including patient

history and contextual information in clinical dermatology to enhance diagnostic accuracy.

For the nuanced task of re-ranking diagnostic candidates, our research indicates that a more refined CoT strategy is necessary. Specifically, prompts that incorporate expert guidelines prove critical in conducting differential diagnoses, yielding superior performance in top-1 and top-2 diagnostic outcomes.

Furthermore, we introduce the novel Multi-Agent Critique (MAC) framework, which incorporates agent-based critique and feedback, and has the potential to perform differential diagnosis and refine its output using feedback.

6 Limitations and Future Work

Our current pipeline does not fully comply with stringent data protection regulations, such as the Health Insurance Portability and Accountability Act (HIPAA). Despite Azure’s availability of a HIPAA-compliant hosting option, our framework has not been fully aligned with these regulatory standards. The imperative to protect patient data privacy and ensure security is paramount in clinical applications. Our findings suggest that deploying a local model might offer a more privacy-centric approach. However, achieving satisfactory performance with local deployment necessitates further research and development. This limitation underscores the critical need to balance privacy considerations with technological efficacy, especially in the sensitive context of healthcare.

An additional dimension of our study pertains to the inherent variability in the performance of the prompting strategies, attributed to the high temperature setting utilized during GPT-4’s open-ended generation tasks. This element of randomness introduces inconsistencies in the model’s responses. We hypothesize that employing over-sampling techniques from GPT-4, coupled with self-consistency prompting, could mitigate these inconsistencies and enhance the overall effectiveness of the diagnostic process.

Furthermore, the MAC framework’s practical application presents challenges, notably in the seamless integration of inter-agent communication. Our observations point to instances where the system failed due to unexpected behaviors during these interactions, highlighting the complexities of implementing such a framework effectively. Additionally, the MAC study has been conducted on a small

set of 15 samples for which the number of retrieved candidates are in the range of 3 to 5 with the MAC system failing for number of candidates greater than 5. This was because of limited context length window with GPT-4 model. Additionally, each call to the agent is financially prohibitive as the number of candidates increase since more rounds of conversations are needed. Such challenges underscore the need for further research and development to refine and optimize the MAC framework for clinical diagnostic applications.

Additionally, given the challenging nature of the dataset with unclear/missing context for a lot of examples, correct bench marking cannot be assumed, but this study can serve as a potential lower bound of GPT-4V’s performance on the complex task of clinical dermatology.

References

- Wei Ba, Huan Wu, Wei W. Chen, Shu H. Wang, Zi Y. Zhang, Xuan J. Wei, Wen J. Wang, Lei Yang, Dong M. Zhou, Yi X. Zhuang, Qin Zhong, Zhi G. Song, and Cheng X. Li. 2022. [Convolutional neural network assistance significantly improves dermatologists’ diagnosis of cutaneous tumours using clinical images](#). *European Journal of Cancer*, 169:156–165.
- Joachim A Behar, Chengyu Liu, Kevin Kotzen, Kenta Tsutsui, Valentina DA Corino, Janmajay Singh, Marco AF Pimentel, Philip Warrick, Sebastian Zanzeder, Fernando Andreotti, et al. 2020. Remote health diagnosis and monitoring in the time of covid-19. *Physiological measurement*, 41(10):10TR01.
- Patricia E Benner, Christine A Tanner, and Catherine A Chesla. 2009. *Expertise in nursing practice: Caring, clinical judgment, and ethics*. Springer Publishing Company.
- Titus Josef Brinker, Achim Hekler, and et al. 2019. [A convolutional neural network trained with dermoscopic images performed on par with 145 dermatologists in a clinical melanoma image classification task](#). *European journal of cancer*, 111:148–154.
- Qihui Chen, Xinyue Hu, Zirui Wang, and Yi Hong. 2023a. [Medblip: Bootstrapping language-image pre-training from 3d medical images and texts](#).
- Zeming Chen, Alejandro Hernández Cano, Angelika Romanou, Antoine Bonnet, Kyle Matoba, Francesco Salvi, Matteo Pagliardini, Simin Fan, Andreas Köpf, Amirkeivan Mohtashami, Alexandre Sallinen, Alireza Sakhaeirad, Vinitra Swamy, Igor Krawczuk, Deniz Bayazit, Axel Marmet, Syrielle Montariol, Mary-Anne Hartley, Martin Jaggi, and Antoine Bosselut. 2023b. [Meditron-70b: Scaling medical pretraining for large language models](#).

