

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

BMJ PUBLISHING GROUP LIMITED,

Petitioner,

v.

EVIDIUM, INC.,

Patent Owner.

Case No. PGR2027-00002

U.S. Patent No. 12,518,883

**DECLARATION OF MARK A. MUSEN, M.D., PH.D. IN SUPPORT OF
PETITION FOR POST GRANT REVIEW OF U.S. PATENT NO. 12,518,883**

TABLE OF CONTENTS

	Page
I. INTRODUCTION.....	1
II. QUALIFICATIONS AND EXPERIENCE	3
A. Overview of Qualifications and Experience.....	3
B. Research and Development	4
C. Teaching Experience	8
D. Service to the Scientific Community.....	9
E. Honors and Awards	10
III. PERSON OF ORDINARY SKILL IN THE ART	10
IV. APPLIED LEGAL PRINCIPLES	12
A. Prior Art	12
B. Anticipation	13
C. Obviousness	14
D. Claim Construction Standard.....	15
E. Patent-Eligible Subject Matter.....	16
F. Written Description	17
G. Burden of Proof	18
V. MATERIALS AND PRIOR ART CONSIDERED.....	18
VI. TECHNOLOGY BACKGROUND.....	22
A. Medical Ontologies and Terminologies	23
1. SNOMED-CT.....	23
2. LOINC.....	26
3. Other Ontologies and Terminologies.....	26
B. Medical Journals	27
C. Knowledge Graphs in Healthcare.....	28
D. Clinical Decision Support Systems	32

E.	APIs for Accessing Medical Knowledge Graphs	34
F.	Reasoning Over Hierarchical Medical Terminologies	35
VII.	OVERVIEW OF THE ‘883 PATENT	36
A.	Summary of the ‘883 Patent	36
B.	Prosecution History of the ‘883 Patent	43
VIII.	CLAIM CONSTRUCTION	48
IX.	OVERVIEW OF PRIOR ART REFERENCES	49
A.	Goldberg (Ex. 1006)	49
B.	Gnanasambandam (Ex. 1007)	50
C.	Dang (Ex. 1008)	50
D.	Datla (Ex. 1005)	51
X.	INVALIDITY ANALYSIS OF THE CHALLENGED CLAIMS	51
A.	Ground 1: Claims 1–14 and 17–34 are Obvious Over Goldberg, Gnanasambandam, and Dang	52
1.	Motivation to Combine	52
2.	Independent Claim 1	54
3.	Dependent Claim 2	97
4.	Dependent Claim 3	103
5.	Dependent Claim 4	104
6.	Dependent Claim 5	107
7.	Dependent Claim 6	111
8.	Dependent Claim 7	116
9.	Dependent Claim 8	120
10.	Dependent Claim 9	122
11.	Dependent Claim 10	124
12.	Dependent Claim 11	126
13.	Dependent Claim 12	134
14.	Dependent Claim 13	137

15.	Dependent Claim 14	141
16.	Dependent Claim 17	144
17.	Dependent Claim 18	146
18.	Dependent Claim 19	148
19.	Dependent Claim 20	150
20.	Dependent Claim 21	151
21.	Dependent Claim 22	153
22.	Independent Claim 23	155
23.	Dependent Claim 24	163
24.	Dependent Claim 25	163
25.	Dependent Claim 26	164
26.	Dependent Claim 27	164
27.	Dependent Claim 28	165
28.	Dependent Claim 29	165
29.	Dependent Claim 30	166
30.	Dependent Claim 31	166
31.	Dependent Claim 32	167
32.	Dependent Claim 33	167
33.	Dependent Claim 34	168
B.	Ground 2: Claims 15-16 are Obvious Over Goldberg, Gnanasambandam, Dang, and Datla	168
1.	Motivation to Combine	168
2.	Dependent Claim 15	172
3.	Dependent Claim 16	177
C.	Ground 3: Claims 1-34 are Directed to Ineligible Subject Matter	178
1.	The Challenged Claims are Directed to an Abstract Idea	178
2.	The Challenged Claims Do Not Recite an Inventive Concept	185
XI.	SECONDARY CONSIDERATIONS	191

XII. ADDITIONAL REMARKS	191
-------------------------------	-----

I, Dr. Mark A. Musen, hereby declare as follows:

I. INTRODUCTION

1. I have been retained on behalf of BMJ Publishing Group Limited (“Petitioner”) to provide my expert opinion on the patentability of claims 1–34 (the “Challenged Claims”) of U.S. Patent No. 12,518,883 (the “’883 Patent”). I understand that Petitioner is filing a Petition for Post Grant Review (“PGR”) of the ’883 Patent before the Patent Trial and Appeal Board (“PTAB” or “Board”) of the United States Patent and Trademark Office (“USPTO”).

2. A detailed description of my professional qualifications, including a listing of my specialties, expertise, and professional activities, is contained in my curriculum vitae, a copy of which is provided as Exhibit 1004. A summary of my professional qualifications is set forth below.

3. I am being compensated for my work in this matter at my standard consulting rate of \$900 per hour. I am also being reimbursed for reasonable expenses that I may incur in connection with this matter. My compensation is in no way contingent upon the results of my study, the substance of my opinions or testimony, or the outcome of this proceeding.

4. Based upon the facts and information available to me to date, I detail my opinions relevant to this Petition for Post Grant Review below. I am prepared to

provide expert testimony on the opinions formed herein if asked about those issues by the Board or others.

5. In forming my opinions, I have relied upon my own knowledge, training, and experience in the fields of biomedical informatics, knowledge representation, medical ontologies, and clinical decision support systems. I may additionally discuss my own work, teachings, and knowledge of the state of the art in the relevant time period. I may also rely on handbooks, textbooks, technical literature, and the like to demonstrate the state of the art in the relevant period and the evolution of the technology at issue.

6. Throughout this Declaration, I refer to specific passages, line numbers, or page numbers of the relevant prior art or other technical documents. The citations are intended to be exemplary as opposed to being exhaustive and are not intended to convey that the citations are the only source of evidence to support the propositions for which they are cited.

7. My analysis is ongoing, and I reserve the right to revise, supplement, and amend my opinions based on new information that may become available, including any arguments or evidence that Patent Owner may present in this proceeding.

II. QUALIFICATIONS AND EXPERIENCE

A. Overview of Qualifications and Experience

8. I am currently a Professor of Biomedical Data Science and a Professor of Medicine (Computational Medicine) at Stanford University School of Medicine, where I also serve as Chief of the Division of Computational Medicine. I have held my appointment at Stanford since 1988. In 2022, I was appointed the Stanford Medicine Professor of Biomedical Informatics Research, an endowed chair.

9. My career at Stanford spans more than four decades. From 1995 to 2002, I served as Associate Professor of Medicine (General Internal Medicine, Medical Informatics). From 1989 to 2011, I held a concurrent appointment as Assistant Professor, then Associate Professor, and then Professor of Computer Science (by courtesy). From 1988 to 1995, I was an Assistant Professor of Medicine (General Internal Medicine). Before joining the Stanford faculty, I served as a Clinical Instructor in the Department of Medicine at Stanford University School of Medicine from 1984 to 1987, during which time I was also a Ph.D. candidate in the Medical Information Sciences Program (1983–1987) and a Henry J. Kaiser Family Foundation Fellow in General Internal Medicine at Stanford (1983–1987). I completed my residency in Internal Medicine at Stanford University Hospital (1981–1983) and my internship in Internal Medicine at Stanford University Hospital (1980–1981).

10. I received my Sc.B. in Biology from Brown University in 1977 and my M.D. from Brown University in 1980. I completed my Ph.D. in Medical Information Sciences at Stanford University in 1988. I am a Diplomate of the National Board of Medical Examiners (Certificate No. 227602, 1981) and a Diplomate of the American Board of Internal Medicine (Certificate No. 091133, retired). I hold a California Medical License (No. G044898, inactive).

