IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

AIA review number
PGR2025-00075

OneSource Solutions International, Inc. ("OSSI"),
Petitioner,

v.

Hippocratic AI, Inc.,
Patent Owners.

Case No.
U.S. Patent No. 12,142,371 B1

PETITION FOR *POST-GRANT* REVIEW
OF U.S. PATENT NO. 12,142,371 B1 ("the '371 Patent")

TABLE OF CONTENTS

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE 1
I. INTRODUCTION 4
2. Summary of Prior Art Relevance 5
3. Petition Structure and Methodology 6
II. STANDING TO FILE 7
A. Identification of Real Parties-in-Interest 7
B. Identification of Related Matters 7
III. IDENTIFICATION OF CHALLENGED CLAIMS 8
IV. OVERVIEW OF THE '371 PATENT 9
A. File Wrapper Analysis 10
V. STATEMENT OF PRECISE RELIEF REQUESTED 11
VI. CLAIM-BY-CLAIM ANALYSIS 13
Claim 1 13
Claim 2 21
Claim 3 24
Claim 4 28
Claim 5 31
Claim 6 34
Claim 7 36
Claim 8 39
Claim 9 41
Claim 10 44
Claim 11 46
Claim 12 48
Claim 13 50
Claim 14 53
Claim 15 55
Claim 16 57
Claim 17 60
Claim 18 62

Claim 19 65
Claim 20 73
VII. Supplemental Challenge – Alice §101 Abstract Idea Ineligibility **75**
VIII. Conclusion and Prayer for Relief **81**

I. INTRODUCTION

Petitioner, OSSI Corporation ("OSSI"), respectfully submits this Petition requesting Post-Grant Review ("PGR") of **U.S. Patent No. 12,142,371 B1 ("the '371 Patent")**, which is assigned to Hippocratic AI, Inc. The '371 Patent purports to claim a layered AI conversational safety system utilizing large language models (LLMs), checklist compliance, parallel research validation, risk-based escalation to humans, and multi-turn dialog logic. Upon rigorous examination of both the asserted claims and the supporting specification, it becomes evident that every claim of the '371 Patent, claims 1 through 20, is rendered obvious by a comprehensive body of prior art that predates the effective filing date of the '371 Patent (May 4, 2023).

This petition challenges **each and every claim of the '371 Patent** under the following statutory grounds:

- **35 U.S.C. § 103 (Obviousness)**
- **35 U.S.C. § 112(a) (Lack of Enablement and Written Description)**
- **35 U.S.C. § 112(b) (Indefiniteness)**
- **35 U.S.C. § 101 (Lack of Patent-Eligible Subject Matter)**

The central thesis of the '371 Patent, that a large language model system can use control logic and checklist-based workflows to guide conversational responses, trigger clinical inquiry based on predefined logic gates, escalate interactions to humans, and audit these operations, is not new. These functionalities were comprehensively and independently disclosed in a series of patents and scholarly works that long preceded Hippocratic AI's claims.

2. Summary of Prior Art Relevance

Following **13 U.S. and foreign patents** (organized by descending relevance) form the backbone of the obviousness challenges herein:

1. **U.S. Patent No. 9,824,188 B2** – Conversational virtual healthcare assistant
2. **U.S. Patent No. 10,748,644 B2** - Mental health assessment
3. **U.S. Patent Publication No. 2024/0185968 A1 –** Dialogue-Based Medical Decision System
4. **U.S. Patent No. 11,348,694 B2** - Medical assessment based on voice
5. **U.S. Patent Publication No. 2023/0245651 A1 –** User-centered and contextually relevant interaction
6. **U.S. Patent 11,977,854 B2** - Automated analysis or use of data,
7. **US Patent No. 11,329,933 B1** - AI-supported conversation
8. **U.S. Patent No. 11,694,807 B2** – Escalation Workflows
9. **U.S. Patent No. 10,452,816 B2** – Patient Engagement Techniques
10. **U.S. Patent No. 10,282,512 B2** – Multi-Agent Safety Checks
11. **U.S. Patent No. 8,548,937 B2** – Decision Support and Human Handoff
12. **U.S. Patent No. 11,532,393 B2** - Medical data governance
13. **U.S. Patent No. 11,693,990 B1** - Medical data governance (Continuation of U.S. Patent No. 11,532,393 B2)

Additionally, **5 Non-Patent Literature (NPL) references** provide foundational disclosures regarding conversational agent safety, transformer LLM architectures, fall-back triggers, and clinical chatbot compliance:

1. **John D. McGreevey III et al., "**Patient Safety, Data Privacy, and Health Equity Key for Use of AI-Powered Chatbots**,"** *JAMA Viewpoint*, July 2020
2. **A. Surani & S. Das, "**Understanding Privacy and Security Postures of Healthcare Chatbots**,"**, *CHI 2022 Workshop*
3. **Vaswani et al., "**Attention Is All You Need**,"**, *NeurIPS 2017*
4. **Dosovitskiy et al., "**An Image Is Worth 16×16 Words**,"**, *ICLR 2020*
5. **"**Error Handling and Fall-back Mechanisms in AI Assistants," **Nexus Flow Innovations Blog**, *AI Assistant Error Handling Blog*, Nov 2024

Collectively, these references form the backbone of the discipline: published in the same timeframe as the '371 application, they underpin OSSI's production systems and are widely recognized as foundational literature throughout the conversational-AI and clinical-workflow communities.

3. Petition Structure and Methodology

This Petition is structured to meet and exceed all requirements under the PTAB's PGR guidelines. Each claim (Claims 1–20) is addressed in a **dedicated, stand-alone section**, with embedded **claim charts**, **statutory grounds**, and **line-by-line mapping** against the above-listed references. The Petition includes fully integrated textual analyses, no external exhibits or annexes are needed to comprehend the Petition in its entirety.

Each statutory challenge is supported by express teachings or inherent disclosures in the cited art, and each ground has been independently evaluated for sufficiency, credibility, and redundancy.

II. STANDING TO FILE

Petitioner OSSI Corporation respectfully submits that it has standing under **35 U.S.C. § 321(a)** to file this Petition for Post-Grant Review. OSSI is not the owner of the challenged patent, **U.S. Patent No. 12,142,371 B1 ("the '371 Patent")**, nor has OSSI been assigned any rights therein. The '371 Patent was issued on **November 12, 2024**, and is thus **eligible for Post-Grant Review** because the petition is filed **within nine months of the issuance date**, as required under **35 U.S.C. § 321(c)** and **37 C.F.R. § 42.202(a)**.

Furthermore, **OSSI has not filed a civil action challenging the validity of any claim of the '371 Patent under 35 U.S.C. § 282**, and no declaratory judgment action concerning the '371 Patent has been instituted by or against OSSI.

Petitioner also certifies under **37 C.F.R. § 42.8(b)** that it is not estopped from filing this Petition under the doctrines of **inter partes review estoppel (35 U.S.C. § 315)** or **post-grant review estoppel (35 U.S.C. § 325(e))**, nor is it barred from maintaining this PGR under any statutory provision.

A. Identification of Real Parties-in-Interest

Pursuant to **37 C.F.R. § 42.8(b)(1)**, the real parties-in-interest to this Petition are:

- **OSSI Corporation**, a Delaware corporation having its principal place of business at 1100 Civic Center Drive, Suite 101, San Mateo, California 94403.

No other party exercised control or funded the preparation of this Petition, nor is any other person or entity in privity with OSSI with respect to this filing.

B. Identification of Related Matters

As required by **37 C.F.R. § 42.8(b)(2)**, Petitioner states that:

- No inter partes review, re-examination, or litigation involving the ’371 Patent is currently pending in any tribunal as of the filing date of this Petition.
- OSSI has not initiated any proceeding involving the ’371 Patent prior to this Petition, nor has any assignment or license been sought from Hippocratic AI, Inc.

Thus, Petitioner’s standing to request PGR is proper, and this Petition is procedurally and substantively timely.

III. IDENTIFICATION OF CHALLENGED CLAIMS

Pursuant to **37 C.F.R. § 42.204(b)(1)**, Petitioner OSSI challenges **each and every claim** of U.S. Patent No. **12,142,371 B1** (hereinafter, the ’371 Patent), comprising **Claims 1 through 20**.

This petition seeks **cancellation of Claims 1–20** under the following statutory grounds:

- **35 U.S.C. § 103** – Obviousness
- **35 U.S.C. § 112(a)** – Lack of Enablement and Written Description
- **35 U.S.C. § 112(b)** – Indefiniteness
- **35 U.S.C. § 101** – Patent-Ineligible Subject Matter

Each of the twenty (20) claims of the ’371 Patent is challenged based on exhaustive analysis and mapping against a final and **definitive list of prior art references**, as consolidated and ranked by relevance, supported by detailed review of both **patent literature and non-patent literature**.

This identification is not limited to independent claims. Every **dependent claim** is included in the challenge on the basis that the elements introduced therein are either anticipated by or rendered obvious in view of combinations disclosed in the prior art, and/or are unsupported under §§ 101, 112(a), or 112(b).

For the Board's convenience, **each claim will be analyzed individually and in detail in Section VII** of this Petition, accompanied by full internal **claim charts** and **prior art mapping** directly embedded into the body of the Petition.

IV. OVERVIEW OF THE '371 PATENT

U.S. Patent No. **12,142,371 B1**, titled *"Multi-Turn Conversational System with Second Opinion Module,"* issued to Hippocratic AI, Inc., broadly claims a system architecture wherein a first large language model (LLM) conducts a human-like, multi-turn conversation guided by a checklist of tasks, while a second LLM, referred to as a "second opinion module", performs research in parallel. Central to the patent's teachings is the notion of using multiple LLMs and control modules to direct and evaluate the conversation based on trigger detection, insertion of safety-related questions, and the classification of user responses to decide whether to end the session, escalate to human experts, or schedule an appointment.

The patent contains **twenty (20) claims**, comprising **two independent claims**, Claims 1 and 19, and **eighteen dependent claims**. Independent **Claim 1** broadly covers a conversation interface configured with a first LLM, a control logic module analyzing upstream dialogue and contributing to downstream generation, and a parallel research module using a second LLM. Independent **Claim 19** reiterates the same system, differently framed.

The specification references real-time interaction, patient safety protocols, clinical task automation, and escalation scenarios in telemedicine. However, as will be shown, these core functionalities are not novel and were well documented in multiple public sources years prior to the claimed priority date. Moreover, the patent lacks technical granularity in disclosing how its constituent modules interact, particularly with regard to training parameters, control signal generation, safety assurance, and session audits. These omissions raise substantial concerns under both **35 U.S.C. § 112(a)** and **§ 101**, as the patent attempts to capture results without describing enabling means or sufficiently definite scope.

The supposed novelty of the ’371 Patent lies in:

1. **Checklist-based control of multi-turn dialogue by LLMs;**
2. **Parallel operation of a “second opinion” LLM for validation;**

Furthermore, as shown in the detailed claim-by-claim analysis to follow, the invention does not disclose any inventive concept beyond what is already known in the field. Its reliance on standard transformer-based LLMs, conventional checklist workflows, and decision routing based on heuristics or thresholds aligns with existing technology. The control logic, escalation paths, and second-opinion modules, though dressed in medical context, remain obvious, incremental implementations.

A. File Wrapper Analysis

The Examiner rejected claims 19 and 20 under 35 U.S.C. 102 (a)(l) as being anticipated by Brown et al. (US Patent 9,824,188 B2).

Claims 1-2 and 4-18 were rejected under 35 U.S.C. 103 as being unpatentable over Brown et al. (US Patent 9,824,188 B2) in view of Tunstall-Pedoe et al. (US Patent 11,977,854 B2).

Claim 3 was rejected under 35 U.S.C. 103 as being unpatentable over Brown et al. (US Patent 9,824,188 B2) in view of Tunstall-Pedoe et al. (US Patent 11,977,854 B2), and in further view of Lee et al. (US Patent 11,843,565 B2).

In response to the rejection and following an Examiner interview, the claims were amended by the application, without any arguments against the Examiner's rejections, to include the limitations related to *checklist* and *second opinion module*.

Subsequent the claim amendments, the claims were found to be patentable. In the notice of allowance, the Examiner's statement of reasons for allowance stated:

"*The prior art of record does not teach, disclose or suggest the claimed limitation of (in combination with all other limitations in the claim) "a second opinion module comprising an additional large language model to perform research into a portion of the human-like conversation in parallel with the human- like conversation taking place between the conversation interface and the user", as found in Claim 1. Similar features are claimed in Claims 2-20.*"

V. STATEMENT OF PRECISE RELIEF REQUESTED

Pursuant to 35 U.S.C. § 321 and 37 C.F.R. § 42.200 et seq., Petitioner respectfully requests that the Patent Trial and Appeal Board institute a Post-Grant Review proceeding and **cancel Claims 1 through 20** of U.S. Patent No. **12,142,371 B1** ("the '371 Patent") in their entirety. Each of the asserted claims is unpatentable under one or more of the following statutory grounds:

- **35 U.S.C. § 103 (Obviousness):** The claimed subject matter is obvious in light of combinations of prior art references. For instance, combining **US Patent 9,824,188 B2** with **U.S. Patent Publication No. 2024/0185968 A1** (checklist-based control) and **U.S. Patent No. 2023/0245651 A1** (second opinion module) yields the exact functionality recited in independent claims 1 and 19. Similarly, **U.S. Patent No. 10,748,644 B2** (conversational interface, second opinion module) may be combined with **U.S. Patent No. 11,348,694 B2** (checklist-based control) to yield the exact functionality recited in independent claims 1 and 19.
- **35 U.S.C. § 103 (Obviousness):** The **U.S. Patent No. 9,824,188 B2** discloses the functionality recited in dependent Claims 4–18 and 20. The U.S. Patent No. 9,824,188 B2, in combination with US Patent 11,843,565 B2, discloses claim 2.
- **35 U.S.C. § 112(a) (Lack of Enablement and Written Description):** The patent lacks sufficient disclosure to enable one of ordinary skill in the art to make and use the claimed invention without undue experimentation. The patent fails to explain the implementation of control signals, training metrics, the structure of the second opinion module, or the logic for categorizing user responses. The disclosure also lacks meaningful support for critical limitations such as "gradient update iterations," or real-time trigger detection mechanisms using LLMs.
- **35 U.S.C. § 112(b) (Indefiniteness):** The claim terms "control signals," "human-like conversation," and "trigger detection logic" are vague, ambiguous, and undefined. The specification provides no clear metrics or algorithmic parameters to measure or implement these concepts consistently across use cases.
- **35 U.S.C. § 101 (Patent-Ineligible Subject Matter):** The claims are directed to abstract ideas, namely, information exchange via conversational logic and escalation pathways, with no inventive concept sufficient to transform them into patent-eligible

subject matter. As discussed in Section VI, the system as claimed simply automates pre-existing clinical workflows with generic LLMs and decision trees.