- Angel Alfonso Cruz-Roa, John Edison Arevalo Ovalle, Anant Madabhushi, and Fabio Augusto González Osorio. 2013. A deep learning architecture for image representation, visual interpretability and automated basal-cell carcinoma cancer detection. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2013*, pages 403–410, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Andre Esteva, Brett Kuperl, Roberto A. Novoa, Justin M. Ko, Susan M. Swetter, Helen M. Blau, and Sebastian Thrun. 2017. [Dermatologist-level classification of skin cancer with deep neural networks](#). *Nature*, 542:115–118.
- Giacomo Frisoni, Alessio Cocchieri, Alex Presepi, Gianluca Moro, and Zaiqiao Meng. 2024. To generate or to retrieve? on the effectiveness of artificial contexts for medical open-domain question answering. *arXiv preprint arXiv:2403.01924*.
- Niharika Gouda and J Amudha. 2020. [Skin cancer classification using resnet](#). In *2020 IEEE 5th International Conference on Computing Communication and Automation (ICCCA)*, pages 536–541.
- Yu Gu, Robert Tinn, Hao Cheng, Michael Lucas, Naoto Usuyama, Xiaodong Liu, Tristan Naumann, Jianfeng Gao, and Hoifung Poon. 2020. [Domain-specific language model pretraining for biomedical natural language processing](#).
- Achim Hekler, Jochen Sven Utikal, Alexander H. Enk, Carola Berking, Joachim Klode, Dirk Schadendorf, Philipp Jansen, Cindy Franklin, Tim Holland-Letz, Dieter Krahl, Christof von Kalle, Stefan Fröhling, and Titus Josef Brinker. 2019. [Pathologist-level classification of histopathological melanoma images with deep neural networks](#). *European Journal of Cancer*, 115:79–83.
- Jiaming Ji, Boyuan Chen, Hantao Lou, Donghai Hong, Borong Zhang, Xuehai Pan, Juntao Dai, and Yaodong Yang. 2024. [Aligner: Achieving efficient alignment through weak-to-strong correction](#). *arXiv preprint arXiv:2402.02416*.
- Y.Q. Jiang, J.H. Xiong, H.Y. Li, X.H. Yang, W.T. Yu, M. Gao, X. Zhao, Y.P. Ma, W. Zhang, Y.F. Guan, H. Gu, and J.F. Sun. 2020. [Recognizing basal cell carcinoma on smartphone-captured digital histopathology images with a deep neural network](#). *British Journal of Dermatology*, 182(3):754–762.
- Omar Khattab and Matei Zaharia. 2020. Colbert: Efficient and effective passage search via contextualized late interaction over bert. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*, pages 39–48.
- Chunyuan Li, Cliff Wong, Sheng Zhang, Naoto Usuyama, Haotian Liu, Jianwei Yang, Tristan Naumann, Hoifung Poon, and Jianfeng Gao. 2023a. [Llava-med: Training a large language-and-vision assistant for biomedicine in one day](#).
- Jian Li, Xi Chen, Weizhi Liu, Li Wang, Yingman Guo, Mingke You, Gang Chen, and Kang Li. 2023b. [One is not enough: Multi-agent conversation framework enhances rare disease diagnostic capabilities of large language models](#).
- Renqian Luo, Liai Sun, Yingce Xia, Tao Qin, Sheng Zhang, Hoifung Poon, and Tie-Yan Liu. 2022. [BioGPT: generative pre-trained transformer for biomedical text generation and mining](#). *Briefings in Bioinformatics*, 23(6). Bbac409.
- Sean E Mangion, Tai A Phan, Samuel Zagarella, David Cook, Kirtan Ganda, and Howard I Maibach. 2023. Medical school dermatology education: a scoping review. *Clinical and Experimental Dermatology*, 48(6):648–659.
- Daniel McDuff, Mike Schaeckermann, Tao Tu, Anil Palepu, Amy Wang, Jake Garrison, Karan Singhal, Yash Sharma, Shekoofeh Azizi, Kavitaulkarni, et al. 2023. [Towards accurate differential diagnosis with large language models](#). *arXiv preprint arXiv:2312.00164*.
- Jing Miao, Charat Thongprayoon, Wisit Cheungpasitporn, and Lynn D Cornell. 2024. [Performance of GPT-4 Vision on kidney pathology exam questions](#). *American Journal of Clinical Pathology*, page aqa030.
- Prakamy Mishra, Zonghai Yao, Parth Vashisht, Feiyun Ouyang, Beining Wang, Vidhi Dhaval Mody, and Hong Yu. 2024. [Synfac-edit: Synthetic imitation edit feedback for factual alignment in clinical summarization](#).
- Harsha Nori, Nicholas King, Scott Mayer McKinney, Dean Carignan, and Eric Horvitz. 2023a. [Capabilities of gpt-4 on medical challenge problems](#).
- Harsha Nori, Yin Tat Lee, Sheng Zhang, Dean Carignan, Richard Edgar, Nicolo Fusi, Nicholas King, Jonathan Larson, Yuanzhi Li, Weishung Liu, Renqian Luo, Scott Mayer McKinney, Robert Osazuwa Ness, Hoifung Poon, Tao Qin, Naoto Usuyama, Chris White, and Eric Horvitz. 2023b. [Can generalist foundation models outcompete special-purpose tuning? case study in medicine](#).
- Dimitrios P. Panagoulas, Evridiki Tsourelis-Nikita, Maria Virvou, and George A. Tsihrantzis. 2024. [Dermacen analytica: A novel methodology integrating multi-modal large language models with machine learning in tele-dermatology](#).
- Archiki Prasad, Peter Hase, Xiang Zhou, and Mohit Bansal. 2022. Grips: Gradient-free, edit-based instruction search for prompting large language models. *arXiv preprint arXiv:2203.07281*.
- Reid Pryzant, Dan Iter, Jerry Li, Yin Tat Lee, Chenguang Zhu, and Michael Zeng. 2023. [Automatic prompt optimization with "gradient descent" and beam search](#). *arXiv preprint arXiv:2305.03495*.