11. I have spent over four decades conducting research and teaching in the fields of biomedical informatics, knowledge representation, ontology engineering, and clinical decision support. My research has centered on the development of methods and tools for creating and managing biomedical knowledge bases and ontologies—the very technologies that are at the heart of the '883 Patent.

B. Research and Development

12. My research career has focused on biomedical knowledge representation, ontology development, and clinical decision support—all topics directly relevant to the technology claimed in the '883 Patent. Among my major research endeavors:

13. **The Protégé Project (1984–present):** Since 1984, I have led the Protégé Project at Stanford, which has produced one of the world's most widely used platforms for ontology development and knowledge-base construction. Protégé is an open-source tool that enables researchers and practitioners to create, visualize, and

manage ontologies and knowledge graphs. The Protégé platform is used by hundreds of thousands of users worldwide and has been central to the development of biomedical ontologies, including those used for medical terminology standards.

14. The National Center for Biomedical Ontology (2005–2017): I served as Principal Investigator of the National Center for Biomedical Ontology (“NCBO”), an NIH-supported National Center for Biomedical Computing. Through the NCBO, I led the development of BioPortal, which became the world’s largest repository of biomedical ontologies, providing access to hundreds of biomedical ontologies including SNOMED-CT and other medical terminologies referenced in the ’883 Patent. This work is directly relevant to the ’883 Patent’s claims regarding knowledge graphs linked to medical ontologies.

15. EON Project (1988–2006): I led the EON project, which developed computer-based decision support systems for clinical protocols and guidelines. This work involved the creation of knowledge representations and reasoning methods for clinical decision-making—technology closely related to the medical reasoning capabilities claimed in the ’883 Patent.

16. The Center for Expanded Data Annotation and Retrieval (2014–present): I am Principal Investigator of CEDAR, an NIH Big-Data-to-Knowledge Initiative Center of Excellence focused on metadata standards and biomedical data integration.

17. **WHO Collaborating Centre (2014–present):** I serve as Director of the Stanford University Collaborating Centre for Classification, Terminology, and Standards for the World Health Organization. In this role, I have been directly involved in the development and maintenance of international medical classification and terminology systems, including work related to ICD-11, the International Classification of Diseases. I served as Chair of the ICD-11 Health Informatics and Modeling Topic Advisory Group (2008–2016) and as a member of the ICD-11 Revision Steering Committee (2008–2016). This work is directly relevant to the '883 Patent's claims concerning medical ontologies.

18. **BioPortal (2021–present):** I am currently Principal Investigator of the NIH-funded BioPortal project, an expansive knowledge base of biomedical entities and relations that provides tools for searching, browsing, and integrating biomedical ontologies and terminologies. The BioPortal platform hosts and provides API-based access to hundreds of biomedical ontologies, including SNOMED-CT and LOINC—the medical terminology systems specifically referenced in the '883 Patent.

19. **RADx Data Hub (2022–2026):** I served as Principal Investigator of the Rapid Acceleration of Diagnostics (RADx) Data Hub, a national data repository supported by over \$44 million in NIH funding to harmonize and integrate COVID-related diagnostic studies.

20. Over the course of my career, I have secured tens of millions of dollars in research funding from the National Institutes of Health, the National Science Foundation, the World Health Organization, DARPA, and other sponsors for projects directly related to biomedical knowledge representation, ontologies, knowledge graphs, and clinical decision support.

21. I have authored or co-authored more than 155 refereed journal articles, numerous conference papers, and multiple book chapters and books. My publications have been cited extensively, and I have an h-index of 107 as measured by Google Scholar. My published work spans topics including biomedical ontology development, knowledge graph construction and management, clinical decision support, medical terminology standards (including SNOMED-CT), knowledge representation and reasoning, and Semantic Web technologies in biomedicine—all areas directly relevant to the technology of the '883 Patent.

22. An h-index of 107 means that I have authored at least 107 publications that have each been cited at least 107 times by other researchers. This metric is widely used in academia to measure both the productivity and the citation impact of a scholar's body of work. For context, an h-index above 40 is generally considered outstanding in most scientific fields, and an h-index above 60 is exceptional and is typically associated with senior researchers who have made sustained, field-shaping contributions. An h-index of 107 places my scholarly output among the most highly

cited researchers in the fields of biomedical informatics and computer science, reflecting the broad influence and adoption of my work on ontologies, knowledge representation, and clinical decision support over the course of more than four decades of research.

C. Teaching Experience

23. I have taught courses at Stanford University for over 35 years in the areas of biomedical informatics, knowledge representation, ontology engineering, and clinical decision support. I served as Professor of Computer Science (by courtesy) at Stanford from 1989 to 2011, in addition to my appointments in the School of Medicine. I have also served as a member of the Executive Committee of the Biomedical Informatics Graduate Program at Stanford from 1988 to 2026, and I served as Director of Admissions for the Medical Information Sciences Program from 1989 to 1997.

24. I have supervised the doctoral dissertations and postdoctoral research of more than 37 graduate students and postdoctoral researchers in biomedical informatics, knowledge representation, and ontology engineering. I have also served as an external examiner for 13 doctoral dissertations and habilitations at international universities, including Linköping University, Université Paris-Sud, Technical University of Vienna, University of Leipzig, Technical University of Berlin,

University of the Basque Country, Graz University of Technology, University of Barcelona, Amsterdam University Medical Center, and University of Leiden.

D. Service to the Scientific Community

25. I have served the scientific community extensively through editorial, advisory, and leadership roles. I am a Founding Co-editor-in-chief of the journal *Applied Ontology* (2005–2016; Associate Editor 2017–present). I have served on the editorial boards of numerous journals in the fields of artificial intelligence and medical informatics, including *Artificial Intelligence in Medicine*, the *Journal of the American Medical Informatics Association*, the *Journal of Web Semantics*, the *International Journal of Human-Computer Studies*, and *Learning Health Systems*, among others.

26. I have served in leadership roles in major national and international scientific organizations. I was General Chair of the International Semantic Web Conference (2005), General Chair of the International Conference on Knowledge Capture (2011), and Program Chair of the American Medical Informatics Association Annual Symposium (2003). I have served on numerous NIH study sections and advisory boards, including the National Advisory Council on Biomedical Imaging and Bioengineering (2012–2016) and the NIH Big Data to Knowledge Centers Steering Committee (2015–2018).

E. Honors and Awards

27. I have been recognized with numerous honors and awards for my contributions to biomedical informatics and knowledge representation. I was elected a Member of the National Academy of Medicine in 2016—one of the highest honors in the fields of health and medicine. I received the Donald A.B. Lindberg Award for Innovation in Informatics (2006) and a Leadership Award (2007) from the American Medical Informatics Association. I was elected a Founding Fellow of the International Academy of Health Sciences Informatics (2017) and received the Classic Paper Award from the Association for the Advancement of Artificial Intelligence (2018). In 2025, I received the Lifetime Achievement Award from the Knowledge Graph Conference and was named a Distinguished Fellow of the American College of Medical Informatics. I have received honorary doctorates from the University of A Coruña, Spain (2021) and the University of Fribourg, Switzerland (2023). I was elected a Fellow of the American College of Medical Informatics in 1989, a member of the American Society for Clinical Investigation in 1997, and a member of the Association of American Physicians in 2010. I also received a National Science Foundation Young Investigator Award in 1992.

III. PERSON OF ORDINARY SKILL IN THE ART

28. In evaluating the prior art references and other materials, I have used the perspective of a person of ordinary skill in the art (“POSITA”) to which the ’883

Patent is related as of the time of the patent's earliest claimed effective filing date. I am informed by Counsel that a POSITA is a hypothetical person presumed to be aware of all pertinent prior art and the conventional wisdom in the art, and who possesses ordinary creativity—not that of an automaton.