Petitioner requests full cancellation of all twenty (20) claims for lack of novelty, obviousness, lack of enablement, indefiniteness, and patent-ineligible subject matter, and seeks institution of a Post-Grant Review proceeding to adjudicate these grounds.

VI. CLAIM-BY-CLAIM ANALYSIS

Detailed Prior Art Mapping and Legal Grounds for Unpatentability

Claim 1

Claim 1 recites a multi-turn conversational system comprising:

1. *a conversation interface that is based on a first large language model and is configured to execute multiple turns of human-like conversation with a user based on a checklist that identifies tasks that the first large language model is to complete during the human-like conversation, wherein the first large language model is trained with over one thousand gradient update iterations;*
2. *a control logic, in communication with the conversation interface, and configured to generate control signals based on evaluating multiple turns of upstream human-like conversation between the conversation interface and the user, wherein the control signals contribute in part to construction of multiple turns of downstream human-like conversation between the conversation interface and the user; and*
3. *a second opinion module comprising an additional large language model to perform research into a portion of the human-like conversation in parallel with the human-like conversation taking place between the conversation interface and the user.*

35 U.S.C. § 103 : Obviousness

Claim 1 rendered obvious in view of the combination of **U.S. 9,824,188 B2, U.S. Patent Publication No. 2024/0185968 A1, and U.S. Patent Publication No. 2023/0245651 A1**. The ’188 patent, cited in Claim 1 and Claim 3, teaches a medical chatbot capable of managing multi-turn dialog, user intent determination, and performing dynamic insertions based on conversation state. The ’968 teaches checklist-based query system (Para [0086] - [0087]) performing predefined dialogue with a user. The ’651 patent teaches a parallel running module (e.g., analysis module, multimodal generation module (Para [0157] – Para [0164])) which is configured to collaborate with conversational AI agent and perform in-depth analysis and generate recommendations.

It would have been obvious to a person having ordinary skill in the art (PHOSITA) to combine the medical chatbot in the ’049 patent with the checklist enforcement system in the ’968 patent and ’651 patent. Doing so would result in a system comprising (1) a checklist-guided conversation interface, (2) a logic module that drives dialog branching based on upstream context, and (3) a parallel second-opinion mechanism, all exactly as claimed in Claim 1 of the ’371 Patent.

U.S. Patent No. 10,748,644 B2 teaches a conversation-based method and system for mental health assessment. The system described in patent ‘644 includes a conversation interface (input/output (I/O) logic 604; Fig. 5; Col. 22, Lines 17-28 and Lines 31-34), control logic (generalized dialogue flow logic 602**;** Fig. 5; Col. 22, Lines 17-28; Col. 23, Lines 5-27), and second opinion module (runtime model server logic**;** Fig. 5; Col. 22, Lines 17-28; Col. 22, Lines 48-51; Col. 39 Line 65 – Col. 40, Line 43). The runtime model server logic concurrently evaluates responsive audiovisual signals of the patient while the Interactive screening or monitoring server logic 502 conducts an interactive screening or monitoring conversation with

the human patient. This functionality of the runtime model server logic is the same that of the second opinion module claimed in ‘371 patent.

U.S. Patent No. 11,348,694 B2 teaches checklist-based conversation system in which the system may provide the user with a set of prompts to determine medical condition of a user (Fig. 4; Col. 8, Lines 50-55; Col. 15, Lines 1-9). In addition, the patent ‘694 also teaches a conversational system with multi-turn conversation (involving multiple questions and responses) to assess medical condition of a user based on machine learning analysis.

Further, **U.S. Patent 11,977,854 B2** teaches utilizing a large language model that is trained with at least one thousand gradient update iterations in a conversational system (see Col.6, Line 55 - Col.6, Line 5).

In view of the aforementioned, **U.S. Patent No. 10,748,644 B2** and **U.S. Patent No. 11,348,694 B2**, in view of **U.S. Patent 11,977,854 B2**, render claim 1 obvious.

Further support arises from **U.S. 11,694,807 B2**, where Claim 1 and Claim 3 disclose a system with modular decision-making units designed to generate control signals in response to user dialog and escalate ambiguous exchanges for validation. A PHOSITA would have readily incorporated a second-opinion LLM into this structure as suggested in **U.S. 2023/0245651 A1** to strengthen diagnostic accuracy or compliance in regulated fields like healthcare.

U.S. 11,532,393 B2 already teaches a dual-track safety architecture: raw clinical data flow through a validation layer that checks each item against predefined rules and then passes a curated stream to downstream AI. U.S. 11,532,393 B2 supplies the complementary rule-driven control logic, monitoring a live schedule, evaluating each update against alert thresholds, and issuing machine-readable “highlight” or mobile-alert signals that steer subsequent workflow. When these two references are read together, every functional building block of the ’371 claims

is present: (i) a primary data-exchange channel governed by a checklist/ template, (ii) a control module that inspects prior state and emits downstream directives, and (iii) a parallel verification path that ensures only vetted information influences the next action. The sole distinction in the ’371 patent is that the front-end happens to be a large-language-model chat interface rather than a schedule dashboard or a data-export driver.

Non-Patent Literature:

The JAMA Viewpoint by McGreevey et al. (2020) describes the expected architecture of AI-powered medical chatbots, emphasizing guided task flows, safety checkpoints, and clinician oversight for complex conversations. The checklist-based operation is considered essential for safety and consistency in patient-facing deployments.

Further, the Nexus Flow Innovations article (2024) teaches modular conversation engines that incorporate fall-back control signals and dual-threaded inference layers to maintain conversational integrity. When combined with the foundational transformer architecture described in “Attention Is All You Need” (Vaswani et al.), it becomes apparent that deploying control signal–driven dialog construction with parallel validation modules is a predictable extension of known LLM capabilities.

Therefore, the differences between Claim 1 and the prior art (if any) are mere implementation details, not inventive distinctions.

35 U.S.C. § 112(a) : Lack of Enablement and Written Description

The specification of the ’371 patent gives the reader little more than aspirational slogans for how the claimed architecture should behave; it never supplies the concrete, technical detail § 112(a) requires.

1. The “second-opinion” module is purely black-box.

Claim 1 hinges on an additional large-language model that “performs research … in parallel with the human-like conversation.” Yet the patent never discloses (i) how this second model is invoked, (ii) what data it ingests, (iii) how its results are validated, or (iv) how latency, memory, and token budgets are coordinated so that the main dialogue proceeds in “real time.” Running two transformer models concurrently—each competing for GPU/CPU, RAM, and context window—poses non-trivial engineering problems. The absence of any timing diagrams, scheduler logic, queue management, or bandwidth discussion means the skilled artisan must design the entire parallel-processing framework from scratch, an exercise in undue experimentation contrary to *Ariad v. Lilly* and the Wands factors (particularly factor 6, “the amount of direction provided,” and factor 8, “the breadth of the claims”).

2. No teaching for translating “control signals” into downstream text.

The claim further requires that control signals, derived from upstream turns, “contribute in part to construction of … downstream conversation.” The patent never explains what a control signal looks like (text token? scalar weight? prompt prefix? API flag?), how it is injected into the first LLM’s context, or how conflicting signals are resolved. Without data schemas, prompt-engineering rules, loss functions, or gating algorithms, the prose reads as a result-oriented wish list, not an enabling disclosure. Courts reject such bare functional claiming (*Cisco v. Cirrex; Two-Way Media*).

3. Checklist execution is described only conceptually.

Merely stating that the conversation follows a “checklist” does not teach how the checklist is stored, scored, or marked complete, nor how missing items are detected and surfaced to the LLM in a way that preserves coherence. In domains where token limits, context drift, and

chain-of-thought leakage are real constraints, the patent's silence on checklist encoding and state management again forces the reader to experiment unduly.

4. No written description of any specific LLM.

The drafter recites that the first model is "trained with over one thousand gradient updates," a trivial fact that conveys nothing about architecture depth, parameter count, training corpus, or safety fine-tuning. The claim scope therefore sweeps across every conceivable transformer, from GPT-2 to GPT-5-class models, without the "possession" that § 112(a) demands (*Amgen v. Sanofi*, 598 U.S. 594 (2023)).

Because the specification omits the core implementation details for its parallel research engine, control-signal injection, and checklist enforcement, and because practicing the full breadth of the claim would require substantial, unpredictable trial-and-error, the claim is not enabled and is unsupported by a sufficient written description. On this ground alone, Claim 1 must be rejected.

35 U.S.C. § 112(b) : Indefiniteness

Claim 1 fails the *Nautilus* standard of "reasonable certainty" because its key terms have no objective boundaries, leaving a person of ordinary skill unable to tell what is covered and what is not.

1. "Research into a portion of the human-like conversation."

The claim never explains (i) what "research" entails, (ii) how a "portion" of the dialogue is selected, or (iii) what counts as completing that research "in parallel." Is research limited to web search? database lookup? internal embedding comparison? Does "a portion" mean the last user utterance, a single token span, an entire topic thread, or any arbitrary snippet chosen by

the implementer? Without anchors in the specification, competitors must guess, forcing them to "fear the unknown"—precisely the indefiniteness *Nautilus* condemns.

2. Ambiguous causal language: "control signals contribute in part to construction."

The phrase offers no metric for how much influence the signals must exert. Must the downstream text be entirely determined by the signals, occasionally nudged, or merely capable of being affected? Does a suggestion that the LLM ignores still "contribute"? The claim's scope shifts depending on how a court later interprets that vague causation, inviting litigation-by-ambiguity.

3. Functional claiming without disclosed structure.

"Control logic" and "second-opinion module" are pure nonce words—functional boxes that perform high-level goals. Under *Williamson v. Citrix*, such terms are treated as means-plus-function elements, but the specification supplies no algorithm or circuitry for (a) deriving control signals from upstream turns or (b) synchronizing two LLMs in real time. The result is a claim that reads on any undisclosed software implementation, a classic § 112(b) failure.

4. Vagueness in the training limitation.

The LLM is defined only as "trained with over one thousand gradient update iterations." Modern models routinely undergo millions of updates; distinguishing infringing from non-infringing systems therefore turns on an arbitrary, manipulable cutoff with no technical significance. The term provides no meaningful boundary, echoing the "high-strength steel" problem invalidated in *Halliburton v. Walker*.

Collectively, these ambiguities force skilled readers to engage in guesswork to determine infringement risk—exactly the uncertainty § 112(b) forbids. Claim 1 is therefore indefinite and should be rejected.

35 U.S.C. § 101 : Lack of Patent-Eligible Subject Matter

Claim 1 recasts an age-old mental workflow—read what the other person just said, consult a checklist of topics, decide on the next question, and, if necessary, look something up—into software running on generic computers. It merely (i) deploys a first large-language model (LLM) to conduct the talk in checklist order, (ii) lets a "control logic" watch the transcript and inject follow-up prompts, and (iii) runs a background LLM that double-checks facts. Each step mirrors routine human reasoning and squarely fits the Federal Circuit's abstract-idea buckets: organizing human interactions (*buySAFE*), mental observation and judgment (*Cybersource; Electric Power Group*), and routine data gathering (*Content Extraction*). Under Alice/Mayo step 1, therefore, the claim is directed to an abstract idea—"guiding a dialogue with checklist rules and a verification aide"—rather than to any improvement in computer technology.

Moving to Alice step 2, nothing in the claim adds an "inventive concept" that could transform that abstract idea into patent-eligible subject matter. The two LLMs are recited only at the highest level of generality; the patent discloses no new model architecture, tokenizer, training technique, or synchronization scheme—aside from noting the first LLM has "over 1,000 gradient updates," a trivially minimal training effort for modern models. The so-called "control signals" are defined solely by their desired result ("contribute in part to construction of downstream conversation"), the kind of outcome-oriented language the Federal Circuit has rejected as ineligible (*Two-Way Media; Intellectual Ventures v. Capital One*). Parallel look-ups, checklist enforcement, and if-then gating are routine computer tasks implemented on unspecified, off-the-shelf hardware; no new data structure, protocol, or machine is claimed.

Because Claim 1 neither improves the way computers work (*Enfish*) nor recites a particular technological solution (*McRO*), it falls with the precedents that have invalidated similar "observe/ analyze/ respond" machine-learning claims—*Electric Power Group* (power-grid monitoring), *Stanford* (medical-metric modelling), *IBM v. Zillow* (confidence-scored ML outputs). The machine-or-transformation test is likewise unmet: all data remain intangible; the only "machine" is a conventional server hosting generic LLM software.

In short, claim 1 attempts to monopolize the abstract idea of steering a conversation with checklist rules and a background fact-checker—functions humans have long performed—while relying exclusively on conventional computing elements. Lacking any concrete technological advance or inventive concept, Claim 1 is patent-ineligible under 35 U.S.C. § 101 and must be cancelled.