- Sudha Rao and Joel Tetreault. 2018. Dear sir or madam, may i introduce the gyafc dataset: Corpus, benchmarks and metrics for formality style transfer. *arXiv preprint arXiv:1803.06535*.
- Stephen Robertson, Hugo Zaragoza, et al. 2009. The probabilistic relevance framework: Bm25 and beyond. *Foundations and Trends® in Information Retrieval*, 3(4):333–389.
- Akhilesh Kumar Sharma, Shamik Tiwari, Gaurav Aggarwal, Nitika Goenka, Anil Kumar, Prasun Chakrabarti, Tulika Chakrabarti, Radomir Gono, Zbigniew Leonowicz, and Michał Jasiński. 2022. Dermatologist-level classification of skin cancer using cascaded ensembling of convolutional neural network and handcrafted features based deep neural network. *IEEE Access*, 10:17920–17932.
- Karan Singhal, Shekoofeh Azizi, Tao Tu, S. Sara Mahdavi, Jason Wei, Hyung Won Chung, Nathan Scales, Ajay Tanwani, Heather Cole-Lewis, Stephen Pfohl, Perry Payne, Martin Seneviratne, Paul Gamble, Chris Kelly, Nathaneal Scharli, Aakanksha Chowdhery, Philip Mansfield, Blaise Aguerre y Arcas, Dale Webster, Greg S. Corrado, Yossi Matias, Katherine Chou, Juraj Gottweis, Nenad Tomasev, Yun Liu, Alvin Rajkomar, Joelle Barral, Christopher Semturs, Alan Karthikesalingam, and Vivek Natarajan. 2022. Large language models encode clinical knowledge.
- Karan Singhal, Shekoofeh Azizi, Tao Tu, S Sara Mahdavi, Jason Wei, Hyung Won Chung, Nathan Scales, Ajay Tanwani, Heather Cole-Lewis, Stephen Pfohl, et al. 2023a. Large language models encode clinical knowledge. *Nature*, 620(7972):172–180.
- Karan Singhal, Tao Tu, Juraj Gottweis, Rory Sayres, Ellery Wulczyn, Le Hou, Kevin Clark, Stephen Pfohl, Heather Cole-Lewis, Darlene Neal, et al. 2023b. Towards expert-level medical question answering with large language models. *arXiv preprint arXiv:2305.09617*.
- Dan Su, Mostofa Patwary, Shrimai Prabhumoye, Peng Xu, Ryan Prenger, Mohammad Shoeybi, Pascale Fung, Anima Anandkumar, and Bryan Catanzaro. 2022. Context generation improves open domain question answering. *arXiv preprint arXiv:2210.06349*.
- Mingxu Tao, Dongyan Zhao, and Yansong Feng. 2024a. Chain-of-discussion: A multi-model framework for complex evidence-based question answering.
- Mingxu Tao, Dongyan Zhao, and Yansong Feng. 2024b. Chain-of-discussion: A multi-model framework for complex evidence-based question answering. *arXiv preprint arXiv:2402.16313*.
- Tao Tu, Anil Palepu, Mike Schaekermann, Khaled Saab, Jan Freyberg, Ryutaro Tanno, Amy Wang, Brenna Li, Mohamed Amin, Nenad Tomasev, et al. 2024. Towards conversational diagnostic ai. *arXiv preprint arXiv:2401.05654*.
- Anca Loredana Udriștoiu, Ariana Elena Stanca, Alice Elena Ghenea, Corina Maria Vasile, Mihaela Popescu, Stefan Udriștoiu, Andreea Valentina Iacob, Ștefan Cristian Castravete, Lucian Gheorghe Gruionu, and Gabriel Gruionu. 2020. Skin diseases classification using deep learning methods. *Current Health Sciences Journal*, 46:136 – 140.
- Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Shaokun Zhang, Erkang Zhu, Beibin Li, Li Jiang, Xiaoyun Zhang, and Chi Wang. 2023. Auto-gen: Enabling next-gen llm applications via multi-agent conversation framework. *arXiv preprint arXiv:2308.08155*.
- Zhichao Yang, Zonghai Yao, Mahbuba Tasmin, Parth Vashisht, Won Seok Jang, Feiyun Ouyang, Beining Wang, Dan Berlowitz, and Hong Yu. 2023. Performance of multimodal gpt-4v on usmle with image: Potential for imaging diagnostic support with explanations. *medRxiv*.
- Zonghai Yao, Ahmed Jaafar, Beining Wang, Yue Zhu, Zhichao Yang, and Hong Yu. 2023. Do physicians know how to prompt? the need for automatic prompt optimization help in clinical note generation. *arXiv preprint arXiv:2311.09684*.
- Zonghai Yao and Hong Yu. 2021. Improving formality style transfer with context-aware rule injection. *arXiv preprint arXiv:2106.00210*.
- Wenhao Yu, Dan Iter, Shuohang Wang, Yichong Xu, Mingxuan Ju, Soumya Sanyal, Chenguang Zhu, Michael Zeng, and Meng Jiang. 2022. Generate rather than retrieve: Large language models are strong context generators. *arXiv preprint arXiv:2209.10063*.
- Yunxiang Zhang, Muhammad Khalifa, Lajanugen Logeswaran, Moontae Lee, Honglak Lee, and Lu Wang. 2023. Merging generated and retrieved knowledge for open-domain qa. *arXiv preprint arXiv:2310.14393*.
- Juexiao Zhou, Xiaonan He, Liyuan Sun, Jiannan Xu, Xiuying Chen, Yuetan Chu, Longxi Zhou, Qingyu Liao, Bin Zhang, and Xin Gao. 2023. Skingpt-4: An interactive dermatology diagnostic system with visual large language model.