29. I understand that the '883 Patent claims priority to U.S. Provisional Application No. 63/331,526, filed on April 15, 2022. For the purposes of my analysis, I have been asked to assume that April 15, 2022 is the effective filing date of each of the Challenged Claims. I am informed by Counsel that the relevant prior art is assessed as of the effective filing date of the patent.

30. In my opinion, a POSITA in the relevant field of the '883 Patent, as of the April 15, 2022 priority date, would have had at least: (a) a master's degree in biomedical informatics, computer science, health informatics, or a closely related field; and (b) at least two years of experience working with medical knowledge representation systems, biomedical ontologies, knowledge graphs, or clinical decision support systems. Alternatively, a person with a bachelor's degree in one of these fields could qualify as a POSITA with additional years of relevant industry or research experience in these areas. Additional education could compensate for less professional experience, and vice versa.

31. A POSITA as of the relevant date would have had a good working knowledge of: biomedical ontologies and terminologies (including SNOMED-CT

and LOINC); knowledge graph construction, management, and querying; clinical decision support methodologies; knowledge representation and reasoning techniques, including hierarchical and graph-based reasoning; application programming interfaces (APIs) for accessing structured medical knowledge; and the integration of medical evidence from multiple sources into computational systems.

32. I meet and exceed the criteria described above. I am, and have been since well before April 2022, a person with at least ordinary skill in the art pertaining to the '883 Patent. I would have qualified as such a person for many years preceding the earliest claimed priority date of the '883 Patent.

IV. APPLIED LEGAL PRINCIPLES

33. I am not an attorney, and I do not offer legal opinions. However, I have been informed by Counsel for the Petitioner about the following legal standards, which I have applied throughout my analysis.

A. Prior Art

34. I am informed by Counsel that a patent claim is invalid if it is not novel or is obvious in view of the prior art. I understand that under the Leahy-Smith America Invents Act ("AIA"), which applies to the '883 Patent, prior art includes, among other things, patents and printed publications that were available before the effective filing date of the claimed invention.

35. I am informed by Counsel that in this PGR proceeding, the prior art that may be evaluated to demonstrate unpatentability includes patents and printed publications, as well as other categories of prior art available in a PGR proceeding.

36. I understand that a reference qualifies as prior art under 35 U.S.C. § 102(a)(1) if it was patented, described in a printed publication, in public use, on sale, or otherwise available to the public before the effective filing date of the claimed invention. I also understand that a reference qualifies as prior art under 35 U.S.C. § 102(a)(2) if it was described in a patent or patent application that was effectively filed before the effective filing date of the claimed invention.

B. Anticipation

37. I am informed by Counsel that a patent claim is anticipated—and thus unpatentable—under 35 U.S.C. § 102 if each and every element of the claim, as properly construed, is found, either expressly or inherently, in a single prior art reference. The prior art reference need not use the identical words or phrases used in the claim, but the reference must describe each element of the claim with sufficient specificity that a person of ordinary skill in the art would recognize the disclosure of each element.

38. I understand that a feature is considered to be “inherently” disclosed in a prior art reference if the feature is necessarily present in the reference, even if it is not expressly described. I am further informed that the standard for anticipation is

the same as the standard for literal infringement, and that a prior art reference that would literally infringe the claim also anticipates the claim.

C. Obviousness

39. I am informed by Counsel that a patent claim is unpatentable under 35 U.S.C. § 103 if the differences between the claimed invention and the prior art are such that the claimed invention as a whole would have been obvious to a person of ordinary skill in the art at the time the invention was made (or, under the AIA, before the effective filing date of the claimed invention). Obviousness is assessed on a claim-by-claim basis.

40. I understand that the obviousness analysis should not be limited to a rigid application of the “teaching, suggestion, or motivation” test. Rather, the analysis should take into account common sense, the knowledge and creativity of a person of ordinary skill in the art, and the inferences and connections that such a person would draw from the prior art.

41. I understand that a claim may be obvious if, among other things: (a) the combination of known elements yields predictable results; (b) simple substitution of one known element for another produces predictable results; (c) known techniques are used to improve similar devices or methods in the same way; (d) a known technique is applied to a known device or method ready for improvement to yield predictable results; (e) the solution was “obvious to try” because it was chosen from

a finite number of identified, predictable solutions with a reasonable expectation of success; or (f) the work of the art led to variations of the prior art and market forces or design incentives prompted the adaptation.

42. I am informed by Counsel that the obviousness inquiry considers four factors: (1) the scope and content of the prior art; (2) the differences between the prior art and the claims at issue; (3) the level of ordinary skill in the pertinent art; and (4) objective indicia of nonobviousness, sometimes referred to as secondary considerations, such as commercial success, long-felt but unsolved needs, failure of others, copying, unexpected results, praise by others, and licensing.

43. I understand that a prior art reference may “teach away” from a claimed combination if the reference criticizes, discourages, or otherwise discredits the claimed combination. However, mere preference for an alternative embodiment does not constitute teaching away. A reference does not teach away merely because it suggests an alternative approach or expresses a preference for a different solution.

D. Claim Construction Standard

44. I am informed by Counsel that in a PGR proceeding before the PTAB, claim terms are given their ordinary and customary meaning as would be understood by a person of ordinary skill in the art in light of the specification and prosecution history of the patent. I understand that the claims of the patent are read in view of

the specification, although limitations from the specification are not to be imported into the claims.

45. Unless I expressly state otherwise, I have applied the ordinary and customary meaning to each claim term in this Declaration, as understood by a person of ordinary skill in the art in view of the specification and prosecution history of the '883 Patent.

E. Patent-Eligible Subject Matter

46. I am informed by Counsel that 35 U.S.C. § 101 provides that a patent may be obtained for “any new and useful process, machine, manufacture, or composition of matter, or any new and useful improvement thereof.” I understand that the Supreme Court has held that laws of nature, natural phenomena, and abstract ideas are not patentable under Section 101.

47. I understand that there is a two-step framework for determining whether a claim is directed to patent-eligible subject matter. At step one, the court determines whether the claim is directed to a patent-ineligible concept (i.e., an abstract idea, law of nature, or natural phenomenon). If so, at step two, the court examines the elements of the claim—both individually and as an ordered combination—to determine whether the claim contains an “inventive concept” sufficient to transform the nature of the claim into a patent-eligible application.

48. I note that the claims of the '883 Patent were initially rejected by the patent examiner under 35 U.S.C. § 101 and were subsequently allowed only after the applicant amended the claims to add specific limitations relating to API-based querying and outputting results to a software application for generating medical determinations. The relevance of the prosecution history to the scope and validity of the claims is discussed further below.

F. Written Description

49. I understand that 35 U.S.C. § 112 imposes certain requirements on the specification of a patent, including a written description requirement.

50. I understand that to satisfy the written description requirement, the patent specification must convey with reasonable clarity to those skilled in the art that, as of the filing date sought, the patentee was in possession of the claimed invention. I understand that the test for written description requires an objective inquiry into the four corners of the specification from the perspective of a person of ordinary skill in the art.

51. I understand the specification must actually or inherently disclose each and every element of a claim for that claim to meet the written description requirement. I understand that, based on that inquiry, the specification must describe an invention understandable to that skilled artisan and show that the inventor actually

invented the claimed invention. It is insufficient that undisclosed subject matter would have been obvious to a person of ordinary skill in the art.

52. I understand the written description requirement cannot be satisfied if the specification expresses a mere wish or plan for obtaining the claimed invention. I understand that actual possession or reduction to practice outside of the specification is not enough. Rather, I understand that it is the specification itself that must demonstrate possession of the claimed subject matter.

G. Burden of Proof

53. I am informed by Counsel that in this PGR proceeding, the petitioner bears the burden of proving that the challenged claims of the '883 Patent are unpatentable by a preponderance of the evidence.