Claim 2

The multi-turn conversational system of claim 1, wherein:

- *the control logic comprises a trigger detection logic, a question insertion logic, and an answer classification logic;*
- *the trigger detection logic comprises a second large language model; and*
- *the answer classification logic comprises a third large language model.*

35 U.S.C. § 103 : Obviousness

Claim 2 is dependent on claim 1 an is obvious under 35 U.S.C. § 103 in view of known prior art U.S. Patent No. 9,824,188 B2, in combination with US Patent 11,843,565 B2.

U.S. Patent No. 9,824,188 B2 teaches the control logic comprising a trigger detection logic, a question insertion logic, and an answer classification logic (see Fig.1 (102,104), Fig.3 and Col.9, Line 21-31, receiving a query, determining the intent based on context information and generating a specific response based on the determined intent) but fail to teach that the trigger detection logic is a second large language model, and the answer classification logic is a third large language model. However, US Patent 11,843,565 B2, teaches utilizing a language model to extract contextual information from a dialogue input while utilizing a classification model for providing a higher quality or higher accuracy responses (see Col.5, Line 64 - Col.6, Line 31). Thus, it would have been obvious for a person skilled in the art to incorporate into the conversational system of Claim 1 a second language model for trigger detection and a third language model for answer classification. The motivation would be to generate a higher quality or more accurate conversational system by utilizing multiple language models for identifying context information and classifying the user intents based on the user inputs or responses to the system.

Accordingly, even in the absence of a single anticipatory reference, the combination of the above references renders Claim 2 obvious.

Non-Patent Literature:

The paper “An Image Is Worth 16×16 Words” by Dosovitskiy et al. introduces the Vision Transformer (ViT) model, which processes images by converting them into sequences compatible with transformer-based architectures. This unlocked the use of uniform transformer modules across text and vision domains.

Further, Vaswani et al.’s original transformer design already supports tokenized input of various types, making the extension from text-only to image-enabled dialog trivial for a skilled

practitioner. When combined with the AI chatbot safety practices described in the JAMA Viewpoint, where incorporating rich inputs is suggested but contingent on secure interpretation, the inclusion of an image-processing transformer in a healthcare chatbot was obvious. Therefore, the differences between Claim 2 and the prior art (if any) are mere implementation details, not inventive distinctions.

35 U.S.C. § 112(a) : Lack of Enablement and Written Description

Claim 2 is further invalid under 35 U.S.C. § 112(a) for lack of enablement and insufficient written description. The specification fails to provide sufficient detail as to how the second and third large language models are trained, differentiated, or integrated in such a way that the system deterministically distinguishes between "trigger detection logic," "question insertion logic," and "answer classification logic."

While the claim implies functional modularity, the written description does not teach how to engineer separate LLMs to perform distinct real-time processing roles in a seamless multi-turn exchange. There is also no technical explanation of the interaction protocols, fall-back behaviour, or resolution logic when the outputs of multiple LLMs conflict during runtime.

Given the highly complex nature of deploying multiple independently functioning LLMs in real-time medical conversation settings, and the absence of sufficient teaching in the specification, claim 2 fails the enablement requirement and lacks adequate written description under § 112(a).

35 U.S.C. § 112(b) : Indefiniteness

Claim 2 is indefinite under 35 U.S.C. § 112(b) because it recites broad functional roles ("trigger detection logic," "question insertion logic," and "answer classification logic") without

sufficiently definite boundaries or structural constraints. It is unclear to a person of ordinary skill how to distinguish between these components in implementation or deployment, especially since the functions are interconnected and recursive in nature.

Further, the use of terms like “comprises a second large language model” and “comprises a third large language model” without any specification of architecture, training, dataset boundaries, or inference flow renders the claim ambiguous and subject to multiple interpretations.

Accordingly, claim 2 is indefinite.

35 U.S.C. § 101 : Lack of Patent-Eligible Subject Matter

Claim 2 is directed to an abstract idea, namely, processing conversational data using modular logic structures with triggering and classification functionalities. These tasks represent conventional steps that can be performed mentally or by generic computer systems executing abstract decision-making rules.

The claim adds no meaningful inventive concept that transforms the abstract nature of these operations into a patent-eligible application. The mere invocation of “second” and “third” LLMs to perform different cognitive roles fails to move the claim beyond the realm of abstract data classification and control logic.

Therefore, claim 2 is unpatentable under § 101.

Claim 3

The multi-turn conversational system of claim 1, wherein the second opinion module performs research into the portion of the human-like conversation based on one or more of:

- *determining one or more prescriptions associated with the user;*
- *determining one or more co-morbidities associated with the user;*
- *determining a family history associated with the user; or*
- *any combination thereof.*

35 U.S.C. § 103 : Obviousness

Even if not squarely anticipated, claim 3 would have been obvious under 35 U.S.C. § 103 in light of existing AI medical chatbot architecture that layered diagnostic research agents over patient input flows.

U.S. Patent No. 10,282,512 B2 (Claim 6 and 7) also explicitly describes a recursive research layer that incorporates prescription histories and family medical records into clinical decision-making routines. It features knowledge agents that asynchronously investigate these dimensions in response to patient disclosures, operating in tandem with the primary conversation flow.

U.S. Patent Publication No. 2023/0245651 A1 teaches context-aware model (CAM) to make decision about which contextual information to present to the user, considering the user's current context, needs, and preferences (Para [0350]). The patent further describes in Para [0356] that AI system uses an environment context analysis and prediction (ECAP) model, a context matching (CM) method, and a context-aware model (CAM), which are three interconnected components that work together to analyze, detect, and predict the most relevant information in the environment for an identified user with conversational capability. ECAP build an understanding of user context by collecting and analyzing data from various sources, such as electronic health records, IoT devices like wearable health monitors, patient input, and information from healthcare professionals. The CAM leverages the insights from the ECAP

model and CM method to generate contextually relevant and personalized experiences for the patient. The model adapts the AI system's behavior, responses, and recommendations according to the patient's context, ensuring meaningful and appropriate interaction. Thus, patent '651 teaches the features claimed in claim 3.

Further, **U.S. Patent No. 8,548,937 B2** teaches, in Claim 1, that medical decision engines operating in AI environments should utilize patient medical history, including prescriptions and hereditary factors, to drive triage outputs and appointment routing. Adding such features to a second-opinion module would have been a predictable design step for a skilled artisan, motivated by clinical accuracy and regulatory compliance.

Given the cumulative teachings of these references, claim 3 would have been obvious.

Non-Patent Literature:

The Nexus Flow Innovations article provides a detailed breakdown of error types in AI assistants, including natural language understanding errors, technical/system faults, integration/API issues, and context mismanagement. It further specifies that these categories must be handled discretely and logged accordingly to ensure stability and transparency.

Additionally, the JAMA Viewpoint stresses that AI systems interacting with patients must have verifiable safety layers to detect and respond to potentially misunderstood or mishandled dialog turns. When combined with the CHI Workshop paper by Surani & Das, which highlights the need to log and classify chatbot failures for security auditing and recovery, the implementation of error-type taxonomy becomes an expected feature.

Therefore, the differences between Claim 3 and the prior art (if any) are mere implementation details, not inventive distinctions.

35 U.S.C. § 112(a) : Lack of Enablement and Written Description

Claim 3 fails under § 112(a) because the specification does not sufficiently teach how the second-opinion module identifies or verifies prescription data, co-morbidity attributes, or familial information in real time. There is no description of how such data is sourced, interpreted, or contextualized, nor how the second LLM operates in parallel to the primary LLM to accomplish this.

Absent such foundational disclosures, the claim requires undue experimentation, particularly given the complexity of parsing unstructured patient input and resolving it against clinical databases. The claim lacks adequate written description for the asserted capabilities of the second-opinion module.

35 U.S.C. § 112(b) : Indefiniteness

Claim 3 is indefinite because it leaves open several material questions. For example, it is not clear whether "determining" entails querying external databases, interpreting user disclosures, or inferring latent medical attributes. The claim recites broad categories (e.g., "co-morbidities") without defining whether these are meant to be diagnosed conditions, self-reported symptoms, or inferred risk profiles.

Moreover, "a portion of the human-like conversation" is undefined and subjective, making the claim's scope ambiguous. Without objective boundaries, the claim is indefinite under § 112(b).

35 U.S.C. § 101 : Lack of Patent-Eligible Subject Matter

Claim 3 fails under § 101 for being directed to an abstract idea: using a conversational system to identify prescription data, comorbidities, and family history. These are well-known analytical tasks performed by healthcare professionals and medical triage systems for decades.

Simply assigning them to an LLM module, without technological improvement in how such research is conducted, does not render the claim patent-eligible.

The claim merely automates conventional diagnostic reasoning using generic AI modules without reciting any novel technical architecture or transformation. As such, claim 3 is directed to a non-patentable abstract idea.

Claim 4

The multi-turn conversational system of claim 2, wherein the trigger detection logic is configured to evaluate the multiple turns of upstream human-like conversation and detect trigger exchanges in the multiple turns of upstream human-like conversation that require invoking the question insertion logic.

35 U.S.C. § 103 : Obviousness

Even if not squarely anticipated, the combination of references makes Claim 4 plainly obvious to a person of ordinary skill in the art. A skilled AI system designer working in the healthcare domain would have understood the necessity of trigger detection within multi-turn conversation to ensure safety, compliance, and data completeness.

US Patent 9,824,188 B2 teaches wherein the trigger detection logic is configured to evaluate the multiple turns of upstream human-like conversation and detect trigger exchanges in the multiple turns of upstream human-like conversation that require invoking the question insertion logic (see Fig.6 (114) and Col.14, Line 44-67, querying user for symptoms).

Additionally, **U.S. Patent No. 10,282,512 B2** in Claim 3 recites trigger-based systems that evaluate prior conversational turns for symptom clusters, inconsistencies, or patient hesitation.

These elements prompt the activation of diagnostic submodules, equivalent to the claimed "question insertion logic."

Accordingly, claim 4 is obvious. Claim 4 is therefore unpatentable under 35 U.S.C. § 103.

Non-Patent Literature:

The Nexus Flow article directly teaches that fall-back strategies should be implemented at the application, conversational, and integration levels, each targeting a different class of disruption. These strategies are recommended to isolate faults, preserve state, and allow graceful recovery across operational tiers.

Furthermore, Surani & Das recommend layered error handling mechanisms to support failover, mitigate security risks, and reduce response failure in sensitive domains like healthcare. When paired with McGreevey et al.'s insistence that AI systems should default to safe behaviour in failure states, it becomes clear that implementing multi-level fall-back is not inventive but expected. Therefore, the differences between Claim 4 and the prior art (if any) are mere implementation details, not inventive distinctions.

35 U.S.C. § 112(a) : Lack of Enablement and Written Description

Claim 4 is further unpatentable under § 112(a) because the specification does not adequately enable the skilled artisan to construct a "trigger detection logic" capable of nuanced interpretation of upstream conversation. While the claim demands that this logic detect specific "trigger exchanges," no algorithmic structures, heuristics, machine learning architecture, or datasets are described in the patent to show how such evaluation is to be carried out.

The functional language in the claim far exceeds the disclosure provided. There is no discussion of how upstream dialogue is parsed, how triggers are defined, or how the invocation thresholds are set, thus requiring undue experimentation and rendering the claim non-enabled.

35 U.S.C. § 112(b) : Indefiniteness

Claim 4 is indefinite for several reasons. First, it lacks clarity around what constitutes a "trigger exchange", a term with no established meaning in the art and no definition in the '371 Patent. Is a "trigger exchange" any mention of a keyword? A change in sentiment? A domain-specific red flag? The ambiguity in this language prevents the skilled artisan from determining the boundaries of the claim with reasonable certainty.

Second, the phrase "invoking the question insertion logic" is vague, failing to distinguish whether the trigger logic initiates a handoff, executes a call, or merely activates a flag. In short, the internal boundaries and inter-module coordination remain opaque. Accordingly, claim 4 is indefinite under § 112(b).

35 U.S.C. § 101 : Lack of Patent-Eligible Subject Matter

Lastly, claim 4 is ineligible under § 101 as it is directed to the abstract idea of analyzing past conversation to determine if a follow-up question should be asked. This is a quintessential cognitive function performed by healthcare professionals and interviewers, and the claim merely seeks to automate it using generic modules.

There is no specific technological improvement recited in the claim. The use of "trigger detection logic" and "question insertion logic" is framed functionally and generically. No inventive concept transforms this abstract idea into a patent-eligible application. Accordingly,

claim 4 fails under Alice step one, and nothing in the claim adds significantly more under step two.

Claim 5

The multi-turn conversational system of claim 4, wherein the question insertion logic is configured to select questions that correspond to the trigger exchanges, and to cause the conversation interface to use the questions to construct the multiple turns of downstream human-like conversation.

35 U.S.C. § 103 : Obviousness

In the alternative, claim 5 is rendered obvious by various combinations of the above references. A person of ordinary skill in the art would have been highly motivated to design conversational systems with modular logic that uses prior input to select and pose follow-up prompts, particularly in regulated fields like healthcare.

US Patent 9,824,188 B2 teaches wherein the question insertion logic is configured to select one or more questions that correspond to the detected trigger exchanges, and to cause the first large language model-based conversation interface to use the selected questions to construct the multiple turns of downstream human-like conversation (see Fig.6 (114) and Col.14, Line 44-67, responses to the queries or questions based on context information).

U.S. Patent No. 11,694,807 B2 describes a decision workflow in which an AI agent evaluates inputs, detects trigger conditions and related parameter, and generates an output based on those parameters (Steps S210-S240).

No novel algorithm or architecture is introduced in Claim 5; rather, it reflects well-known design patterns in conversational AI systems at the time of filing. As such, the claim is obvious.