A Appendix

A.1 Multi-Agent Conversation (MAC) Case Study

In this section, we present an exemplar case study of a debate facilitated by the Multi Agent Chat (MAC) system. The dialogues exemplified in Table 7 to Table 12 illustrate the dynamic interaction between diagnostic specialists and admin agents within our MAC framework.

The ensuing discussion is prefaced with prompts which have successfully generated the anticipated

outcomes, showcasing the MAC system’s adeptness. We provide a series of prompts, as detailed in Table 6 and the expert guidelines to ensure that the communication trajectory remains aligned with the system’s strategic objectives.

Through the presented case study, we aim to elucidate the capabilities of the MAC system’s in the context of clinical diagnosis and the efficacy of its prompts in steering the group chat among various agents to achieve coherent, goal-oriented dialogue.

Custom Agent (Diagnostic Specialist) Prompt is to guide the model in adopting the role of a diagnostic specialist. It aims to facilitate the identification of salient features of a skin condition based on clinical observations that align with the designated skin disease. Utilizing these features and the provided details regarding the disease progression, the model is tasked with generating evidence to substantiate the diagnosis of the specified skin disease.

Coordinator Prompt is utilized to guide the model to play the role of a coordinator, orchestrating the conversation among the agents, collecting the generated evidences and critiques, consolidating them and passing them to the Admin agent for further analysis and diagnosis. This agent is essential to ensure smooth transition in between the agents.

Admin Prompt is designed to instruct the model to play the role of Admin, the admin is a head doctor, who first meticulously assesses the quality of generated evidences and critiques. Then the Admin is tasked with guiding the relevant specialists to refine their evidence in light of the critiques received, thereby enhancing the robustness of the diagnostic case. This iterative process of evaluation and refinement underscores the Admin’s pivotal role in ensuring the accuracy and integrity of the final diagnosis, highlighting the significance of expert oversight in collaborative medical diagnostics.

A.2 Guidelines and Instructions for Multi-Agent Chat

The instructions and expert guidelines are most crucial for the MAC module, since it is a comprehensive briefing of the objectives that the system aims to fulfill. This prompt delineates not only the sequence of actions requisite for task execution but also the intricacies of inter-agent transitions,

Diagnostic Specialist Prompt

As ‘Rick’, you are a medical practitioner specializing in dermatology. You are provided with an image description and assigned a specific skin condition, your role is to tune the image description to match with the disease.

Once done you can use the image description to generate a detailed report providing evidence that supports this diagnosis. Afterwards critique each of the other probable diseases by explaining why they do not fit the case study as well as your assigned diagnosis does. Ensure clarity and comprehensiveness in your analysis and critiques.

Coordinator Prompt

As the ‘Coordinator’, your primary responsibility is to oversee the diagnostic process. You will receive a clinical observation of a skin disease, a case study along with a list of probable diseases. Your task is to assign each Diagnostic Specialist a unique probable disease to advocate for, based on the provided details. You’ll ensure that each specialist receives all necessary information to perform their analysis effectively. Finally, gather the tuned image descriptions, evidences and critiques from the specialists and present them to the Admin for final evaluation. Your role is crucial for maintaining efficient communication and organization among the specialists.

Admin Prompt

As the ‘Admin’, your objective is to evaluate the evidence and critiques provided by the Diagnostic Specialists majorly based and aligned to the image description since solely depending on the case study can be tricky to determine the most probable disease for a given case study. Initiate your process by assessing the quality of each critique. Seek consensus among the critiques to strengthen the evidence for a particular diagnosis. You may need to instruct Diagnostic Specialists to refine their evidence based on your findings. Through a structured discussion with the Coordinator and the Diagnostic Specialists, lead the team towards agreeing on a final, most suitable diagnosis for the case study.

Table 6: Prompts for Multi Agent Chat.

thereby charting the entirety of procedural flow. Furthermore, it encompasses a set of critical guidelines mandating adherence to principles of clarity and precision, alongside the seamless exchange of information among pertinent agents. These directives are imperative to avoid miscommunication and ensure that all interactions remain aligned with the task’s end goals. The Task Prompt is furnished to the GroupChatManager, serving as the catalyst for activating dialogues among specialized agents within the framework. This structured approach is pivotal in harmonizing the collective efforts of diverse agents, thus optimizing the overall functionality and efficacy of the MAC system.

A.3 Multi-Agent Chat - Example

The entire chat is accessible in our GitHub repository for reference. In this section, we present selected excerpts from the Multi-Agent Chat to illustrate the flow and demonstrate the system’s

capabilities:

Set-up of the Chat:

Patient Query (with the context of disease progression): "The skin condition, as shown in the images, presents widespread erythematous patches with violaceous hues across the leg. The patient has multiple crusted plaques and erosions, with sizes varying from a few millimeters to several centimeters. Some lesions have a serpiginous border, suggesting an active edge. The skin's texture looks lichenified in some places, indicating chronicity, and scaling is evident across various regions, signaling some level of dryness and exfoliation. Some patches have merged, forming a larger area of affected skin. Signs of excoriations are present, most likely due to itching, and scattered pustules can also be observed."