V. MATERIALS AND PRIOR ART CONSIDERED

54. In forming my opinions, I have considered the '883 Patent and its prosecution history, the prior art references discussed below, and the other materials listed in this section. I have also relied upon my own knowledge, training, and decades of experience in the fields of biomedical informatics, knowledge representation, medical ontologies, knowledge graphs, and clinical decision support systems.

55. A list of the materials I have considered includes, but is not limited to, the following:

- The ‘883 Patent (Ex. 1001);
- The prosecution history of the ‘883 Patent (Ex. 1002);
- U.S. Patent Publication No. 2019/0252074 (“Datla,” Ex. 1005);
- U.S. Patent No. 11,636,949 (“Goldberg,” Ex. 1006);
- U.S. Patent No. 12,248,826 (“Gnanasambandam,” Ex. 1007);
- U.S. Patent Publication No. 2022/0208373 (“Dang,” Ex. 1008);
- U.S. Patent No. 11,194,860 (“Zhong,” Ex. 1009);
- Eryu Xia et al., “Mining Disease-Symptom Relation from Massive Biomedical Literature and Its Application in Severe Disease Diagnosis,” IBM Research (2018) (“Xia,” Ex. 1010);
- U.S. Patent Publication No. 2021/0233658 (Ex. 1011, “Van Assel”);
- U.S. Patent Publication No. 2022/0310260 (Ex. 1020, “Ullaskrishnan”);
- U.S. Patent Publication No. 2019/0026437 (Ex. 1021, “Syeda-Mahmood”);
- U.S. Patent No. 12,512,193 (Ex. 1022, “Marcjan”);
- U.S. Patent Publication No. 2023/0377748 (Ex. 1023, “Yang”);
- U.S. Patent No. 11,488,713 (Ex. 1024, “Velez”);
- U.S. Patent Publication No. 2018/0108443 (Ex. 1025, “Li”);
- U.S. Patent No. 12,014,135 (Ex. 1026, “Potts”);
- U.S. Patent No. 7,260,480 (Ex. 1027, “Brown”);

- Dan McCreary, Knowledge Graphs: “The Third Era of Computing” (Mar. 23, 2019) (Ex. 1028, “McCreary”);
- “DocKG: A Knowledge Graph Framework for Health with Doctor-in-the-Loop,” Ming Sheng et al. (Ex. 1029, “Sheng”);
- “Learning a Health Knowledge Graph from Electronic Medical Records,” Maya Rotmensch et al. (July 20, 2017) (Ex. 1030, “Rotmensch”);
- “Semantic Health Knowledge Graph: Semantic Integration of Heterogeneous Medical Knowledge and Services,” Longxiang Shi et al. (Feb. 12, 2017) (Ex. 1031, “Shi”);
- “Clinical Decision-Support Systems,” Mark A. Musen et al. (2021) (Ex. 1032, “Musen Textbook”);
- “Artificial Intelligence in Medical Diagnosis,” Peter Szolovits et al. (1988) (Ex. 1033, “Szolovits”);
- “*Internist-I*, An Experimental Computer-Based Diagnostic Consultant for General Internal Medicine,” Randolph A. Miller et al. (1982) (Ex. 1034, “Miller”);
- “INTERNIST-I Properties: Representing Common Sense and Good Medical Practice in a Computerized Medical Knowledge Base,” Fred E. Masarie et al., Computers and Biomedical Research (1985) (Ex. 1035, “Masarie”);

- “Towards computer-assisted maintenance of medical knowledge bases,” Dario A. Giuse, *Artificial Intelligence in Medicine* (1990) (Ex. 1036, “Giuse”);
- “Automated Encoding of Clinical Documents Based on Natural Language Processing,” Carol Friedman et al. (2004) (Ex. 1037, “Friedman”);
- SNOMED CT Starter Guide (Version 2014-12-02) (Dec. 2014) (https://web.archive.org/web/20160718140111/http://doc.ihtsdo.org/download/doc_StarterGuide_Current-en-US_INT_20141202.pdf) (Ex. 1038);
- About LOINC, LOINC from Regenstrief (<https://web.archive.org/web/20220125211233/https://loinc.org/about>) (Ex. 1039);
- Mark Musen et al., “The National Center for Biomedical Ontology,” *Journal of the American Medical Informatics Association* (2012) (Ex. 1040);
- Natalya Noy et al., “BioPortal: ontologies and integrated data resources at the click of a mouse,” *Nucleic Acids Research*, 2009, Vol. 37 (May 2009) (Ex. 1041);
- History of The BMJ (<https://www.bmj.com/about-bmj/history-of-the-bmj>) (Ex. 1042);

- About – PubMed

(<https://web.archive.org/web/20200309080055/https://pubmed.ncbi.nlm.nih.gov/about/>) (Ex. 1043);

- BioPortal API Documentation

(<https://web.archive.org/web/20140725213752/http://data.bioontology.org/documentation>) (Ex. 1044).

56. I reserve the right to rely on additional materials in the future, including materials that may become available during the course of this proceeding or in response to any arguments or evidence presented by Patent Owner.

VI. TECHNOLOGY BACKGROUND

57. The '883 Patent, titled “Translation of Medical Evidence Into Computational Evidence and Applications Thereof,” describes transforming unstructured medical evidence, stored in a computer system, into a computational form to create a knowledge graph including medical condition nodes, finding nodes, and investigation nodes. In this section, I provide an overview of the relevant technology that was well known to a POSITA before the earliest claimed priority date of the '883 Patent (i.e., April 15, 2022). All of the technology concepts discussed below were well-established in the art before that date.

A. Medical Ontologies and Terminologies

58. A medical ontology is a structured, formal representation of medical knowledge that defines concepts such as medical conditions, findings, and investigations, and the relationships among them. Medical ontologies have been in use for decades and serve as foundational tools for organizing, integrating, and reasoning over medical information. Long before the priority date of the ‘883 Patent, medical ontologies were widely used in clinical information systems, electronic health records, and clinical decision support applications.

1. SNOMED-CT

59. The Systematized Nomenclature of Medicine—Clinical Terms (“SNOMED-CT”) database is one of the most comprehensive and widely used clinical terminology systems in the world. *See* Ex. 1038; Ex. 1020 at ¶ [0061] (“SNOMED is considered to be the most comprehensive, multilingual clinical healthcare terminology in the world”), ¶ [0019] (“SNOMED ontology is . . . the most widely used and most complete clinical terminology.”). SNOMED-CT was first released in January 2002 and has been continuously maintained and updated since then. It is derived from prior versions of the terminology that date back to the 1960s. SNOMED is used in healthcare systems across more than 40 countries.

60. In or around 2013, the U.S. Government (specifically, the Office of the National Coordinator for Health Information Technology (“ONC”)) required that

SNOMED CT be included in EHR systems for them to be certified for Stage 2 “Meaningful Use.”


61. SNOMED-CT organizes clinical concepts including medical conditions, findings, and clinical observations. *See, e.g.*, Ex. 1006 at 1:65-2:1 (SNOMED “holds sets of relations and qualifiers that support detailed description of medical concepts, as well as names of findings and disorders that can be grouped in families”); Ex. 1020 at ¶ [0061] (“SNOMED comprehensive coverage includes: clinical findings, symptoms, diagnoses....”); Ex. 1022 at 11:7-11 (SNOMED “is used to standardize clinical findings, disorders, body structures, procedures, microorganisms, allergies, and various other clinical domains”).

62. By the time of the ’883 Patent’s priority date, SNOMED-CT contained hundreds of thousands of clinical concepts organized across hierarchies of clinical findings, disorders, procedures, body structures, and other clinical domains. *See, e.g.*, Ex. 1006 at 1:54–56 (SNOMED covers “hundreds of thousands of medical terms arranged as an ontology”), 1:65–2:1; Ex. 1021 at ¶ [0030] (reporting 393,073 concepts in SNOMED CT 2011 standard); Ex. 1020 at ¶ [0061]; Ex. 1023 at ¶ [0074] (reporting 340,000 concepts in SNOMED).