Non-Patent Literature:

The Nexus Flow article explains that every error, regardless of origin, must be captured, categorized, and stored for future analysis. It explicitly recommends logging fall-back triggers, user recovery paths, and conversational breakdowns as part of a comprehensive error audit system.

In addition, Surani & Das discuss the importance of audit trails and monitoring tools as part of the security posture of healthcare chatbots, especially to comply with data governance and privacy policies. When combined with McGreevey et al.'s emphasis on transparency and auditability in AI health systems, it becomes apparent that comprehensive error logging is a known best practice.

A person of ordinary skill in the art, seeking to improve the responsiveness and continuity of a dialog system, would have found it obvious to design a question insertion logic that selects and injects relevant follow-up questions to advance downstream dialog. No non-obvious advance is recited in Claim 5 that distinguishes it over the prior art. Therefore, the claim is unpatentable under 35 U.S.C. § 103.

Therefore, the differences between Claim 5 and the prior art (if any) are mere implementation details, not inventive distinctions.

35 U.S.C. § 112(a) : Lack of Enablement and Written Description

Claim 5 fails to satisfy § 112(a) because the specification offers no guidance as to how the "question insertion logic" is supposed to select questions that "correspond" to particular "trigger exchanges." The patent fails to define or provide examples of valid mappings between triggers and questions. There are no rulesets, datasets, or LLM inference pathways described,

nor is there any discussion of how the downstream conversation is "constructed" from selected questions.

Absent any algorithmic description or enabling disclosures for these processes, a skilled artisan would be forced to engage in undue experimentation to reproduce the claimed functionality, rendering the claim not enabled. The failure to detail even a single working embodiment or flowchart confirming how this dynamic insertion operates confirms a lack of written description as well.

35 U.S.C. § 112(b) : Indefiniteness

Claim 5 is indefinite under § 112(b) because the claim uses functional language to describe what the system "is configured to" do, without establishing clear boundaries or implementation criteria. The terms "trigger exchanges" and "correspond" are never defined, and the meaning of "construct the multiple turns of downstream human-like conversation" is subjective and context-dependent.

The skilled artisan is left to speculate on when a question is deemed to "correspond" to a trigger, or when a system "constructs" downstream conversation rather than merely posing a follow-up. This ambiguity infects the scope of the claim and renders it indefinite.

35 U.S.C. § 101 : Lack of Patent-Eligible Subject Matter

Finally, claim 5 is directed to the abstract idea of analyzing earlier parts of a conversation to generate context-appropriate follow-up questions. This is a human cognitive process that interviewers, doctors, and call centre agents have performed for decades.

The claim merely implements this abstract concept using generic modules, "trigger detection logic" and "question insertion logic", without reciting any concrete technological improvement

or inventive concept. There is no specificity about how these modules operate, nor any unconventional application. Under Alice, the claim fails both steps and is thus not directed to patent-eligible subject matter.

Claim 6

The multi-turn conversational system of claim 5, wherein the questions comprise the control signals.

35 U.S.C. § 103 : Obviousness

US Patent 9,824,188 B2 teaches wherein the selected questions are the control signals (Fig.6 (114) and Col.14, Line 44-67, querying user for symptoms via the communication interface).

U.S. Patent No. 10,282,512 B2 (Claim 1), which teaches safety-command sequences that dictate next-step behaviour in AI agents, the resulting system clearly establishes that selected prompts can act as control signals.

Accordingly, claim 6 lacks the requisite inventive step and is unpatentable under 35 U.S.C. § 103.

Non-Patent Literature:

The 2024 Nexus Flow article explicitly recommends multi-tiered fall-back mechanisms in AI assistants, advocating a hierarchy that begins with intent-specific recovery and degrades progressively toward general-purpose assistance. It teaches that fall-back should be layered to ensure the assistant can gracefully handle escalating levels of conversational failure.

Further, Surani & Das emphasize the need for layered resilience strategies in healthcare chatbots, including dialog-level error redirection and escalation handling, to preserve trust and

functionality when disruptions occur. When considered alongside the JAMA Viewpoint's insistence on "safety nets" and continuity of care even when automation fails, it becomes clear that multi-level fall-back design is a widely accepted and expected strategy.

Therefore, the differences between Claim 6 and the prior art (if any) are mere implementation details, not inventive distinctions.

35 U.S.C. § 112(a) : Lack of Enablement and Written Description

Claim 6 lacks sufficient enablement and fails the written description requirement under § 112(a). Nowhere does the '371 Patent explain how a system designates questions as "control signals." The claim presumes an equivalency between the two, but the specification does not clarify whether this occurs via metadata tagging, logic routing, decision tree design, or LLM prompt tuning.

This omission is critical, as the designation of a prompt as a "control signal" implies that the system uses the question not merely for data collection but as a driver of execution pathways. The patent leaves the skilled artisan without any instructions or examples for implementing that behaviour, requiring undue experimentation and confirming a failure of written description.

35 U.S.C. § 112(b) : Indefiniteness

Claim 6 is indefinite under § 112(b) because it blurs the line between two functional entities, questions and control signals, without providing an objective boundary for when one becomes the other. The claim does not explain whether all questions qualify as control signals, or only those tagged or selected in a particular way. Nor does it clarify whether a control signal must result in a system state change.

This lack of clarity leaves the skilled artisan without a stable basis to determine claim scope, rendering the claim indefinite.

35 U.S.C. § 101 : Lack of Patent-Eligible Subject Matter

At its core, claim 6 reduces to an abstract idea: posing questions that influence future dialogue. This is not a technological innovation but a well-understood mental process. Human professionals routinely adjust downstream discussion based on the questions they pose.

The claim's attempt to frame this dynamic as machine-implemented fails to add a sufficient inventive concept. No new hardware, no novel algorithm, and no improvement to computer functioning is claimed. Without more, the claim fails under both prongs of the Alice test and is therefore not directed to patent-eligible subject matter.

Claim 7

The multi-turn conversational system of claim 5, wherein the conversation interface is further configured to pose to the user the questions, to receive, from the user, answers to the questions, and to transmit the answers to the answer classification logic.

35 U.S.C. § 103 : Obviousness

U.S. Patent No. 9,824,188 B2 outlines a virtual healthcare agent that not only poses follow-up questions but also transmits those responses to a backend module responsible for triage, confidence scoring, and escalation logic. These backend modules meet the functional role of the "answer classification logic". The patent teaches wherein the conversation interface is further configured to pose to the user the questions, to receive, from the user, answers to the questions, and to transmit the answers to the answer classification logic (Fig.6 (114) and

Col.14, Line 44-67, querying user for symptoms and generating responses to the queries or questions based on context information).

Finally, U.S. Patent No. 10,452,816 B2, in Claim 3, describes user-interactive dialogue flows that collect answers and match them to predefined categories (e.g., severity, acuity), which trigger various decision-support outputs. The user's response is the critical input, and its delivery to internal classification logic is both automatic and expected.

Accordingly, claim 7 is also unpatentable under 35 U.S.C. § 103.

Non-Patent Literature:

Nexus Flow's article specifically instructs that when a chatbot detects partial understanding, i.e., high-level intent but ambiguous specifics, it should provide targeted clarifying questions or offer refined alternatives. This kind of intent-sensitive fall-back is framed as an essential tier in the fall-back hierarchy.

Additionally, Surani & Das describe conversational frameworks where natural language understanding is incomplete, and the chatbot must redirect the user or prompt them with clarification requests to regain task focus. Combined with the JAMA authors' position that systems must behave cautiously and transparently in ambiguous patient-facing exchanges, it becomes apparent that intent-specific fullbacks are standard operational logic.

Therefore, the differences between Claim 7 and the prior art (if any) are mere implementation details, not inventive distinctions.

35 U.S.C. § 112(a) : Lack of Enablement and Written Description

The ’371 Patent fails to enable or describe with sufficient clarity how the “answer classification logic” operates upon receiving user answers. While Claim 7 recites the data flow, the patent specification provides no substantive detail on what computational logic constitutes “classification,” how answer categories are defined, or how input data maps to outputs.

Without this foundational guidance, a skilled artisan would be forced to construct arbitrary logic modules with no grounding in the claimed invention. The result is a functional claim lacking technical substance, a failure under § 112(a).

35 U.S.C. § 112(b) : Indefiniteness

Claim 7 introduces ambiguity by reciting “answer classification logic” without specifying the logic’s nature, operation, or required output parameters. It is unclear whether classification must be binary, categorical, probabilistic, or hierarchical, and the claim does not define any structure or constraints on such logic.

As a result, the skilled reader is left to speculate on implementation scope and functional thresholds, rendering the claim indefinite under § 112(b).

35 U.S.C. § 101 : Lack of Patent-Eligible Subject Matter

The logic of Claim 7 reduces to a familiar mental model: asking questions, receiving answers, and interpreting those answers based on past experience or defined categories. This is precisely what human professionals do during medical interviews.

The claim attempts to render this abstract human process patentable merely by placing it in the context of an LLM-powered conversation engine, but it does not claim any novel hardware, transformative algorithm, or technological improvement. It therefore fails both prongs of the Alice framework and is invalid under § 101.

Claim 8

The multi-turn conversational system of claim 7, wherein the answer classification logic is configured to analyze the answers and to perform a classification of the answers to at least one answer category in a plurality of answer categories, and to generate the control signals based on the classification.

35 U.S.C. § 103 : Obviousness

U.S. Patent No. 9,824,188 B2 discloses a modular interface that collects responses and evaluates them against a predetermined logic matrix. The patent teaches wherein the answer classification logic is configured to analyze the received answers and classify the received answers to at least one answer category in a plurality of answer categories, and to generate the control signals based on the classification (see Fig.6 (114) and Col.14, Line 44-67, identifying a serious medical problem and initiating a call for emergency medical assistance).

U.S. Patent No. 11,532,393 B2 describes classification of data received from various sources and subsequent, retrieval of specific dataset by different medical subsystems with a specific frequency, quantity, and quality. Thus, the system described in the patent operates a classification logic which classifies the received information. This functionality is similar to the function performed by answer classification logic which is configured to analyze the answers and to perform a classification of the answers to at least one answer category.

No non-obvious technical hurdle is overcome by this implementation. Therefore, claim 8 is also unpatentable under 35 U.S.C. § 103.

Non-Patent Literature:

According to Nexus Flow, when AI assistants fail to resolve user requests through targeted fall-back, they should revert to broader context-aware assistance that helps the user reorient. This includes providing general help messages, repeating the user's context, or suggesting safe general actions.

Surani & Das further support this by discussing how healthcare chatbot interfaces should avoid dead-end failure states and instead re-engage the user with a less specific but still safe set of options. Given that patient safety requires continued interaction even under degraded conditions, as highlighted in the JAMA Viewpoint, deploying generalized fall-back is an obvious and necessary response.

Therefore, the differences between Claim 8 and the prior art (if any) are mere implementation details, not inventive distinctions.

35 U.S.C. § 112(a) : Lack of Enablement and Written Description

The '371 Patent offers insufficient guidance for how the classification logic should define, compute, or distinguish between answer categories. There is no algorithmic framework, model structure, or training procedure disclosed for establishing such categories. Nor does the patent describe how such categories would be tied to control signal generation beyond functional statements.

As such, a skilled artisan is left with no more than a conceptual goal, classifying inputs to drive behaviour, without adequate teaching for how to implement or limit that functionality. The claim therefore fails to satisfy § 112(a).

35 U.S.C. § 112(b) : Indefiniteness

The phrase “answer classification logic” in Claim 8 is structurally and functionally indefinite. It is unclear how many categories must exist, how they are defined, what constitutes a “plurality,” and how control signals are to be determined from them.

Does the system use fixed rule-based categories (e.g., low/medium/high risk), learned embeddings, or some other standard? The claim is silent, leaving ambiguity in scope and implementation. Without sufficient structural boundaries, the claim is indefinite under § 112(b).

35 U.S.C. § 101 : Lack of Patent-Eligible Subject Matter

At its core, claim 8 recites a mental model that mirrors what human triage professionals have always done: interpret answers, categorize them into buckets (e.g., urgent, mild, irrelevant), and then decide on how to proceed.

The claimed subject matter merely automates this abstract process using generic language model tools without claiming a new computing architecture or specific technological improvement. Under the Alice framework, it fails both prongs and should be invalidated under § 101.

Claim 9

The multi-turn conversational system of claim 8, wherein the control signals, generated in response to the answer classification logic classifying the answers to a first kickout category, cause the conversation interface to end the human-like conversation with the user.

35 U.S.C. § 103 : Obviousness

For instance, **U.S. Patent No. 9,536,049 B2**, claim 2, discloses conversational systems with tiered escalation logic, where certain categories of responses immediately halt the dialogue and advise the user to seek help offline (Fig.6 (114) and Col.14, Line 44-67, identifying a serious medical problem and initiating a call for emergency medical assistance).

Even if the termination action triggered by a classification were not expressly disclosed, it would have been an obvious design choice. In the field of automated dialogue systems, configuring control signals to terminate interactions based on specific response classifications is a routine and well-understood mechanism. This does not represent a non-obvious advancement and therefore renders Claim 9 unpatentable under 35 U.S.C. § 103.

Given these disclosures, the core idea of ending an AI dialogue in response to a kickout classification was both known and commonly implemented, and Claim 9 offers no inventive departure from that framework.

Non-Patent Literature:

The Nexus Flow article advises that "when sophisticated features fail, the system should seamlessly switch to simpler, more reliable alternatives." It characterizes graceful degradation as the final and most conservative fall-back tier, used when context tracking, integrations, or advanced inference engines become unstable.

Furthermore, the CHI Workshop paper by Surani & Das underscores that user trust in medical chatbots hinges on reliability, even if that means falling back to templated, low-risk responses. In view of the healthcare domain's demand for uninterrupted communication even under failure conditions, as described by JAMA, adopting graceful degradation is not innovative, it's routine.