Probable Diseases: "Prurigo nodularis, Chronic eczema, Psoriasis, Lichen simplex chronicus, Allergic or irritant contact dermatitis"

Clinical Observation of the Skin Condition: "The skin condition, as shown in the images, presents widespread erythematous patches with violaceous hues across the leg. The patient has multiple crusted plaques and erosions, with sizes varying from a few millimeters to several centimeters. Some lesions have a serpiginous border, suggesting an active edge. The skin's texture looks lichenified in some places, indicating chronicity, and scaling is evident across various regions, signaling some level of dryness and exfoliation. Some patches have merged, forming a larger area of affected skin. Signs of excoriations are present, most likely due to itching, and scattered pustules can also be observed."

Ground Truth: Chronic Eczema

Excerpt 1 - GroupChat Initialization: The group chat is initiated by the admin. The task, meticulously crafted for our use case, is provided to the chat manager who then follows the outlined steps.

Excerpt 2 - Evidence & Critiques: After the coordinator assigns probable diseases to each diagnostic specialist, they are sequentially called to generate supportive evidence and critique other possibilities.

Excerpt 3 - Consolidated Evidence & Critiques: Once the diagnostic specialists have processed their assigned diagnoses, the coordinator gathers and consolidates the evidence and critiques

for each disease.

Excerpt 4 - Admin Refinement Instructions:

This consolidated evidence and critiques are reviewed by the admin, who assesses them and may request further information. The admin provides instructions for specialists to refine their evidence, aiming for a more accurate diagnosis.

Excerpt 5 - Enhanced Evidence: Based on the refinement instructions from the Admin, the designated agent is tasked with enhancing their evidence to better support their diagnosis.

Excerpt 6 - Final Diagnosis: Based on the refined evidence, the admin determines the most relevant final diagnosis.

A.4 Dermatology Guidelines

When a dermatologist evaluates a skin condition, they typically follow a systematic approach that involves several areas.

- **Patient History:** Look at the "User Query" to extract relevant context that will help in accurate diagnosis of skin conditions.
- **Visual Inspection:** The initial step involves a thorough visual examination of the affected area.

For visual inspection, the dermatologist looks at the following features and creates a list of possible skin conditions that match the visual features.

1. **Size:** What is the size of the skin lesions? Is it small or large?
2. **Shape:** What is the shape of the lesions?
3. **Color:** What is the color of the skin lesions?
4. **Location:** Where is the skin lesion or rash located?
5. **Distribution Pattern:** What is the distribution pattern, is it localized or widespread?
6. **Existence of symmetry:** Are the lesions symmetric?
7. **Borders:** Do the edges of the lesion appear sharp, irregular, or blurred?
8. **Elevation:** Is the lesion flat, raised, or depressed below the skin surface?
9. **Texture:** Does the surface quality look smooth/scaly/rough/soft/hard.

Pattern Recognition: Dermatologists are trained in recognizing patterns that certain skin conditions commonly present. These patterns, combined with the other collected information, help in forming a

Excerpt - 1:

****Chat Initialisation with the Admin:****

Admin (to chat_manager):

The objective is to facilitate a systematic discussion to diagnose a patient's skin condition, based on the clinical observation of the skin condition (The skin condition, as shown in the images, presents widespread erythematous patches with violaceous hues across the leg. The patient has multiple crusted plaques and erosions, with sizes varying from a few millimeters to several centimeters. Some lesions have a serpiginous border, suggesting an active edge. The skin's texture looks lichenified in some places, indicating chronicity, and scaling is evident across various regions, signaling some level of dryness and exfoliation. Some patches have merged, forming a larger area of affected skin. Signs of excoriations are present, most likely due to itching, and scattered pustules can also be observed.), case study (Please help take a look, what kind of skin disease is this? Suffering from the disease for more than 10 years. It is recurrent and is very itchy! It happens wherever I scratch in some places.) and a list of probable diseases (['prurigo nodularis', 'chronic eczema', 'psoriasis', 'lichen simplex chronicus', 'allergic or irritant contact dermatitis']). Give equal emphasis to the clinical observation as well as the symptoms and the disease progression to ensure accuracy. Assume that no further information or biopsy will be conducted.

""

****Process Overview:****

""

""

****Detailed Steps:****

""

""

****Rules for Effective Process Management:****

""

****Once "TERMINATE" has been returned no further conversation must happen****

Table 7: Excerpt 1 : Group chat initialisation

preliminary diagnosis.

Differential Diagnoses: Based on the evaluation, the dermatologist will consider a list of possible conditions (differential diagnoses) and rule them out one by one, based on the evidence and test results.

Excerpt - 2:

Rick (to chat_manager):

As Diagnostic Specialist Rick,

Assigned Diagnosis: Prurigo Nodularis.

Supporting Evidence for Prurigo Nodularis: The indications of chronic scratching or rubbing like lichenification suggest that the rash could be Prurigo Nodularis. The patient's description of the condition as being very itchy and recurrent over a span of more than a decade also aligns with this diagnosis. In addition, the presence of scattered pustules can also be seen in cases of Prurigo Nodularis.

Critiques for Other Diseases:

1. **Chronic Eczema:** This condition also presents an itchy rash that can become lichenified from chronic scratching, but usually has a more defined pattern of eruption that is not described here.
2. **Psoriasis:** While this condition presents chronic plaques, they usually exhibit a characteristic silvery scale. No such description is provided here.
3. **Lichen Simplex Chronicus:** While this condition is characterized by lichenification, it generally affects a specific region rather than having a widespread distribution as described here.
4. **Allergic or Irritant Contact Dermatitis:** These conditions generally present rapidly after contact with an offending substance, which does not fit the chronic, decade-long presentation described here.