63. SNOMED-CT is organized such that concepts are connected through hierarchical (“is-a”) and other relationships. *See, e.g.*, Ex. 1020 at ¶ [0022]; Ex. 1006 at 9:50–60, FIGS. 6–7. This approach enables traversal of ancestor-descendant

relationships among these concepts. *See, e.g.*, Ex. 1020 at ¶ [0022]; Ex. 1006 at 10:35–46.

64. For example, the figure below shows a SNOMED concept tree showing “is-a” edges including “Bacterial Pneumonia is-a Infective Pneumonia is-a Pneumonia.” Ex. 1006 at FIG. 6.



Ex. 1006 at FIG. 6

65. The ‘883 Patent acknowledges that SNOMED-CT was a known medical terminology ontology. Ex. 1001 at 4:12–18 (“It is realized that multiple clinical ontologies exist, including SNOMED-CT, but that these by themselves do not provide the end-to-end computational representation necessary to make evidence and associated administrative guidelines operational. And although SNOMED-CT provides taxonomies of clinical terms (e.g., fever, heart valve replacements, etc.) it

can only perform limited reasoning.”); *see also id.* at 5:13–19, 7:8–19, 8:40–55, 11:1–10, 12:15–35.

2. LOINC

66. The Logical Observation Identifiers, Names, and Codes (“LOINC”) database is a standard terminology for identifying medical laboratory observations, clinical measurements, and other health-related data. LOINC was first developed in 1994 by the Regenstrief Institute and has been adopted by healthcare organizations worldwide. *See* Ex. 1039.

67. LOINC provides a universal coding system for laboratory tests, clinical observations, and related investigation results. *See, e.g.,* Ex. 1022 at 11:65-12:10 (“LOINC: Logical Observation Identifiers Names and Codes was created specifically to standardize the identification and reporting of medical laboratory observations, including measurements. It has been expanded to standardize clinical observations as well.”). By 2022, LOINC included codes for tens of thousands of laboratory and clinical observations. *See* Ex. 1039.

68. The ‘883 Patent acknowledges that LOINC was a known medical terminology. Ex. 1001 at 5:13–19, 7:11–12, 20:29–31.

3. Other Ontologies and Terminologies

69. In addition to SNOMED-CT and LOINC, numerous other medical ontologies and terminologies have been developed and widely used, including the

International Classification of Diseases (“ICD”), RxNorm (for medications), the Gene Ontology, and the Medical Subject Headings (“MeSH”). *See, e.g.*, Ex. 1022 at 11:50-60, 12:5-15; Ex. 1021 at ¶ [0030]; Ex. 1005 at ¶ [0054]; Ex. 1024 at 4:55–65, 6:5–10.

70. My own work to develop the NCBO’s BioPortal system in the mid-2000s made hundreds of such ontologies publicly available and searchable, serving as a critical resource for the biomedical research community. *See* Ex. 1040; Ex. 1041.

71. By the time of the ‘883 Patent’s priority date, the landscape of biomedical ontologies was vast and well-established.

B. Medical Journals

72. Medical journals are peer-reviewed publications that share, among other things, clinical research findings, medical studies, and healthcare guidelines. Example medical journals include The British Medical Journal (“The BMJ”), the Journal of the American Medical Association (“JAMA”), The Lancet, and The New England Journal of Medicine (“NEJM”). The BMJ is one of the world’s oldest medical journals, having published its first weekly edition in 1840. Ex. 1042.

73. PubMed is a free resource maintained by the National Library of Medicine that provides access to tens of millions of citations and abstracts from

medical journals. PubMed provides links to full text articles when available, e.g., through PubMed Central (PMC). *See* Ex. 1043.


C. Knowledge Graphs in Healthcare

74. A knowledge graph is a structured representation of real-world entities (represented as “nodes”) and the relationships among them (represented as “edges”). Knowledge graphs organize information in a graph structure that enables efficient traversal, querying, and reasoning. *See, e.g.*, Ex. 1025 at ¶ [0003]; Ex. 1026 at 17:30–18:30.

75. In the healthcare domain, knowledge graphs have been used to represent medical conditions, findings (such as symptoms, signs, and test results), investigations (such as medical tests and procedures), procedures, medications, and the relationships among these entities. *See, e.g.*, Ex. 1005 at ¶ [0031]; Ex. 1006 at 6:10–30; Ex. 1011 at ¶ [0115].


76. Medical knowledge graphs typically organize nodes into categories corresponding to different types of medical concepts, including medical condition nodes, finding nodes (which correspond to symptoms, signs, or test results associated with one or more medical conditions), and investigation nodes (which correspond to medical tests or medical procedures associated with one or more

medical conditions). For example, the graph below includes multiple node types that are connected by edges. Ex. 1008 at FIG. 1A.



Ex. 1008 at FIG. 1A

77. As another example, the graph below shows how certain nodes (e.g., symptom nodes) can be linked via edges to multiple other nodes (e.g., a disease node).

100B**FIG. 1B**

Ex. 1008 at FIG. 1B

78. The relationships among these nodes, including ancestor-descendant relationships and associative relationships (e.g., “is-finding-of” or “is-test-for”), enable traversal and reasoning across the graph. *See, e.g.*, Ex. 1006 at 6:40–65, 8:5–30, 10:5–20, 10:25–50, FIGS. 6–7, 9–11; Ex. 1008 at ¶¶ [0027]–[0031], [0047]–[0050]; FIGS. 1A, 1B.

79. The use of knowledge graphs in healthcare was well-established before the ‘883 Patent’s priority date. Before that date, the field was already addressing the

problem of transforming unstructured medical evidence into computational form; medical evidence has been historically represented primarily as prose documents in books and journals, and computational architectures based on text alone are not sufficient to represent, process, or reason controllably or reliably about patient-specific situations. *See, e.g.*, Ex. 1025 at ¶ [0034]; Ex. 1031 at pp. 2-3; Ex. 1030 at pp. 1-2; Ex. 1029 at p. 1; Ex. 1006 at 1:10-30; Ex. 1034 at 468-69; Ex. 1035 at pp. 458-60.

80. The construction of medical knowledge graphs from multiple evidence sources, including the medical literature (including journals like BMJ), clinical guidelines, medical ontologies (e.g., SNOMED and LOINC), electronic health records (“EHR”), and other medical evidence, was a widely practiced approach in the field. There were multiple ways to construct a knowledge graph. First, subject matter experts could manually create the knowledge graphs. Ex. 1006 at 6:5-30, 9:35-50, 10:50-60; Ex. 1034 at p. 469; Ex. 1035 at pp. 458-60.

81. Second, the knowledge graph could be created using techniques like natural language processing (“NLP”). *See, e.g.*, Ex. 1025 at ¶¶ [0032], [0034]; Ex. 1030 at pp. 1, 9; Ex. 1037 at pp. 392-93. For example, McCreary describes using machine learning to harvest relevant information to identify new knowledge for a knowledge graph that “becomes a repository of semantically precise vertices and

relationships with confidence weights retained from the machine learning processes.” Ex. 1028 at 5.




Figure 4: The Knowledge Graph Era: where machine learning continuously reads raw data, combines this with existing knowledge and produces new knowledge, answers and explanations

82. Another approach combined a human-in-the-loop with machine learning to construct knowledge graphs. *See* Ex. 1029 at pp. 1–2, 11; Ex. 1036 at pp. 21-22, 25-26, and 29-31.

83. In sum, by 2022, the construction and use of medical knowledge graphs for integrating clinical evidence, supporting clinical reasoning, and enabling computerized decision support was already well known and widely used.