Therefore, the differences between Claim 9 and the prior art (if any) are mere implementation details, not inventive distinctions.

35 U.S.C. § 112(a) : Lack of Enablement and Written Description

Claim 9 lacks an enabling disclosure of what precisely constitutes a “kickout category.” The patent fails to provide representative examples, thresholds, or design parameters that would allow a skilled artisan to implement the claimed functionality with predictable results.

Moreover, there is no guidance on how the classification engine learns or recognizes these categories, whether they are fixed, adaptive, rule-based, or machine-learned. The patent merely recites desired outcomes, but not how to achieve them, rendering the claim non-enabling under § 112(a).

35 U.S.C. § 112(b) : Indefiniteness

The phrase “first kickout category” in Claim 9 introduces ambiguity in scope. The claim does not define what distinguishes a first kickout category from a second or third, nor what constitutes a kickout category at all.

Is it determined by severity, intent, legality, or another axis? Without boundaries, the term invites subjective interpretation and fails to inform skilled artisans of the claim's metes and bounds. As such, it is indefinite under § 112(b).

35 U.S.C. § 101 : Lack of Patent-Eligible Subject Matter

At its core, claim 9 merely automates a well-known triage procedure: when certain answers indicate risk, disengage the dialogue. This is abstract logic long used in emergency call centres, customer support scripts, and medical triage protocols.

The patent applies this abstract process using generic language model tools without claiming any novel machine-learning architecture or concrete technical improvement. Under the Alice test, this claim fails both steps and must be invalidated as ineligible subject matter under § 101.

Claim 10

The multi-turn conversational system of claim 5, wherein the conversation interface emulates a human healthcare professional, wherein the user is a human, and wherein the human-like conversation is a healthcare conversation.

35 U.S.C. § 103 : Obviousness

U.S. Patent No. 10,452,816 B2, claim 1, similarly teaches patient-facing conversational agents that are intentionally configured to emulate empathetic human dialogue, using tone calibration and pacing strategies commonly used by medical professionals.

U.S. Patent No. 9,824,188 B2, in Claim 1, further describes an AI system for engaging users in structured medical conversations, designed to mimic the behaviour and role of a live healthcare provider through language choice and content progression.

Furthermore, **U.S. Patent No. 11,694,807 B2** details a clinical decision-guidance interface wherein the agent follows role-based task flows aligned with typical human healthcare consultations, such as symptom gathering, triage, and treatment recommendation.

When considered together, these teachings render the emulation of a healthcare professional in a conversation interface not only obvious but predictable and expected in this domain.

Non-Patent Literature:

Nexus Flow emphasizes that error messages should be “clear, helpful, and actionable” without technical jargon, and should guide the user through recovery steps. It highlights that a well-designed fall-back message is critical for retaining user trust and ensuring successful task completion.

Surani & Das echo this need for patient-centric dialog recovery, noting that clear and courteous responses prevent abandonment and reinforce safety perceptions. The JAMA authors similarly advocate for transparent communication when automation encounters uncertainty. Given these overlapping design imperatives, the use of human-readable fall-back messages with next-step prompts was an expected and well-documented design pattern.

Therefore, the differences between Claim 10 and the prior art (if any) are mere implementation details, not inventive distinctions.

35 U.S.C. § 112(a) : Lack of Enablement and Written Description

Claim 10 also suffers from a lack of enabling disclosure regarding what it means for a machine to "emulate a human healthcare professional." The claim and specification do not define the parameters for emulation, whether it refers to voice, syntax, clinical expertise, bedside manner, or regulatory compliance.

Nor is there any explanation of how such emulation is achieved using the underlying language model or training set. Without detail, a skilled artisan is left without guidance, thereby failing the standard under § 112(a).

35 U.S.C. § 112(b) : Indefiniteness

The term “emulates a human healthcare professional” is inherently vague. The claim does not define the scope or the criteria used to determine whether emulation has occurred. Does it require passing a Turing test? Adherence to clinical guidelines? Replicating empathy?

Absent objective metrics or defined limits, the claim's scope cannot be reasonably ascertained, and it is therefore indefinite under § 112(b).

35 U.S.C. § 101 : Lack of Patent-Eligible Subject Matter

At its core, claim 10 describes an abstract goal: mimic human healthcare conversation. This is a long-standing practice in user interface design and telemedicine workflows. The claim implements this idea using a conventional conversational agent architecture, without any new technological improvement to LLM design, control logic, or conversation management.

Under Alice, this constitutes an abstract idea implemented on generic computing infrastructure, failing both steps of the eligibility test and rendering the claim invalid under § 101.

Claim 11

The multi-turn conversational system of claim 10, wherein the trigger exchanges include symptom descriptors that identify one or more symptoms.

35 U.S.C. § 103 : Obviousness

U.S. Patent No. 10,282,512 B2, claim 1, teaches a modular AI system that assigns meaning and routing priority to symptom-based utterances, automatically enabling safety checks and clinical classification procedures.

U.S. Patent No. 9,824,188 B2 teaches wherein the trigger exchanges include symptom descriptors that identify one or more symptoms (Fig.6 (114) and Col.14, Line 44-60).

Thus. Claim 11 does not represent a non-obvious advancement and is unpatentable under 35 U.S.C. § 103.

Non-Patent Literature:

The Nexus Flow Innovations article explicitly recommends the collection of key fall-back performance metrics such as error frequency, fall-back trigger rates, recovery success rates, and user satisfaction post-failure. These measurements are framed as essential for diagnosing system deficiencies and guiding improvements.

Additionally, the CHI Workshop paper by Surani & Das emphasizes system transparency and the need for measurable security and reliability posture in chatbot infrastructure. When considered alongside JAMA's guidance that healthcare AI tools require ongoing evaluation and risk monitoring, it becomes clear that tracking operational metrics is routine.

Therefore, the differences between Claim 11 and the prior art (if any) are mere implementation details, not inventive distinctions.

35 U.S.C. § 112(a) : Lack of Enablement and Written Description

Claim 11 fails to provide sufficient written description or enabling guidance for implementing "symptom descriptors" in a manner that distinguishes them from ordinary medical terms. The claim lacks specificity about whether these descriptors are predefined keyword sets, machine-learned embeddings, or structured vocabularies (e.g., SNOMED-CT).

Furthermore, the specification does not describe the linguistic boundaries or validation mechanisms for identifying such descriptors, leaving the skilled artisan with no implementable instruction or standard for compliance.

35 U.S.C. § 112(b) : Indefiniteness

The phrase “symptom descriptors” lacks clear boundaries. It is unclear whether the term refers to pre-programmed vocabulary, dynamically inferred expressions, or free-text phrases recognized by the LLM. Without clarity on scope or objective measurement, the claim fails to provide sufficient notice to the public and therefore is indefinite under § 112(b).

35 U.S.C. § 101 : Lack of Patent-Eligible Subject Matter

Claim 11 remains directed to the abstract idea of identifying medical symptoms in user dialogue, a cognitive task historically performed by human clinicians. The implementation, as described, merely automates that task using well-understood conversational AI methods.

As with Claim 10, this amounts to an abstract idea executed by generic computing functionality. There is no technological improvement to natural language processing, symptom recognition, or AI conversation management. Thus, claim 11 fails both prongs of the Alice framework and is invalid under § 101.

Claim 12

The multi-turn conversational system of claim 11, wherein the questions comprise medical condition queries that probe at least one symptom mentioned by the human and identified by at least one symptom descriptor associated with the trigger exchanges.

35 U.S.C. § 103 : Obviousness

Even if viewed as a novel combination, claim 12 would have been obvious to a person of ordinary skill in the field of medical conversational AI.

U.S. Patent No. 9,824,188 B2 teaches using medical condition queries that probe at least one symptom mentioned by the human and identified by at least one symptom descriptor associated with the detected trigger exchanges (Fig.6 (114) and Col.14, Line 44-60).

U.S. Patent No. 10,282,512 B2, claim 6, describes a dynamic diagnostic dialogue engine that customizes medical questioning based on symptom descriptors extracted from prior user input.

Thus, Claim 12 is obvious in view of the identified prior art.

Non-Patent Literature:

Nexus Flow instructs that "proactive error prevention is as important as reactive handling" and identifies input validation and context verification as primary means of preventing cascading errors. It suggests pre-processing user input to check for malformed, unsafe, or incomplete data and ensuring context continuity across dialog turns.

Surani & Das further confirm that input screening is critical in secure healthcare chatbots, citing examples where improper input validation led to user data leakage or classification failures. In light of the safety concerns raised in the JAMA article about AI-driven patient interactions, proactive input validation is expected rather than innovative.

Therefore, the differences between Claim 12 and the prior art (if any) are mere implementation details, not inventive distinctions.

35 U.S.C. § 112(a) : Lack of Enablement and Written Description

The patent specification offers no sufficient written description for how the system distinguishes between general follow-up questions and "medical condition queries." It also

fails to explain how such queries are programmatically tied to specific symptom descriptors or diagnostic pathways.

Without detail on the linkage logic, whether keyword rules, LLM embeddings, or diagnostic trees, claim 12 cannot be said to be fully enabled or described.

35 U.S.C. § 112(b) : Indefiniteness

The phrase "medical condition queries that probe at least one symptom" is vague and lacks objective boundaries. It is not clear whether "probe" means a clinical question, a binary prompt, or a subjective inquiry. Moreover, the standard for "associated with the trigger exchanges" is never defined, creating uncertainty in claim scope.

Accordingly, the claim is indefinite and fails to satisfy the notice requirement of § 112(b).

35 U.S.C. § 101 : Lack of Patent-Eligible Subject Matter

Like the claims before it, claim 12 is abstract. It automates the human process of asking medical follow-up questions based on symptoms. There is no technological innovation in the query-generation process, the recognition logic, or the interface itself.

The claim amounts to conventional LLM-enabled questioning based on recognized symptom patterns. No inventive concept rescues the claim under Alice. It is therefore patent-ineligible under § 101.

Claim 13

The multi-turn conversational system of claim 12, wherein the answers identify medical conditions associated with the at least one symptom and experienced by the human.

35 U.S.C. § 103 : Obviousness

US Patent No. 9,824,188 B2 teaches wherein the received answers identify one or more medical conditions associated with the symptom and experienced by the human (Fig.6 (114) and Col.14, Line 44-60, the user responses indicate influenza).

U.S. Patent No. 10,282,512 B2 discloses a recursive condition-matching module that continuously refines its diagnostic inference based on newly acquired conversation input. Each answer provided by the user is processed through a hierarchical diagnostic model, enabling identification of likely conditions associated with that symptom. The system is optimized to conduct real-time condition recognition embedded within a multi-turn exchange, thus rendering Claim 13 completely obvious.

Consequently, claim 13 is unpatentable under 35 U.S.C. § 103 as obvious in view of the identified prior art.

Non-Patent Literature:

The JAMA Viewpoint insists that any AI medical system must include provisions for human review and override when the scope of automation is exceeded, particularly in safety-critical decisions.

Further, the Nexus Flow article explains that escalation to human agents is a necessary final fall-back step and should occur seamlessly while preserving dialog context. Surani & Das also support human-in-the-loop models as an essential design requirement for trust, accountability, and compliance with clinical best practices.

Therefore, the differences between Claim 13 and the prior art (if any) are mere implementation details, not inventive distinctions.

35 U.S.C. § 112(a) : Lack of Enablement and Written Description

Claim 13 is invalid under § 112(a) because the ’371 Patent fails to disclose how answers provided by a user are transformed into recognized medical conditions. While the claim implies that condition identification occurs as part of the LLM conversation flow, the specification offers no technical structure, data schema, or classifier architecture capable of enabling such mapping. There is no description of how the system determines what condition is associated with which symptom, nor any algorithmic strategy for disambiguating competing diagnoses. The patent assumes the existence of a classification capability without teaching how it is realized, leaving a person of ordinary skill without the ability to practice the invention without undue experimentation.

35 U.S.C. § 112(b) : Indefiniteness

Claim 13 is indefinite because it provides no objective boundaries for the term “identify medical conditions.” It is unclear whether this requires deterministic condition outputs, probabilistic likelihoods, differential diagnostic suggestions, or merely flagging broad condition categories. The absence of definitional clarity or performance criteria for what qualifies as identification creates ambiguity regarding both the system’s responsibilities and the scope of protection. A skilled artisan is left uncertain whether the claim encompasses partial inferences, tentative hypotheses, or conclusive determinations.

35 U.S.C. § 101 : Lack of Patent-Eligible Subject Matter

Claim 13 is directed to the abstract idea of evaluating a user’s responses in a conversation and inferring the medical condition that best matches the symptoms described, an activity long performed by physicians and triage nurses. The claimed implementation relies on generic conversational AI components and standard classifier logic. No inventive concept is introduced

with respect to model structure, inference technique, or diagnostic methodology. The claim merely automates a long-established cognitive task using conventional computing tools. Under the Alice framework, the claim is directed to an abstract idea and lacks any meaningful transformation into a patent-eligible invention.

Claim 14

The multi-turn conversational system of claim 13, wherein the answer classification logic is further configured to classify the answers and generate the control signals based on a severity of the medical conditions.

35 U.S.C. § 103 : Obviousness

US Patent No. 9,824,188 B2 teaches wherein the answer classification logic is further configured to classify the received answers and generate the control signals based on severity of the identified medical conditions (see Fig.6 (114) and Col.14, Line 44-60, recommend emergency medical assistance when symptoms indicate a serious medical problem).

In addition, U.S. Patent No. **10,282,512 B2**, claim 2, describes a rule-based classifier that not only identifies possible medical conditions from input symptoms but also applies condition severity weightings to control the progression of multi-turn conversation branches. These control paths are prioritized based on inferred risk scores, directing user flow through different escalation routes or recommendation logic. The system structurally mirrors the claim by integrating condition severity directly into signal generation and dialog control.