****CALL COORDINATOR** to validate completion.**

Table 8: Excerpt 2: Evidence & Critiques by Diagnostic Specialist Rick.

Excerpt - 3:

Coordinator (to chat_manager):

Thank you for your findings, Michael.

As the Coordinator, I acknowledge the completion of Michael's analysis regarding 'allergic or irritant contact dermatitis'.

****Compiling Findings****

I'll now compile and categorize the generated evidences and critiques for each potential disease:

1. Prurigo Nodularis

- **Supporting Evidence:** Chronic scratching or rubbing and recurrent itchiness over a decade. Presence of scattered pustules.

- **Consolidated Critiques:** More defined pattern of eruption for Chronic Eczema is missing. No silver scales as in Psoriasis. More localized presentation expected as in Lichen Simplex Chronicus. No acute flare-ups post exposure as in Contact Dermatitis.

2. Chronic Eczema

- **Supporting Evidence:** Widespread and merging erythematous patches, lichenified skin at some regions and itchiness wherever scratched.

- **Consolidated Critiques:** Prurigo Nodularis typically presents hard, itchy lumps. Silvery scales typical to Psoriasis missing. Lichen Simplex Chronicus is usually localized.

3. Psoriasis

- **Supporting Evidence:** Widespread plaques, lichenified skin and scaling.

- **Consolidated Critiques:** In Prurigo Nodularis itching is more generalized. No typical eczema characteristics like weeping, oozing. Lichen Simplex Chronicus usually is localized.

4. Lichen Simplex Chronicus

- **Supporting Evidence:** Lichenification of skin due to chronic rubbing, intense itchiness, especially in localized areas, plus the serpiginous border.

- **Consolidated Critiques:** Larger areas of affected skin unlike Prurigo Nodularis. No oozing or crusting unlike Chronic Eczema. Silvery scale of Psoriasis missing. Contact Dermatitis usually presents acute symptoms post exposure.

5. Allergic or Irritant Contact Dermatitis

- **Supporting Evidence:** Chronic dryness, exfoliation, and itching response to certain irritants.

- **Consolidated Critiques:** Doesn't describe nodules typical to Prurigo Nodularis. Not enough weeping and crusting for Chronic Eczema. No silvery scaly plaques like Psoriasis. Lichen Simplex Chronicus is usually localized.

As the Coordinator, I present the compiled evidence to the Admin for a final evaluation on this patient's skin condition.

Table 9: Excerpt 3: Consolidated Evidences & Critiques

Excerpt - 4:

Admin (to chat_manager):

As the Admin,

Thank you, Coordinator, for compiling the findings.

Firstly, it's noticeable that the given evidences for "Allergic or Irritant Contact Dermatitis" and "Psoriasis" are relatively weak compared to other diagnoses. I agree with the critiques that the chronicity and absence of typical defining characteristics such as acute flare-ups in the case of dermatitis and silvery scales in the case of psoriasis make these diagnoses less likely.

Focusing on "Prurigo Nodularis", the evidence is plausible but the image description does not mention the typical nodular lumps that are characteristic of this disease.

"Lichen Simplex Chronicus" has supportive evidence that strongly aligns with chronic itching and lichenification, yet the widespread presence of the disease contradicts the usual localized occurrence of this condition.

"Chronic Eczema" also aligns well with the majority of described symptoms, including itching, chronicity, and lichenification. The absence of strong objections in the critiques and aforementioned aligning symptoms gives strength to this diagnosis.

Considering all compiled evidence and critiques, it seems that **Chronic Eczema** might be the most likely diagnosis given the available information. However, some uncertainty remains due to overlapping symptoms with other conditions and the lack of additional diagnostic tests.

Diagnostic Specialist Sam, I would like you to enhance your evidence for 'Chronic Eczema' in light of the critiques provided by other specialists. Please refer to the following critiques and **provide more specific details that distinguish Chronic Eczema from other conditions:**

1. Prurigo Nodularis: **Your analysis could benefit from addressing whether or not the absence of nodules is indeed conclusive evidence against this diagnosis.**
2. Lichen Simplex Chronicus: **Can you further explain the distinguishing factors between these two conditions? Specifically, consider the details regarding distribution and impact of itch-induced scratching.**
3. Identification or ruling out of 'Allergic or Irritant Contact Dermatitis': **Please provide more info that can make this differentiation clearer.**

Table 10: Excerpt 4: Admin Refinement Instructions

Excerpt - 5:

Sam (to chat_manager):

As Diagnostic Specialist Sam,

[Enhanced Evidence for 'Chronic Eczema'](#):

To revisit and strengthen my diagnosis supporting Chronic Eczema, let's address the critiques:

1. [Absence of nodules in Prurigo Nodularis](#): Prurigo Nodularis is characterized by hard, itchy nodules which may be paired with lichenification because of chronic scratching. However, such nodules are not explicitly reported in this clinical presentation. Instead, we note widespread erythematous patches with various sizes and serpiginous borders, a pattern more consistent with chronic eczema.