D. Clinical Decision Support Systems

84. Clinical decision support systems (“CDSS”) are computer-based tools that provide clinicians, patients, and other healthcare stakeholders with person-specific, intelligently filtered, and timely information to enhance health-related decisions. Ex. 1032 at p. 803. CDSS have been in use since the 1970s. *Id.* at p. 806; Ex. 1033 at p. 80. Among the earliest and most influential systems were MYCIN (developed at Stanford), which used rule-based reasoning to diagnose and treat

bacterial infections, and INTERNIST-1, which performed diagnostic decision support in a manner that attempted to mimic the reasoning methods of an expert physician. Ex. 1032 at p. 821; Ex. 1033 at pp. 80, 83; Ex. 1034 at pp. 468-69.

85. INTERNIST-1 was well known for modeling medical conditions and their associated findings using data structures that made explicit many of the knowledge-base distinctions described in the '883 patent, including the separate representation of medical conditions and their findings, the association of individual findings with multiple medical conditions, and the indication that certain medical conditions could occur only in the setting of specific patient characteristics (demographics). *See* Ex. 1034 at pp. 469–70; Ex. 1035 at pp. 458–60, 462, 469–70. The program used its knowledge base to rank order possible medical conditions that might explain a patient's signs, symptoms, and laboratory-test results, and to suggest, through an interactive dialog, medical findings—if their values were to become known—that might allow the system to refine its list of possible medical conditions. *See* Ex. 1034 at pp. 469-70; Ex. 1035 at pp. 458-60, 462, 469-70.

86. By the time of the '883 patent's priority date, CDSS had evolved substantially to incorporate knowledge graphs and ontology-based reasoning. Ex. 1032 at pp. 825–27. Modern CDSS leverage structured medical knowledge bases—including knowledge graphs linked to standard terminologies—to support diagnostic

reasoning, treatment recommendations, and other clinical decisions. *Id.* at pp. 811, 825–27.

E. APIs for Accessing Medical Knowledge Graphs

87. An application programming interface (“API”) is a set of protocols, routines, and tools that allow different software applications to communicate with one another. The use of APIs to access medical knowledge bases and ontologies was well-established before the ‘883 Patent’s priority date. *See, e.g.*, Ex. 1007 at 100:24–34; Ex. 1027 at 3:25–45.

88. For example, the BioPortal system that I developed at Stanford provides RESTful APIs that allow external software applications to search, browse, and query hundreds of biomedical ontologies, including SNOMED-CT and LOINC, and it has done so since it was first deployed in 2006. Ex. 1044. Similarly, the UMLS (Unified Medical Language System) maintained by the National Library of Medicine has long provided API-based access to its terminology services. The FHIR (Fast Healthcare Interoperability Resources) standard, published by HL7 International, includes terminology service APIs for accessing medical vocabularies. *See* Ex. 1032 at pp. 813, 832. By 2022, API-based querying of medical knowledge graphs to return structured results to software applications for generating healthcare determinations was not only well known in the art but also standard practice.

89. In these systems, APIs enable a computer system to receive, via an application programming interface (API), a medical query from a software application for querying said knowledge graph, while providing the interface through which the external software application receives an output in response.

F. Reasoning Over Hierarchical Medical Terminologies

90. Medical terminologies such as SNOMED-CT are organized as hierarchies that support traversal of ancestor-descendant relationships. In the '883 patent's language, the knowledge graph is "configured to support reasoning capability by allowing traversal of ancestor-descendant relationships among one or more of said medical condition nodes, said finding nodes, or said investigation nodes." For example, a query about a specific disease can be generalized to its parent concepts (ancestors) to find broadly relevant findings or specialized to its child concepts (descendants) to find more specific diagnostic information.

91. The use of the '883 patent's "reasoning capability" through "traversal of ancestor-descendant relationships among one or more of said medical condition nodes, said finding nodes, or said investigation nodes" over SNOMED-CT hierarchies was standard practice in biomedical informatics and was well-established long before the '883 patent's priority date. SNOMED-CT's classification structure, for example, explicitly supports such traversal-based reasoning by organizing clinical concepts into a polyhierarchy where each concept

may have multiple parent concepts and multiple child concepts. Reasoning systems have long exploited these relationships to generalize and specialize clinical queries, identify relevant medical concepts, and support clinical decision-making.

VII. OVERVIEW OF THE ‘883 PATENT

A. Summary of the ‘883 Patent

92. The ‘883 patent claims priority to U.S. Provisional Application No. 63/331,526, filed April 15, 2022. It issued from a continuation of Application No. 18/446,954, filed August 9, 2023, which itself is a continuation of the application that issued as U.S. Patent No. 11,869,674 (“the ‘674 Patent”). The ‘883 patent is in the same patent family as U.S. Patent No. 12,518,882 (“the ‘882 Patent”), which was issued on the same date.

93. The ‘883 patent is titled “Translation of Medical Evidence Into Computational Evidence and Applications Thereof.” Ex. 1001 at 1:1–3. The specification describes a platform that extracts clinical concepts from medical evidence sources and organizes them into a graph database of interlinked computational elements — including conditions, diagnostic factors, risk factors, and investigations — that can be queried by software applications. *Id.* at 2:43–62.

94. The specification begins by acknowledging that the problem it purports to address — making medical evidence accessible and actionable — has long been recognized. The specification concedes that “[c]urrent content providers, i.e.

publishers, such as DynaMed, BMJ, UpToDate, etc., provide encyclopedia-style medical content look-up for clinicians covering diagnosis, and management of diseases” and that these providers “assemble the published evidence through expert, human-centric activity—looking at published papers and published society guidelines, sometimes supported by tooling to help human-led assembly activity.” Ex. 1001 at 1:59–65. The specification further acknowledges that “[c]urrent best practice takes what are essentially narrative paper documents, even if they are in a digital form, and boils them down to other paper documents, again in narrative form, that require human beings to interpret their contents.” *Id.* at 1:66–2:5. In other words, the patent recognizes that the underlying process of extracting diagnostic information from medical evidence and organizing it for clinical use was already being performed before the alleged inventions, albeit through human-centric methods.

95. The ‘883 Patent discloses a knowledge graph, which is a database of interlinked nodes representing medical conditions, findings (diagnostic factors), and investigations (laboratory tests and examinations). The specification describes the knowledge graph as a graph database 50 that stores “extracted information from the medical evidence sources after that information has been translated into computational elements.” Ex. 1001 at 8:10–15. Finding nodes are called “elemental diagnostic factors” (EDFs) and investigation nodes are called “elemental diagnostic

investigations” (EDIs). *Id.* at 8:4–9, 9:61–10:10; 11:63–12:3. As the specification explains, “[a] diagnostic factor is a factor that the physician can observe, detect or measure quickly or have the patient describe to him during an office visit (traditional signs and symptoms), while an investigation is a laboratory test or an examination performed by a specialist.” *Id.* at 9:21–26.

96. The specification illustrates linkages between these node types in, for example, FIG. 4A, which shows conditions linked to factor groups containing EDFs and investigation groups containing EDIs. Ex. 1001 at 18:28–19:24.




FIG. 4A

Ex. 1001 at FIG. 4A

97. The specification discloses that the knowledge graph nodes are extracted from existing medical evidence sources and linked to pre-existing ontologies. The specification lists evidence sources from which data are extracted, including: “National Library of Medicine, JAMA Rational Clinical Evidence, BMJ

Best Practice (or other British Medical Journal databases), SNOMED-CT . . . LOINC . . . CPT . . . StatPearls content, and laboratory test reference ranges, drug interactions . . . and other online sources such as Merck Manuals, WebMD, NORD . . . the CDC . . . the WHO, etc.” Ex. 1001 at 7:4–22. Each of these sources was well known and publicly available before the ‘883 Patent’s priority date.