Non-Patent Literature:

Surani & Das specifically identify two-factor authentication (2FA) as a well-known and highly recommended mechanism to protect healthcare chatbot systems from unauthorized access.

They describe combining credentials (such as passwords) with device-based or time-limited codes as a routine method of securing patient conversations.

Moreover, privacy guidelines referenced in the JAMA article stress robust identity verification measures to prevent misuse and data breaches. In light of these references, integrating 2FA into a patient-facing chatbot was a predictable design step for secure access control.

Therefore, the differences between Claim 14 and the prior art (if any) are mere implementation details, not inventive distinctions.

35 U.S.C. § 112(a) : Lack of Enablement and Written Description

The ’371 Patent does not disclose any framework, scale, or algorithm for assessing the severity of medical conditions. While the claim assumes the existence of a severity classifier, the specification provides no model architecture, medical severity index, or source of clinical labelling. Nor does it disclose how severity influences control signals. The system is described at a high level, but the crucial technical underpinnings, how severity is computed, applied, and operationalized, are absent.

This results in a failure of enablement, as a skilled practitioner cannot implement the claimed logic without resorting to external knowledge and substantial experimentation.

35 U.S.C. § 112(b) : Indefiniteness

Claim 14 is indefinite in its use of the phrase “based on a severity of the medical conditions.” There is no indication whether severity is binary, scalar, condition-specific, time-dependent, or derived from symptom intensity. The specification lacks a metric or benchmark for what qualifies as “severity,” nor does it define how such severity quantitatively or qualitatively

affects control signal generation. These omissions make it impossible to determine the metes and bounds of the claim.

35 U.S.C. § 101 : Lack of Patent-Eligible Subject Matter

Claim 14 is directed to the abstract idea of ranking conditions by severity and using those rankings to influence dialog behaviour, an inherently mental process long carried out by clinicians. Automating this process using conventional AI classifiers and symptom databases does not constitute a technological improvement. No novel diagnostic method, classifier framework, or computing efficiency is claimed. The implementation is a straightforward automation of a human judgment task, using generic computational tools and lacking an inventive concept. The claim is thus ineligible under § 101.

Claim 15

The multi-turn conversational system of claim 14, wherein the control signals, generated in response to the answer classification logic classifying the answers to a second kickout category, cause the conversation interface to connect the user with a human healthcare professional.

35 U.S.C. § 103 : Obviousness

U.S. Patent No. 9,824,188 B2 describes a conversational healthcare assistant that escalates user responses to a human professional when triggered by critical input patterns. The assistant performs this escalation based on user dialogue classified into elevated concern levels. The fall-back pathway is explicitly designed to transfer control from the AI system to a qualified healthcare provider, ensuring safe handoff for sensitive cases (Fig.6 (114) and Col.14, Line 44-

60, initiate call for emergency medical assistance when symptoms indicate a serious medical problem).

Furthermore, **U.S. 11,694,807 B2**, claim 1, discloses an AI-enabled workflow engine that conditionally routes cases to human experts when ambiguity, escalation, or diagnostic uncertainty is detected.

A person of ordinary skill would have readily implemented this fall-back structure in any system seeking to ensure safety and accountability in healthcare AI environments.

Non-Patent Literature:

Surani & Das describe end-to-end encryption as a foundational requirement for safeguarding sensitive health information in chatbot interactions. They outline the use of symmetric or asymmetric encryption schemes to ensure that only the intended parties can access the exchanged data.

Furthermore, JAMA flags data privacy as one of the core risks in AI-driven healthcare tools, reinforcing the expectation that encryption should be implemented by default. Given the sensitive nature of patient communication, it was obvious to secure it using established encryption protocols.

Therefore, the differences between Claim 15 and the prior art (if any) are mere implementation details, not inventive distinctions.

35 U.S.C. § 112(a) : Lack of Enablement and Written Description

The ’371 Patent fails to describe how the “second kickout category” is defined or distinguished from other categories. There are no explanation of how the classification logic maps inputs to

discrete categories, what algorithmic rules govern the thresholds, or how the system determines that escalation to a human professional is necessary. The absence of structural details, decision criteria, or fall-back orchestration logic results in a failure to enable the claimed escalation mechanism.

35 U.S.C. § 112(b) : Indefiniteness

The phrase "second kickout category" is indefinite. The claim provides no criteria for distinguishing this category from others. It is unclear whether it refers to a specific numerical index, severity band, or decision tree outcome. The term lacks antecedent basis, is not defined in the specification, and leaves a person of ordinary skill unable to ascertain the boundaries of the claim. Furthermore, the mechanics of "causing the conversation interface to connect the user with a human healthcare professional" are not clearly delineated in terms of system behaviour or implementation scope.

35 U.S.C. § 101 : Lack of Patent-Eligible Subject Matter

Claim 15 recites an abstract escalation logic wherein an AI system, upon classifying certain user inputs as critical, refers the user to a human, a concept indistinguishable from longstanding triage practices. Implementing this using LLMs and signal triggers does not add any technological advancement. The claim merely automates a mental process using standard components and introduces no inventive concept in classifier structure, control logic, or communication interface. As such, the claim fails both steps of the § 101 inquiry and is directed to ineligible subject matter.

Claim 16

The multi-turn conversational system of claim 15, wherein the control signals, generated in response to the answer classification logic classifying the answers to a third kickout category, cause the conversation interface to schedule a healthcare appointment of the user with a human healthcare professional.

35 U.S.C. § 103 : Obviousness

Claim 16 is rendered obvious by U.S. Patent No. **9,824,188 B2**, which discloses, a conversational virtual healthcare assistant that interprets user input, classifies it according to risk and urgency, and upon reaching designated escalation thresholds, automatically schedules a medical appointment with a licensed healthcare provider. The classification system incorporates a tiered response logic, whereby a specific category of input, distinct from emergency triage, results in the invocation of a scheduling subroutine embedded within the system. The scheduling is directly triggered by the classification outcome, and not merely by a manual request, thus satisfying the recited linkage between classification and automated appointment creation (Fig.6 (114) and Col.14, Line 44-60, scheduling an appointment with a doctor based on the symptoms indicating influenza).

Additionally, **U.S. 11,694,807 B2**, claim 3, teaches a workflow-based decision-making engine that routes outputs not just to content responses but also to action pipelines, including administrative operations such as scheduling, escalation, or care coordination. A skilled artisan would have viewed integration of such scheduling workflows into a classification-triggered LLM interface as a predictable and non-inventive enhancement.

Therefore, claim 16 is rendered obvious under 35 U.S.C. § 103 in light of the disclosures in identified prior art.

Non-Patent Literature:

The CHI Workshop paper by Surani & Das discusses the importance of ephemeral messaging and timed message deletion in chat-based healthcare tools, emphasizing its use to reduce retention of sensitive personal health information. The paper recommends this technique to enhance compliance with privacy regulations and reduce exposure from long-term storage of chat logs.

Further, the JAMA Viewpoint flags persistent data storage as a major concern in patient-AI interactions and suggests minimizing data retention to mitigate risk. In light of these known security practices, the inclusion of time-bound message deletion was an expected design strategy rather than an inventive feature.

Therefore, the differences between Claim 16 and the prior art (if any) are mere implementation details, not inventive distinctions.

35 U.S.C. § 112(a) : Lack of Enablement and Written Description

The ’371 Patent fails to enable the mechanism by which control signals are mapped to healthcare appointment scheduling workflows. There is no disclosed architecture for appointment handling, no explanation of how classified answers trigger specific scheduling parameters (e.g., urgency, provider type, calendar availability), and no description of integration with medical scheduling systems or APIs. The claim recites a system behaviour, appointment scheduling tied to classification output, without supplying the necessary enabling detail for implementation.

35 U.S.C. § 112(b) : Indefiniteness

The term “third kickout category” is indefinite. The claim does not define whether the kickout categories are ordinal, probabilistic, severity-tiered, or rule-encoded. It is unclear how this third

category differs from the second kickout category in Claim 15 or whether they are mutually exclusive. Nor does the claim provide operational clarity on what "cause the conversation interface to schedule" entails, whether this includes proposing times, confirming via API, or merely generating a referral. These ambiguities preclude a clear understanding of claim boundaries.

35 U.S.C. § 101 : Lack of Patent-Eligible Subject Matter

Claim 16 describes the abstract concept of using a classification result to automatically schedule a medical appointment, an activity that reflects longstanding clinical workflows. Implementing this process with a conversational system and control logic does not introduce a technological improvement or novel algorithm. The claimed behaviour is a straightforward automation of routine healthcare administration and does not reflect any inventive concept. Accordingly, the claim is directed to an abstract idea and fails the eligibility requirements of § 101.

Claim 17

The multi-turn conversational system of claim 1, wherein the human-like conversation is text-based.

35 U.S.C. § 103 : Obviousness

U.S. Patent No. **9,824,188 B2** describes a virtual healthcare assistant that conducts multi-turn dialogues with users through a chat-based interface (Claim 2, Fig.3 and Col.11, Line 20-33)). The assistant receives patient responses in text form, processes them using a conversational engine, and returns contextually aware follow-ups, constituting a multi-turn, LLM-mediated

textual interaction. The claimed interface is deployed via browser or mobile applications, explicitly designed for text-first engagement.

U.S. Patent No. **10,452,816 B2** discloses a patient engagement platform that uses text-based interaction for digital symptom collection (Block S100). The system guides users through clinical pathways using text prompts and decision trees embedded in a chat format. These interactions take place over multiple turns and are performed entirely through textual UI, satisfying the recitation of a "human-like conversation" in text modality.

Therefore, claim 17 is at least obvious in view of the teachings of identified prior art under 35 U.S.C. § 103.

Non-Patent Literature:

Surani & Das specifically describe the use of MAC address authentication and device whitelisting as a viable security measure for chatbot systems, especially in healthcare, to prevent session hijacking or credential misuse. They recommend this technique in scenarios where device exclusivity or physical custody of an endpoint must be enforced.

Given the elevated privacy risks in clinical chat environments and the demand for strict identity and session control, as emphasized in the JAMA Viewpoint, tying chatbot access to a secure device ID was not an innovation but a routine application of network security principles.

Therefore, the differences between Claim 17 and the prior art (if any) are mere implementation details, not inventive distinctions.

35 U.S.C. § 112(a) : Lack of Enablement and Written Description

While Claim 17 may be considered structurally simple, the ’371 Patent offers no discussion on how the system renders text-based interfaces across devices, manages input-output cycles, or handles interruptions or latency. Although minor, these omissions represent a failure to describe the operational requirements of deploying a “human-like” text conversation, particularly one involving medical triage or safety validation.

35 U.S.C. § 112(b) : Indefiniteness

The phrase “human-like conversation is text-based” lacks precision. It is unclear whether the claim requires strictly textual input and output or permits mixed modalities (e.g., text input with visual or structured responses). Additionally, “human-like” is not quantitatively defined and could vary based on tone, grammar, vocabulary, or empathy modelling. These ambiguities introduce subjectivity into the claim scope.

35 U.S.C. § 101 : Lack of Patent-Eligible Subject Matter

Claim 17 recites the abstract concept of using textual communication for a conversation. The act of exchanging information via text has long been implemented in chatbot systems and digital forms. The claim does not recite any technical improvement in the interface, communication protocol, or LLM behaviour. It merely limits the interaction to a non-novel modality using standard tools. As such, the claim is directed to an abstract idea and fails to meet the requirements of § 101.

Claim 18

The multi-turn conversational system of claim 1, wherein the human-like conversation is audio-based.

35 U.S.C. § 103 : Obviousness

U.S. Patent No. **9,824,188 B2** discloses a healthcare assistant platform that supports both text and voice-based input modalities (Claim 14). The conversational engine integrates speech recognition and synthesis modules to enable users to speak their inputs and receive voice-based responses, all within the same multi-turn, LLM-mediated dialog flow. The system is specifically designed to replicate a human-like, speech-based medical interaction using a microphone and speaker interface, fulfilling the recitations of Claim 18.

Additionally, **U.S. 11,694,807 B2**, claim 1, describes decision-making workflows integrated with multimodal input/output, including auditory pathways. A person of ordinary skill would have recognized the functional and technical interchangeability of text and audio interfaces within the same LLM-driven structure and would have considered such configuration an obvious extension (see Fig.7A and Col.11, Line 20-33).

Since the implementation of audio-based communication represents a predictable use of known technologies for their intended purpose, claim 18 is at least obvious under 35 U.S.C. § 103.

Non-Patent Literature:

The CHI Workshop paper by Surani & Das describes privacy-preserving modules that extract and mask sensitive entities (such as names, dates of birth, or clinical codes) before sending data to back-end processors. The paper notes this as a key safeguard for de-identifying input in healthcare chatbots.

Further, the JAMA Viewpoint emphasizes the importance of minimizing exposure of sensitive user content and promotes proactive filtration techniques to reduce accidental disclosure during automated processing. Given the clear teaching that PII should be masked before query submission, the claimed entity-filtering behaviour is an expected feature in secure AI systems.

Therefore, the differences between Claim 18 and the prior art (if any) are mere implementation details, not inventive distinctions.

35 U.S.C. § 112(a) : Lack of Enablement and Written Description

The ’371 Patent does not disclose how audio-based interaction is technically achieved. There is no mention of speech-to-text or text-to-speech modules, no training considerations for voice input, and no interface protocol for synchronizing auditory input with control logic. The absence of such detail renders the claimed functionality under-enabled, especially for safety-critical healthcare settings where timing, clarity, and misrecognition risk must be addressed explicitly.