2. [Distinguishing features between Chronic Eczema and Lichen Simplex Chronicus](#): Though both conditions show lichenification due to chronic scratching, they do have differing behaviors. Lichen Simplex Chronicus usually exhibits itself in one or two specific regions of the body, whereas Chronic Eczema can affect larger, more widespread areas as described in this clinical case. Thus, the widespread distribution here lends more credence to a diagnosis of Chronic Eczema rather than the typically localized Lichen Simplex Chronicus.

3. [Differentiating between Chronic Eczema and Allergic or Irritant Contact Dermatitis](#): Contact Dermatitis generally surfaces as an acute flare-up following exposure to a particular substance and often resolves once the irritant or allergen is avoided, whereas Chronic Eczema's cause is multifactorial - influenced not only by external irritants but also by internal factors, such as the patient's immune response. Furthermore, Chronic Eczema exhibits a distinctive pattern of flares and subsiding inflammation over time. This history of enduring for over a decade and recurrent nature of the skin condition directs more towards Chronic Eczema.

Overall, despite sharing common symptoms like itching and skin alterations with the mentioned conditions, Chronic Eczema appears to fit best given the specifics of the condition's distribution and chronicity.

Table 11: Excerpt 5: Enhanced Evidence

Excerpt - 6:

Admin (to chat_manager):

As the Admin,

Thank you, Sam, for the enhanced evidence and Coordinator for your facilitation.

Taking into account the evidences and critiques from all Diagnostic Specialists and the enhanced evidence provided by Sam, I conclude that in the absence of any further diagnostic tests or additional information, the most accurate diagnosis among the provided probable diseases for this case is [Chronic Eczema](#). The patient's long term history, reported symptoms like itching wherever the patient scratches, the recurrent nature of the condition, and clinical information like widespread erythematous patches all point towards Chronic Eczema.

Please ****TERMINATE**** the conversation, Coordinator.

Table 12: Excerpt 6: Final Diagnosis

| | |
|--|---|
| <p>Naive CoT</p> | <p>Act as an expert dermatologist. You are provided with a dermatology case with images. You are tasked to create a list of possible skin conditions for the given case.</p> <p><i>Instructions:</i></p> <p>1: Look at the images and the medical query and see what relevant information you can extract from the medical query that can be useful in diagnosis. 2: Create a possible list of skin conditions.</p> |
| <p>Medical Guidelines based CoT</p> | <p>You are provided with a dermatology case with images.</p> <p>For this case, you are provided with some images and additional user query. You are asked to give a diagnosis for this scenario. Act as a dermatologist. Refer the guidelines below and follow the guidelines to generate the diagnosis.</p> <p><i>Guidelines:</i></p> <p>When a dermatologist evaluates a skin condition, they typically follow a systematic approach that involves several key steps.</p> <p><i>Visual Inspection:</i> The initial step involves a thorough visual examination of the affected area. The dermatologist looks at the:</p> <ol style="list-style-type: none"> 1. Size 2. Shape 3. Color: The color (red, brown, black, blue, white) and whether it's uniform. 4. Location of the lesion or rash. 5. Distribution Pattern (localized/widespread) 6. Existence of symmetry (yes or no) 7. Borders: The edges of the lesion—are they sharp, irregular, or blurred? 8. Elevation: Whether the lesion is flat, raised, or depressed below the skin surface. 9. Texture: The surface quality (smooth, scaly, rough, soft, hard). 10. Pattern Recognition: Dermatologists are trained in recognizing patterns that certain skin conditions commonly present. These patterns, combined with the other collected information, help in forming a preliminary diagnosis. 11. Consideration of Differential Diagnoses: Based on the evaluation, the dermatologist will consider a list of possible conditions (differential diagnoses), and rule them out one by one, based on the evidence and test results. 13. Create a list of possible candidates after the above steps. |

Table 13: CoT based Retrieval

Act as an expert dermatologist.
You are provided with a dermatology case with images and associated medical query. You are tasked to choose the most probable skin condition from the set of candidates.

Medical Query:
<query>

Candidates:
<candidates>

Instructions:

- 1: Look at the images and the medical query and see what relevant information you can extract from the medical query that can be useful in diagnosis.
- 2: Give a score to each candidate skin condition in the range of 1 - 10 with 1 being the least probable and 10 being the most probable.
- 3: Choose a single most probable disease. If there is a tie in scores for the most probable conditions, pick a single skin condition between those candidates at random and return.

Table 14: Naive CoT for Re-Ranker

Act as an expert dermatologist. You are provided with a dermatology case. For this case, you are provided with some images a user query and list of candidates.

User Query: *<query>*

Candidates: *<candidates>*

Guidelines:

When a dermatologist evaluates a skin condition, they typically follow a systematic approach that involves several areas.

Patient History: Look at the "User Query" to extract relevant context that will help in accurate diagnosis of skin conditions.

Visual Inspection: The initial step involves a thorough visual examination of the affected area.

For visual inspection, the dermatologist looks at the following features and for each, the dermatologist creates a list of possible skin conditions that show such visual features.

1: Size: What is the size of the skin lesions? Is it small or large?

2: Shape: What is the shape of the lesions?

3: Color: What is the color of the skin lesions?

4: Location: Where is the skin lesion or rash located?

5: Distribution Pattern: What is the distribution pattern, is it localized or widespread?

6: Existence of symmetry: Are the lesions symmetric?

7: Borders: Do the edges of the lesion appear sharp, irregular, or blurred?

8: Elevation: Is the lesion is flat, raised, or depressed below the skin surface?