98. The specification describes an NLP engine 74 that extracts clinical concepts from medical evidence sources and patient charts to construct the knowledge graph. Ex. 1001 at 7:56–8:10. The specification acknowledges that the NLP engine uses conventional techniques: “The NLP processes that may be used include but are not limited to: Word Segmenter, Sentence Segmenter, Negation Cues, Negation Scopes, NER (Named Entity Recognition), Relationship Prediction, NEL (Named Entity Linking), Phrase Extractor, Composite Matcher, Embedding Matcher, Graph Matcher, NLI and Long-to-Long Matcher.” *Id.* at 7:56–8:3. Each of these NLP techniques was well known in the art before the ‘883 Patent.




FIG. 2

Ex. 1001 at FIG. 2


99. The specification describes using “a clinical terminology ontology with pre-coordinated and post-coordinated expressions from SNOMED-CT,” using LOINC “that gives us reasoning for laboratory tests,” and using “RxNORM that gives us reasoning for pharmaceuticals.” Ex. 1001 at 5:13–19. As discussed in my Technology Background, SNOMED-CT and LOINC are pre-existing, third-party controlled vocabularies that were widely used in clinical informatics systems for decades before the ‘883 Patent’s priority date. Regarding SNOMED, the specification states that “it is realized that multiple clinical ontologies exist, including SNOMED-CT, but that these by themselves do not provide the end-to-end computational representation necessary to make evidence and associated administrative guidelines operational.” Ex. 1001 at 4:12–16. The specification

further states that “although SNOMED-CT provides taxonomies of clinical terms (e.g., fever, heart valve replacements, etc.) it can only perform limited reasoning.” *Id.* at 4:17–27. These statements confirm that SNOMED-CT was a known, pre-existing ontology.

100. The specification also describes a “reasoning service 72,” which “may use an internal or external taxonomy or any other suitable medical hierarchy in order to perform this inferential reasoning.” Ex. 1001 at 23:61–24:15. The specification provides an example: if a patient has “acute fever,” the system infers that the patient also has the broader concept “fever” and the still broader concept “change in body temperature” by “using the SNOMED ontology polyhierarchy and walking the relevant ancestor-descendant chains.” *Id.* at 8:48–53. The specification further explains that “[t]he reasoning service 72 also contains a component that understands the ancestor-descendant relationships of clinical tests, e.g. CBC with differential is a parent of CBC (complete blood count) is a parent of Hematocrit and White Blood Count.” *Id.* at 8:54–58. This “reasoning capability” is precisely the hierarchical traversal of SNOMED-CT’s polyhierarchy that was standard practice in biomedical informatics before the ‘883 Patent’s priority date, as I discussed in my Technology Background.

101. The specification describes a “DX Glass API” 40 and associated backend components. Ex. 1001 at 5:62–6:12. The “diagnostic glass” application 24

is described as “a real-time, fully-automated, possibly hands-free software application and user interface, designed specifically to be the system of interaction by which a clinician may enter patient data, and explores diagnoses of a patient either during a patient encounter or when the patient is not present.” *Id.* at 6:13–26. The specification describes the APIs used for communication with external systems as conventional, off-the-shelf components: “FHIR (Fast Healthcare Interoperability Resources) standards,” “USCDI (United States Core Data for Interoperability), CDS Hooks, and vendor-specific APIs.” *Id.* at 5:50–54, 6:4–8.



Ex. 1001 at FIG. 1 (combined)

102. The specification also discloses that “In one embodiment, diagnostic glass 24 is an application integrated into the EMR environment via relevant standards such as a SMART (Substitutable Medical Applications, Reusable Technologies) on FHIR application for EHRs (such as Epic and Cerner).” Ex. 1001 at 6:43–50. SMART on FHIR was a well-established standard for launching clinical applications from within EHR systems before the ‘883 Patent.

103. The system output described in the specification includes “the likely, ranked conditions of the patient, along with for each, unknown factors that should be queried and investigations that should be performed.” Ex. 1001 at 9:5–13. As the specification acknowledges, this output is produced by referencing “conditions, factors and investigations in graph database 50.” *Id.*

104. The ‘883 Patent contains 34 claims, including two independent claims: claims 1 and 23. Claim 1 is directed to a method of providing reasoning capability to integrated medical evidence, and claim 23 is directed to a system for providing the same. The dependent claims add limitations relating to, among other things, the structure of investigation and result nodes, the types of findings, the nature of the querying software application (e.g., an AI application or an NLP engine), iterative querying, and the size of the knowledge graph.

B. Prosecution History of the ‘883 Patent

105. I have reviewed the prosecution history of the ‘883 Patent. Ex. 1002. The prosecution history is relevant to understanding the scope and meaning of the claims, as well as the representations made by the applicant to obtain the patent.

106. On February 11, 2025, the applicant filed Patent Application No. 19/050,872 with 30 claims, including two independent claims (claims 1 and 21) and 28 dependent claims.

107. On April 23, 2025, the Examiner issued a non-final Office Action rejecting all 30 pending claims on three grounds: (1) lack of patent-eligible subject matter under 35 U.S.C. § 101; (2) anticipation under 35 U.S.C. § 102(a)(2); and (3) provisional non-statutory double patenting over a co-pending application. Ex. 1002 at 86-87.

108. First, the Examiner rejected claims 1–30 under 35 U.S.C. § 101 as being directed to an abstract idea without significantly more. The Examiner determined that claims 1–30 recited steps of “observation, evaluation, and judgement that can be practically performed by a human, either mentally or with the use of pen and paper.” Ex. 1002 at 88. The Examiner found that, apart from generic computer components such as “an application programming interface (API)” and “a computer system comprising one or more processors and memory,” nothing in the claim elements precluded the steps from being performed in the human mind. *Id.* at 88-89. The Examiner therefore concluded that the claims fell within the “Mental Processes” grouping of abstract ideas. *Id.* at 89. The Examiner further determined that receiving a medical query via an API and outputting results to a software application amounted to insignificant extra-solution activity and data-gathering steps that did not integrate the abstract idea into a practical application. *Id.* at 89-90. The Examiner found these elements to be well-understood, routine, and conventional activities and concluded that the claims were not patent eligible.

109. Second, the Examiner rejected claims 1–30 under 35 U.S.C. § 102(a)(2) as anticipated by Gnanasambandam et al. (U.S. Patent Application Publication No. 2023/0115939) (“Gnanasambandam ‘939”). Ex. 1002 at 95-96. The Examiner mapped the Gnanasambandam ‘939 reference’s knowledge graph structure—including its root nodes, layer nodes, and edge relationships—onto the claimed knowledge graph with medical condition nodes, finding nodes, and investigation nodes. *Id.* at 96-97. The Examiner found that Gnanasambandam ‘939 disclosed each limitation of independent claims 1 and 21, including receiving a medical query via an API, a knowledge graph with medical condition and finding nodes, investigation nodes linked to medical condition nodes, and outputting a response to the software application. *Id.* at 96-103.

110. On September 16, 2025, the applicant filed a response to the Office Action. The response included amendments to the claims, arguments against the rejections, four new claims (claims 31–34), and an information disclosure statement. Ex. 1002 at 48-49.

111. The applicant amended independent claims 1 and 21, as well as several dependent claims. Critically, claims 1 and 21 were amended to recite that nodes in the knowledge graph “are extracted from a plurality of medical evidence sources and are linked to a medical terminology ontology,” and that the ontology “includes a Systemized Nomenclature of Medicine-Clinical Terms (SNOMED-CT) database, or

a Logical Observation Identifiers, Names, and Codes (LOINC) database.” Ex. 1002 at pp. 50-54. The claims were also amended to require that “said knowledge graph is configured to support reasoning capability by allowing traversal of ancestor-descendant relationships among one or more of said medical condition nodes, said finding nodes, or said investigation nodes.” *Id.* at 50, 54. Third, the applicant added a new outputting limitation requiring “outputting, responsive to said medical query, an output to said software application, wherein said output includes data related to one or more of said medical condition nodes, said finding nodes, or said investigation nodes that are relevant to said medical query in order for said software application to generate a medical determination.” *Id.* at pp. 54-55. The applicant also added four new claims 31–34, which required that the knowledge graph include more than twenty thousand nodes (claims 31 and 33) or more than one hundred thousand nodes (claims 32 and 34). *Id.* at 55-56.