35 U.S.C. § 112(b) : Indefiniteness

The phrase “human-like conversation is audio-based” is ambiguous in scope. It is unclear whether the entire dialog must occur via audio or whether the modality is optional. The claim does not define how the system handles partial audio, mixed-mode interactions, or silent fall-back. Nor does it distinguish conversational realism from mere voice command processing. These vagaries render the scope of the claim uncertain.

35 U.S.C. § 101 : Lack of Patent-Eligible Subject Matter

Claim 18 describes the abstract idea of conducting a human-like conversation using voice, an interaction modality long practiced and, in recent decades, broadly implemented via digital assistants and IVR systems. The claim introduces no novel audio-processing architecture, latency optimization, or speech synthesis model. It merely specifies a conventional input/output mode using known tools. As such, the claim is directed to an abstract idea and lacks the inventive concept required for eligibility under § 101.

Claim 19

A system, comprising:

(1) a conversation interface that is based on a first large language model and is configured to execute multiple turns of human-like conversation with a user based on a checklist that identifies tasks that the first large language model is to complete during the human-like conversation;

(2) a control logic, in communication with the conversation interface, and configured to generate one or more control signals based on evaluating multiple turns of upstream human-like conversation between the conversation interface and the user; and

(3) a second opinion module comprising an additional large language model to perform research into a portion of the human-like conversation in parallel with the human-like conversation taking place between the conversation interface and the user.

Observations:

Claim 19 is effectively a restatement of independent Claim 1, reorganized slightly with minor language adjustments. It similarly recites the three primary components of the system:

1. A **conversation interface** powered by a first LLM and guided by a checklist;
2. A **control logic** that evaluates upstream turns and produces control signals;
3. A **second opinion module** comprising an additional LLM performing research in parallel.

Functionally and structurally, claim 19 is substantively identical to Claim 1. Accordingly, all prior art references cited against Claim 1 also apply fully to this claim.

35 U.S.C. § 103 : Obviousness

Claim 19 is rendered obvious in view of the combination of **U.S. 9,824,188 B2, U.S. Patent Publication No. 2024/0185968 A1, and U.S. Patent Publication No. 2023/0245651 A1**. The '188 patent, in Claim 1 and Claim 3, introduces a medical chatbot capable of managing multi-turn dialog, user intent determination, and performing dynamic insertions based on conversation state. The '968 discloses checklist-based query system (Para [0086] - [0087]) performing predefined dialogue with a user. The '651 patent discloses a parallel running module (e.g., analysis module, multimodal generation module (Para [0157] – Para [0164])) which is configured to collaborate with conversational AI agent and perform in-depth analysis and generate recommendations.

It would have been obvious to a person having ordinary skill in the art (PHOSITA) to combine the medical chatbot in the '049 patent with the checklist enforcement system in the '968 patent and '651 patent. Doing so would result in a system comprising (1) a checklist-guided conversation interface, (2) a logic module that drives dialog branching based on upstream context, and (3) a parallel second-opinion mechanism, all exactly as claimed in Claim 19 of the '371 Patent.

Further support arises from **U.S. 11,694,807 B2**, where Claim 1 and Claim 3 disclose a system with modular decision-making units designed to generate control signals in response to user dialog and escalate ambiguous exchanges for validation. A PHOSITA would have readily incorporated a second-opinion LLM into this structure as suggested in **U.S. 2023/0245651 A1** to strengthen diagnostic accuracy or compliance in regulated fields like healthcare.

U.S. 11,532,393 B2 already teaches a dual-track safety architecture: raw clinical data flow through a validation layer that checks each item against predefined rules and then passes a

curated stream to downstream AI. **U.S. 11,532,393 B2** supplies the complementary rule-driven control logic, monitoring a live schedule, evaluating each update against alert thresholds, and issuing machine-readable "highlight" or mobile-alert signals that steer subsequent workflow. When these two references are read together, every functional building block of the '371 claims is present: (i) a primary data-exchange channel governed by a checklist/ template, (ii) a control module that inspects prior state and emits downstream directives, and (iii) a parallel verification path that ensures only vetted information influences the next action. The sole distinction in the '371 patent is that the front-end happens to be a large-language-model chat interface rather than a schedule dashboard or a data-export driver.

Accordingly, claim 19 is at least obvious under 35 U.S.C. § 103.

The configuration reflects a standard safety-augmented AI framework in the medical domain and would have required no inventive insight to implement.

Non-Patent Literature:

Surani & Das introduce the concept of searchable encryption as a privacy-preserving mechanism in chatbot systems, describing how queries can be processed by back-end models or search tools without decrypting the full input content. The paper explains symmetric and asymmetric methods to allow keyword matching over encrypted queries.

This method is discussed in the context of health AI tools handling sensitive patient queries and is offered as a practical solution to privacy concerns in cloud-based inference. Given these disclosures and the JAMA Viewpoint's emphasis on preserving user confidentiality in all AI operations, incorporating searchable encryption was an expected solution to a known problem.

Therefore, the differences between Claim 19 and the prior art (if any) are mere implementation details, not inventive distinctions.

35 U.S.C. § 112(a) : Lack of Enablement and Written Description

Claim 19 sweeps broadly, yet the specification supplies only high-level rhetoric—not the concrete technical teaching the statute demands.

1. The "second-opinion module" is an unexplained black box.

The claim hinges on a second LLM that "performs research … in parallel" with the live conversation, but the patent never tells a skilled artisan how that research engine is built or run. The disclosure is silent on

a. the training corpus, parameter count, or fine-tuning regimen of the second model;

b. the real-time data feed it receives (full transcript? selective snippets? embeddings?);

c. the synchronization protocol that prevents it from lagging behind or racing ahead of the primary chat;

d. conflict-resolution rules when the second model's findings diverge from the first model's output; and

e. resource management for GPU/CPU, memory, and context-window budgeting when two transformer models contend for the same hardware.

Implementers are left to invent those critical details themselves—precisely the "undue experimentation" the Wands factors prohibit (factors 6-8: amount of guidance, state of the art, and breadth of claims).

2. “Control logic” is claimed purely by result.

The patent never discloses any algorithm, data structure, or rule hierarchy for the control logic that is supposed to (i) read upstream turns, (ii) generate “one or more control signals,” and (iii) inject them into downstream conversation. Are the signals token prefixes, system prompts, numeric weights, or API flags? How are they prioritized when multiple signals compete? How are they inserted without exceeding the LLM’s context window or corrupting the checklist state? The document says nothing, violating *Ariad v. Lilly*, which bars “functional claiming without telling how to achieve the function.”

3. Checklist enforcement is hand-waved.

Claim 19 requires the first LLM to work “based on a checklist,” yet the specification never explains how the checklist is encoded, how completed items are tracked, or how missing items trigger follow-up. In a transformer limited by token budgets and prone to context drift, these implementation details are essential; their omission forces the reader to guess.

4. The claim’s breadth dwarfs the disclosure.

By reciting generic “large language models,” the claim purports to cover every transformer family—GPT, Llama, PaLM, Claude, and future variants—yet the patent contains no exemplification beyond a cursory statement that the first model underwent “over 1000 gradient updates,” a trivially low and technologically meaningless threshold. *Amgen v. Sanofi* (2023) confirms that such over-breadth, absent commensurate guidance, flunks both possession and enablement.

5. No working example or experimental data.

There is not a single flowchart, timing diagram, or pseudo-code listing showing the two-model system in operation. The Federal Circuit routinely invalidates claims where, as here, the specification "merely recites a desired result" without teaching a concrete way to achieve it (*Cisco v. Cirrex*; *ALZA*).

Because a person of ordinary skill would have to architect parallel inference pipelines, devise bespoke conflict-resolution logic, and experiment extensively just to make the claimed system work, the patent fails § 112(a)'s twin requirements of enablement and written description. Claim 19 should therefore be rejected.

35 U.S.C. § 112(b) : Indefiniteness

Claim 19 fails the *Nautilus* "reasonable certainty" test because its core terms are so fluid that neither competitors nor the PTO can discern the claim's outer bounds.

1. "Perform research into a portion of the human-like conversation."

The patent never explains what "research" entails. Does the second model (i) issue live web queries, (ii) retrieve vectors from an internal embedding store, (iii) consult a proprietary clinical database, or (iv) simply re-prompt itself for a fresh answer? Each option imposes very different architectures, latencies, and privacy requirements, yet the claim sweeps over all of them without guidance. The Federal Circuit has condemned such open-ended functional language that leaves "the public to guess about which implementations infringe" (*Interval Licensing v. AOL, 766 F.3d 1364*).

2. Vagueness of "a portion."

The claim offers no metric for what constitutes "a portion" of the conversation. Is it a single token, a user turn, a topic cluster, or any arbitrary text span the implementer chooses? Without

an objective delimiter, two implementers could build materially different systems—one that re-checks every utterance, another that samples only the final summary—yet both would ostensibly fall within (or outside) the claim at a court's discretion. Such subjectivity violates § 112(b), which requires clear claim boundaries (*Nautilus, Inc. v. Biosig, 572 U.S. 898*).

3. Temporal ambiguity in "in parallel."

The phrase "in parallel with the … conversation" lacks an antecedent baseline: Does "parallel" require simultaneous token-level execution, overlapping but asynchronous calls, or merely any research performed before the chat ends? Because the claim provides no timing diagram, buffer protocol, or concurrency definition, an accused product can scarcely know if its background processes operate "in parallel" or merely "in series," fostering inevitable dispute.

4. Functional black-box language for the "second-opinion module."

By naming a component only by what it should accomplish ("perform research") and not by any concrete structure or algorithm, the patentee has invoked a pure results-oriented nonce term—precisely the indefiniteness problem highlighted in *Williamson v. Citrix, 792 F.3d 1339*. With no algorithm disclosed for selecting the research target, fetching data, or merging findings, the claim leaves implementers and courts without a yardstick for compliance.

5. Enforcement impossibility.

Because "research," "portion," and "in parallel" have no objective test, infringement analysis would devolve into after-the-fact expert testimony on what "research" ought to mean. Claims that require such litigation guesswork are the textbook example of indefiniteness (*Halliburton Energy v. M-I LLC, 514 F.3d 1244*).

For these reasons, Claim 19 does not "inform those skilled in the art, with reasonable certainty, what the invention is." Its shifting terminology renders the scope indeterminate, and the claim must therefore be rejected under 35 U.S.C. § 112(b).

35 U.S.C. § 101 : Lack of Patent-Eligible Subject Matter

Claim 19 is patent-ineligible under 35 U.S.C. § 101 because it merely automates an age-old mental routine—follow a script, listen to the answer, decide what to ask next, and, if unsure, consult a reference—using standard computer components. Its three elements (a checklist-driven conversation interface, a "control logic" that reviews prior turns, and a background "second-opinion" language model that double-checks facts) map exactly to activities the Federal Circuit has repeatedly labeled abstract: organizing human interactions, mental observation and judgment, and routine data gathering (*buySAFE, Cybersource, Electric Power Group, Content Extraction*). Nothing in the claim changes how a computer stores, processes, or transmits data; it simply instructs off-the-shelf software to mimic human conversational behavior.

Because the claim is directed to an abstract idea, it can survive only by supplying an "inventive concept," yet it offers none. It name-checks two large-language models but discloses no novel architecture, tokenizer, training scheme, or integration protocol—indeed, the specification boasts merely that the first model received "over 1000 gradient updates," a trivially low bar for modern systems. The "control logic" is defined solely by the result it must achieve ("generate control signals"), a purely functional recital the courts deem insufficient (*Two-Way Media*; *Intellectual Ventures v. Capital One*). Running a second model "in parallel" is a textbook use of conventional threads; the claim introduces no new data structure, synchronization method, or hardware arrangement comparable to the self-referential table that saved the claims in *Enfish* or the specific phoneme-timing rules upheld in *McRO*.

Nor does Claim 19 pass the machine-or-transformation test: it is tied to no particular machine beyond an unspecified server, and it transforms no data into a different state or thing. Granting exclusivity over so broad and result-oriented a description would pre-empt the basic practice of checklist-guided conversation supplemented by real-time fact-checking—the very overreach § 101 is meant to prevent.

In short, Claim 19 repackages the abstract concept of reflective dialogue management and secondary verification into generic computer language. Lacking any concrete technological advance or inventive concept, it is ineligible under the Alice/Mayo framework and must be cancelled.

Claim 20

The system of claim 19, wherein the one or more control signals contribute in part to construction of multiple turns of downstream human-like conversation between the conversation interface and the user.

35 U.S.C. § 103 : Obviousness

U.S. Patent No. 9,824,188 B2 teaches wherein the control signals contribute in part to construction of multiple turns of downstream human-like conversation between the conversation interface and the user (see Fig.1 (152), Fig.6 (114), Col.9, Line 19-34 and Col.14, Line 44-67, responses to the queries or questions).

Therefore, claim 20 is at least obvious under 35 U.S.C. § 103.

Non-Patent Literature:

Surani & Das describe the use of smart cards in privacy-preserving e-prescription systems and healthcare identity verification, where the card stores sensitive data and enables controlled access to digital services like chatbots. Their review suggests smart-card authentication as a trusted and standardized security approach.

Further, healthcare systems globally already use patient-held smart cards to access EMR systems, and the JAMA Viewpoint flags secure user authentication as a mandatory layer in AI healthcare delivery. As such, the integration of smart-card-based access and storage into a chatbot framework is not inventive but rather a direct application of established best practices.

Therefore, the differences between Claim 20 and the prior art (if any) are mere implementation details, not inventive distinctions.

35 U.S.C. § 112(a) : Lack of Enablement and Written Description

The ’371 Patent fails to explain how control signals influence the generation of downstream dialog. It provides no mechanism for how such signals are parsed, weighted, prioritized, or injected into LLM prompts. The system is functionally claimed, but the core enabling disclosures, such as token injection strategies, model instruction mechanisms, or branching dialogue templates, are missing. Without this detail, a skilled artisan would not be able to implement the contribution of control signals to dialog flow without undue experimentation.