9: Texture: Does the surface quality looks smooth/scaly/rough/soft/hard.

Pattern Recognition: Dermatologists are trained in recognizing patterns that certain skin conditions commonly present. These patterns, combined with the other collected information, help in forming a preliminary diagnosis.

Differential Diagnoses: Based on the evaluation, the dermatologist will consider a list of possible conditions (differential diagnoses) and rule them out one by one, based on the evidence and test results.

Instructions:

step 1: Evaluate the medical images based on Visual Inspection Guidelines.

step 2: Evaluate the medical query as Patient History Guidelines.

step 3: Create a case summary using information extracted at step 1 and step 2.

step 4: For each candidate skin condition present in the list of Candidates, give a score on a scale of 1 - 10 (where 1 is the least probable and 10 is the most probable) that describes how likely is the given skin condition as a diagnosis for the case summary.

step 5: Return the two most probable skin candidates based on scores obtained at step 4.

Table 15: Expert Guidelines Grounded CoT with Context for Re-Ranker

You are provided with a dermatology case. For this case, you are provided with some images and list of possible candidates.

Candidates: <*candidates*>

Visual Inspection Guidelines: The initial step involves a thorough visual examination of the affected area. The dermatologists keep a track of 10 visual features.

- 1: Size: What is the size of the skin lesions? Is it small or large?
- 2: Shape: What is the shape of the lesions?
- 3: Color: What is the color of the skin lesions?
- 4: Location: Where is the skin lesion or rash located?
- 5: Distribution Pattern: What is the distribution pattern, is it localized or widespread?
- 6: Existence of symmetry: Are the lesions symmetric?
- 7: Borders: Do the edges of the lesion appear sharp, irregular, or blurred?
- 8: Elevation: Is the lesion is flat, raised, or depressed below the skin surface?
- 9: Texture: Does the surface quality looks smooth/scaly/rough/soft/hard.
- 10: Pattern Recognition: Dermatologists are trained in recognizing patterns that certain skin conditions commonly present. These patterns, combined with the other collected information, help in forming a preliminary diagnosis.

Act as a dermatologist and follow the instructions below:

Instructions:

step 1: For the given images, use the guidelines and generate a visual description.

step 2: For each candidate in the "Candidates", generate the visual description that describes the candidate disease. Also mention distinguishing features based on visual guidelines. Include features like shape, colours, lesion type and area of localization to create a visual description for the disease.

step 3: Compare the visual description which was generated for each candidate skin condition at step b with the image description generated at step a. Give a score in the range of 1 to 10 with 1 being the lowest match and 10 being the highest match.

step 4: Choose the most probable candidate which has the highest score with the images based on step 3.

Table 16: Expert Guidelines Grounded CoT without Context for Re-Ranker

Rules:**1. Simplify and Be Direct**

- Example: “The condition is Chronic Eczema.”

- Explanation: Human expert responses tend to be direct and use simpler language. Avoid overly complex explanations and aim for straightforward answers directly addressing the patient’s inquiry.

2. Diagnosis Confirmation

- Example: “Your diagnosis is a Myxoid Cyst based on the clear image provided.”

- Explanation: Include statements that confirm the diagnosis confidently, as seen in responses like “Chronic Eczema.” or “It is myxoid cyst.” Use assertive language to convey confidence in your diagnosis.

3. Detail Symptom Correlation

- Example: “The semi-spherical cyst near the end of your thumb, as described, leads to a diagnosis of Myxoid Cyst.”

- Explanation: Explicitly connect the diagnosis with observed symptoms or test results when applicable, similar to the detailed descriptions in some valid responses. This helps patients understand why a particular diagnosis is made.

4. Incorporate Treatment Options Clearly

- Example: “For Psoriasis, I recommend oral capsules such as glycyrrhizic acid glycosides, along with transfer factors.” - Explanation: When suggesting treatments, mention specific medications or procedures clearly and concisely, as observed in responses with high completeness. If possible, explain the purpose of each treatment briefly.

5. Mention Commonality or Prevalence

- Example: “Chronic Eczema is quite common and effectively manageable with the right treatment.”

- Explanation: If applicable, include a brief note on how common the condition is or any relevant statistical information that could reassure the patient or provide context, akin to how some expert responses include prevalence information.

6. Use Patient-Friendly Language

- Example: “Based on the photo you provided, it looks like you have a Myxoid Cyst, which is a fluid-filled lump that’s not harmful.”

- Explanation: Ensure the language used is patient-friendly, avoiding unnecessary medical jargon that could confuse the patient. When medical terms are unavoidable, consider providing a brief, simple explanation.

7. Personalization and Empathy

- Example: “I understand that dealing with Chronic Eczema can be frustrating. Regular moisturizing and the treatments we’ve discussed should offer relief.” - Explanation: Whenever possible, personalize the response to the patient’s situation. Display empathy to make your responses feel more human and less robotic.

Table 17: Automatic Prompt Optimization (APO) Rules

Rules:

1: Skin condition A is similar to B if they have same name.

2: Skin condition A is similar to B if B is also known by the name A.

3: Skin condition A is similar to B if both are part of the same root skin condition.

Example: Herpetic Eczema and seborrheic eczema are similar since they have same root, Eczema.

4: Skin condition A is similar to B if they are both have the same effect and share a common cause.

Table 18: Evaluation Guidelines Rules