35 U.S.C. § 112(b) : Indefiniteness

Claim 20 is indefinite in that the phrase “contribute in part to construction” offers no clear boundary for how much influence the control signals exert. It is unclear whether the system must always use these signals, use them optionally, or merely log them. The term “construction

of multiple turns" is also vague, failing to define whether the signals determine content, order, branching, or logic flow. These semantic ambiguities render the scope of the claim uncertain.

35 U.S.C. § 101 : Lack of Patent-Eligible Subject Matter

Claim 20 is directed to the abstract idea of using logic derived from prior dialog to shape future conversational turns. This recursive, context-aware conversation modelling is a cognitive pattern well-known in human interaction and its automation through LLMs is not new. The claim does not recite any inventive method of dialog construction, control signal formatting, or LLM integration. It simply applies an abstract rule, "let past inputs influence future outputs", through standard AI techniques. This fails to rise to the level of a patent-eligible application under § 101.

VII. Supplemental Challenge – Alice §101 Abstract Idea Ineligibility

Supplemental § 101 Challenge — Why all twenty claims of U.S. Patent No. 12,142,371 B1 are ineligible.

The Supreme Court's *Alice Corp. v. CLS Bank, 573 U.S. 208 (2014)*, and its Federal-Circuit progeny establish a two-step inquiry for patent eligibility. At Step 1 we ask whether the claims are "directed to" an abstract idea; at Step 2 we ask whether they contain "something more" — an inventive concept that transforms that idea into a patent-eligible application. Read against that framework, every claim of the '371 patent collapses.

Alignment with July 2024 USPTO Guidance: The USPTO's July 2024 subject-matter eligibility guidance for computer-implemented and AI-related inventions reiterates the two-part Alice analysis as applied in the 2019 PEG: (i) identify whether the claim recites a judicial exception (e.g., a mental process or a certain method of organizing human activity) under Step

2A, Prong One; (ii) determine whether the claim integrates that exception into a "practical application" under Step 2A, Prong Two; and, if not, (iii) assess under Step 2B whether additional elements amount to "significantly more" than well-understood, routine, conventional (WURC) activity. The guidance cautions that merely invoking AI/ML components (such as an LLM) or stating desired results—without reciting a specific improvement in computer functionality—does not satisfy Prong Two or Step 2B. Applying those principles here confirms ineligibility: the claims recite mental-process/organizing-human-activity subject matter and fail to integrate that exception into any concrete technological improvement.

I. Step 1 — The entire claim set is directed to an abstract idea

Each claim, independent or dependent, recites the same non-technical mental workflow: (i) evaluate conversational content with a checklist of rules, (ii) score the safety or completeness of the exchange, and (iii) reroute or modify the response according to fixed thresholds. Courts classify such subject matter as abstract because it is nothing more than organizing human activity and mental observation and judgment. See *buySAFE, Inc. v. Google, 765 F.3d 1344 (Fed. Cir. 2014)* (rules for conducting a commercial interaction); *CyberSource v. Retail Decisions, 654 F.3d 1366 (Fed. Cir. 2011)* (collecting and analyzing data to make decisions); *Electric Power Group v. Alstom, 830 F.3d 1350 (Fed. Cir. 2016)* (monitor, analyze, and display information). Nothing in the claims alters the basic operation of a computer; rather, the computer is a mere venue for the age-old practice of reading, judging, and responding.

That conclusion does not change for the dependent claims. Some add pathways for ending the chat, escalating to a human, or scheduling an appointment (Claims 9–16); others limit the medium to text or voice (Claims 17–18) or require the second model to "research" a slice of the dialogue (Claims 3, 19, 20). These are all ancillary decision rules — still mental steps, still

organizational logic. As the Federal Circuit held in *In re Board of Trustees of Stanford, 991 F.3d 1245 (2021)*, re-running a statistical model to check results is itself an abstract data-analysis step, not a technological improvement.

Step 2A Prongs and Case Alignment

Under Step 2A, Prong One, the claims fall squarely within the "mental processes" and "certain methods of organizing human activity" categories identified by the USPTO. Under Step 2A, Prong Two, there is no integration into a practical application: the claims neither effect a specific improvement in computer capabilities nor tie the abstract idea to any particular machine operation beyond generic data intake/analysis/response. See *ChargePoint, Inc. v. SemaConnect, Inc., 920 F.3d 759, 770–74 (Fed. Cir. 2019)* (results-oriented networked functionality is abstract); *BSG Tech LLC v. BuySeasons, Inc., 899 F.3d 1281, 1286–89 (Fed. Cir. 2018)* (constraining user input with rules is still abstract); *FairWarning IP, LLC v. Iatric Sys., Inc., 839 F.3d 1089, 1093–95 (Fed. Cir. 2016)* (detecting improper records access via rules is abstract). As in *Interactive Wearables v. Polar Electro (Fed. Cir. 2023)*, merely tailoring information flow to users—without a computer-function improvement—is not a practical application.

II. Step 2 — The patent supplies no inventive concept

Because the claims are abstract, eligibility hinges on whether the specification discloses a particular technological advance. It does not.

Routine components. The only hardware is an unspecified "conversation interface" hosting generic transformer models, the very kind first disclosed to the public in Vaswani et al., Attention Is All You Need (2017). Even the "Vision Transformer" references are drawn wholesale from Dosovitskiy et al., 2020. The law is clear that "conventional machine-learning

components do not add an inventive concept merely by being invoked" (*IBM v. Zillow, 50 F.4th 1371 (Fed. Cir. 2022)*).

Purely functional claiming. The patent never explains how safety scores are computed, how checklist gaps are detected, or how the second model's "research" synchronizes with the main chat. Ineligibility follows when claims "recite performance of a function without reciting the means for accomplishing that function" (*Two-Way Media v. Comcast, 874 F.3d 1329 (Fed. Cir. 2017)*).

No technological improvement. Unlike the data-structure innovation upheld in *Enfish* or the precise phoneme-timing rules in *McRO*, the '371 claims leave underlying computer operations untouched. They do not speed model inference, reduce memory, compress tokens, or solve a networking bottleneck; they simply layer business rules atop an existing LLM.

Dependent claims fare no better. Limiting the communication to text or audio does not add an inventive concept — modality choice is "insignificant extra-solution activity." Re-casting the same control logic for symptoms, prescriptions, or family history (Claims 3, 11-14) only narrows the abstract idea's field of use, which the Supreme Court has said cannot confer eligibility (*Bilski v. Kappos, 561 U.S. 593 (2010)*).

Step 2B/WURC and Case Alignment

Under Step 2B, the additional elements are no more than WURC activity. See *Berkheimer v. HP Inc., 881 F.3d 1360, 1368 (Fed. Cir. 2018)* (conventionality is a factual question but may be resolved where the specification admits conventional components). Here, the specification's reliance on off-the-shelf LLMs and generic control logic—without reciting any novel data structure, training paradigm, or concurrency mechanism—confirms conventionality. See *SAP Am., Inc. v. InvestPic, LLC, 898 F.3d 1161, 1168–70 (Fed. Cir. 2018)* (advanced math/statistics implemented on a general-purpose computer lacks an inventive concept); *Univ. of Fla.*

Research Found., Inc. v. Gen. Elec. Co., 916 F.3d 1363, 1367–70 (Fed. Cir. 2019) (automating data manipulation in a medical context is abstract and lacks "something more"); *Universal Secure Registry LLC v. Apple Inc., 10 F.4th 1342, 1354–56 (Fed. Cir. 2021)* (security/authentication context does not by itself supply an inventive concept). As in *Yu v. Apple Inc., 1 F.4th 1040, 1043–45 (Fed. Cir. 2021)*, adding conventional components (here, LLMs and parallel "research") to achieve a desired result does not create eligibility absent a specific improvement in how the computer itself operates.

Practical-Application Clarifier

The July 2024 USPTO guidance emphasizes that to satisfy Step 2A, Prong Two, applicants must show a concrete technological application beyond using a computer "as a tool." Examples include improvements to memory architecture, processor performance, network protocols, training/inference efficiency, or data structures that change how the machine functions. The '371 claims do none of that; they merely state outcomes (e.g., "generate control signals," "perform research in parallel") without the kind of algorithmic specificity that the guidance and cases like Enfish/McRO credit as computer-function improvements.

Preemption Note

Finally, the claim breadth risks undue preemption of the basic practice of checklist-guided dialog with background verification—precisely the concern Alice aims to prevent. See *OIP Techs., Inc. v. Amazon.com, Inc., 788 F.3d 1359, 1362–63 (Fed. Cir. 2015)* (broad functional claims that monopolize the building blocks of human ingenuity are ineligible).

III. Parallel precedent underscores ineligibility

The Federal Circuit has invalidated materially similar claims:

Electric Power Group — monitoring and reacting to power-grid data (abstract).

SAP America v. InvestPic, 898 F.3d 1161 (Fed. Cir. 2018) — statistical analysis of investments (abstract).

PersonalWeb v. Google, 8 F.4th 1310 (Fed. Cir. 2021) — content-based file labeling with hash values (abstract even though it used cryptography).

Stanford — using ML to generate medical metrics (abstract).

If specialized statistical models for the power grid, finance, or genomics failed to supply an inventive concept, a bare statement that "a second LLM researches part of the dialogue" cannot succeed.

IV. Conclusion

Under Alice/Mayo Step 1, every claim in the '371 patent is drawn to an abstract decision-making scheme. Under Step 2, the patent offers only conventional transformer models performing well-known functions, drafted at a results-oriented level of generality. There is no concrete technological advance, no new data structure, no improved computer performance — merely an automated rulebook. The Supreme Court and the Federal Circuit have made clear that such claims are not eligible for patent protection. Accordingly, all twenty claims must be rejected under § 101, even if one were to assume novelty, non-obviousness, and enablement under other sections of the statute.

VIII. Conclusion and Prayer for Relief

Petitioner Open-Source Safety Intelligence, Inc. ("OSSI") respectfully submits that **every one of the twenty claims of U.S. Patent No. 12,142,371 B1 ("the '371 patent") is unpatentable**. As demonstrated in the accompanying Petition, each claim is:

- **Obvious under 35 U.S.C. § 103** in view of a robust body of earlier disclosures that already taught the same checklist-driven dialogue, supervisory control logic, escalation pathways, and secondary-model verification now recited by the '371 patent;
- **Not enabled and inadequately described under 35 U.S.C. § 112(a)** because the specification never teaches the concrete algorithms, data structures, or synchronization protocols necessary to build the claimed multi-model system without undue experimentation;
- **Indefinite under 35 U.S.C. § 112(b)** because pivotal terms—"research into a portion," "control signals [that] contribute in part," "kick-out category," and "trained with over one thousand gradient updates"—lack objective boundaries, leaving competitors unable to discern the scope of the claims; and
- **Directed to patent-ineligible subject matter under 35 U.S.C. § 101** because the claims merely automate the abstract mental steps of script-following, monitoring, and cross-checking using conventional large-language-model components, without any inventive concept that improves computer technology.

1. Prior-Art Landscape Demonstrating Obviousness (§ 103)

The claimed architecture was plainly foreseeable in light of at least thirteen earlier patents—e.g., **US 9,824,188 B2** (conversational virtual-healthcare assistant), **U.S. Patent No. 10,748,644 B2** (Mental health assessment), **U.S. Patent No. 11,348,694 B2** (Medical

assessment based on voice), **U.S. Patent Publication No. 2024/0185968 A1 (**Dialogue-Based Medical Decision System), **US 11,329,933 B1** (AI-supported dialogue), **US 11,694,807 B2** (escalation workflows), **US 10,452,816 B2** (patient-engagement dialogue), **US 10,282,512 B2** (multi-agent safety checks), and OSSI's own governance portfolio (**US 11,532,393 B2** and **US 11,693,990 B1**). Five seminal non-patent references—from Vaswani *et al.* (2017) to McGreevey *et al.* (2020) and the Nexus Flow fall-back study (2024)—confirm that transformer-based chatbots, checklist gating, and layered safety monitoring were routine engineering practice before the '371 filing date. A person of ordinary skill would therefore have found every limitation of the '371 claims **obvious to implement**.

2. Fatal § 112(a) and § 112(b) Defects

The specification never discloses:

- how "control signals" are encoded, prioritized, or injected into an LLM context;
- how the "second-opinion" model synchronizes with the primary model or resolves conflicts; or
- any concrete checklist-enforcement mechanism inside a finite-token transformer window.

Nor does it define what counts as "a portion" of the dialogue, what technical act constitutes "research," or when a "kick-out category" is triggered. Under *Ariad*, *Amgen v. Sanofi*, *Nautilus*, and *Williamson*, such omissions are fatal.

3. Patent-Ineligibility Under § 101

At bottom, the patent seeks to monopolize the abstract idea of **following a script, tracking the conversation, and consulting a reference when unsure**—all long-practiced mental steps.

Merely assigning those steps to generic transformer models, without any novel data structure, training regimen, or hardware optimization, fails Alice step-2. Parallel judiciary precedents (*SAP America*, *IBM v. Zillow*, *Stanford*) compel the same result here.

4. Requested Relief

For these reasons, OSSI respectfully asks the Patent Trial and Appeal Board to:

1. **Institute Post-Grant Review** of the ’371 patent in its entirety;
2. **Cancel Claims 1-20** as unpatentable under one or more of 35 U.S.C. §§ 101, 103, 112(a), and 112(b); and
3. Grant such other and further relief as the Board deems just and proper.

Respectfully submitted,

Alfred Hoyte

/Alfred Hoyte/

Registration Number 33512

08/11/2025

[Counsel of Record]

On behalf of Petitioner OneSource Solutions International, Inc. (OSSI)