112. In response to the § 101 rejection, the applicant argued that the amended claims could not be practically performed by a human and that the API, through which the software application outputs its medical determination, is “a software construct that a human cannot practically perform in his or her mind,” and that “a human cannot implement a software application in his or her mind in order to access a knowledge graph via an API.” Ex. 1002 at 58-59. The applicant further argued that the claims recited a technical solution to a technical problem—integrating medical

evidence from diverse, unstructured sources into a unified knowledge graph that enables reasoning via ancestor-descendant traversal. *Id.* at 57-58.

113. In response to the § 102 rejection, the applicant argued that Gnanasambandam ‘939 did not anticipate the amended claims because it teaches knowledge graphs containing only a single medical condition, whereas the amended claims require a plurality of medical condition nodes with finding and investigation nodes traversable to different medical condition nodes. Ex. 1002 at 59-60. The applicant also argued that Gnanasambandam ‘939 does not disclose the specific medical terminology ontologies required by the claims—SNOMED-CT or LOINC. *Id.* at 60-61. The applicant further argued that Gnanasambandam ‘939 does not disclose traversing ancestor-descendant relationships to support reasoning capability. *Id.* at 61. The applicant also presented arguments for the patentability of specific dependent claims. *Id.* at 61-63.

114. On November 15, 2025, the examiner issued a Notice of Allowance withdrawing the rejections and allowing claims 1–34 over the prior art of record. Ex. 1002 at 16.

115. Regarding the prior § 101 rejection, the examiner agreed that the newly added limitation requiring outputting to the software application for generating a medical determination overcame the rejection, reasoning that the API through which the software application outputs its medical determination “is a software construct

that a human cannot practically perform in his or her mind,” and that “a human cannot implement a software application in his or her mind in order to access a knowledge graph via an API.” Ex. 1002 at 16-17.

116. Regarding the prior § 102 rejection, the examiner found that the prior art of record does not teach or suggest the combination of limitations including receiving a medical query via an API, identifying ancestor-descendant relationships using the reasoning capability, and outputting data to the software application for generating a medical determination. Ex. 1002 at 17.

117. The ’883 Patent issued on January 6, 2026, with 34 claims, including the original claims 1–30 as amended and new claims 31–34. Ex. 1002 at 2-3.

118. As discussed further below, the elements added during prosecution to overcome the Section 101 rejection—particularly the API-based querying and output to a software application for generating medical determinations—were well known in the prior art.

VIII. CLAIM CONSTRUCTION

119. I have been instructed that, unless otherwise noted herein, the Challenged Claims of the ’883 patent are to be given their ordinary and customary meaning as understood by a POSITA at the time of the alleged invention in light of the specification and prosecution history. I have thus applied the plain and ordinary meaning of the claim terms in my analysis below. I reserve the right to supplement

my opinions to the extent that one or both parties later propose claim constructions for any claim terms in the '883 Patent.

IX. OVERVIEW OF PRIOR ART REFERENCES

A. Goldberg (Ex. 1006)

120. Goldberg discloses a hybrid knowledge graph diagnostic system that integrates clinical data with ontology-based knowledge graphs, including SNOMED and LOINC. Goldberg's system includes "a healthcare knowledge graph database for storing a medical hybrid knowledge graph, the hybrid knowledge graph including healthcare data from one or more ontology-based knowledge graphs and clinical data from one on [sic] more sources; a data processing engine running a diagnostic algorithm, designed to process a patient condition data and to process a hybrid knowledge graph database, and select one or more pathways on the knowledge graph to generates [sic] one or more differential diagnoses for suspected disorders." Ex. 1006, Abstract. Goldberg's hybrid knowledge graph includes six distinct node sets: (1) diseases, (2) associated conditions, (3) symptoms, (4) complications, (5) laboratory tests, and (6) treatments. *Id.* at 16:8–24. The system integrates SNOMED's "is-a" hierarchies and LOINC for laboratory test ontologies, and its diagnostic algorithm selects pathways to generate differential diagnoses. *Id.* at 5:51–62, 6:5–28, 7:25–50, 14:38–67, 16:8–24.

B. Gnanasambandam (Ex. 1007)

121. Gnanasambandam discloses an API deployment platform for healthcare knowledge graphs. The system “publish[es] . . . at least one API that uses artificial intelligence to perform analytics based on a knowledge graph pertaining to ontological data of a medical condition.” Ex. 1007 at 1:34–39; claim 1. The platform makes these APIs available to client devices running software applications, which subscribe to and invoke them. *Id.* at 100:24–34, 101:3–24, 102:22–46. Gnanasambandam’s knowledge graphs and logical structures “may be generated and updated continuously or on a periodic basis by an artificial intelligence engine with evidence-based guidelines, physician research, patient notes in EMRs, physician feedback, and so forth.” *Id.* at 8:24–29. The key disclosures of Gnanasambandam include its API architecture for medical knowledge graph access, AI-based analytics, and programmatic access by external software applications.

C. Dang (Ex. 1008)

122. Dang discloses a graph-based clinical diagnosis system employing a three-tier knowledge graph architecture. Dang’s FIG. 1A meta-model knowledge graph expressly includes multiple node types, including “Lab Test & Check” and “Observation Item” as dedicated node types alongside Disease, Symptom, Sign, and other node types. Ex. 1008 at ¶ [0027]; FIG. 1A. FIG. 1B discloses a three-tier operational graph with factual nodes (including “symptom nodes {s1, s2, s3 . . . sm}”;

context nodes $\{c_1 \dots c_p\}$, and test nodes $\{t_1, t_2 \dots t_j\}$ ”), feature nodes $\{f_1, f_2 \dots f_x\}$, and disease nodes $\{d_1, d_2, d_3 \dots d_n\}$. *Id.* at ¶¶ [0027]–[0031]; FIGS. 1A, 1B. The graph includes conditional probabilities $P(f|d)$ on the edges between feature nodes and disease nodes and supports iterative next-inquiry generation for clinical diagnosis. *Id.* at ¶¶ [0028]–[0031]; FIG. 1B.

D. Datla (Ex. 1005)

123. Datla discloses “a natural language processing engine configured to extract the at least one patient symptom and at least one demographic parameter from the received natural language input,” which then queries a knowledge graph to identify “a ranked list of one or more medical conditions, diagnoses, treatments, and/or tests for the patient.” Ex. 1005, claim 1; ¶ [0007].

124. The knowledge graph has “symptoms as root nodes connected by edges to the diseases and medical conditions associated with those symptoms,” and the system generates a diagnosis graph as a subset of the knowledge graph via an activation-decay process that produces “a ranked list of one or more medical conditions, diagnoses, treatments, and/or tests for the patient.” Ex. 1005 at ¶¶ [0006]–[0008], [0028], [0036], [0050]–[0057]; claim 1.

X. INVALIDITY ANALYSIS OF THE CHALLENGED CLAIMS

125. In this section, I provide my opinions regarding why the challenged claims of the ’883 Patent are unpatentable as obvious under 35 U.S.C. § 103 and for

being directed to ineligible subject matter under 35 U.S.C. § 101. In forming these opinions, I have considered the '883 Patent (Ex. 1001), its prosecution history (Ex. 1002), the prior art references identified below, and all other materials listed in Section V. I have also relied on my own knowledge, training, and decades of experience in the fields of biomedical informatics, knowledge representation, medical ontologies, knowledge graphs, and clinical decision support systems.

A. Ground 1: Claims 1–14 and 17–34 are Obvious Over Goldberg, Gnanasambandam, and Dang

126. In my opinion, claims 1–14 and 17–34 of the '883 Patent are obvious over Goldberg, Gnanasambandam, and Dang. As explained below, Goldberg and Gnanasambandam together disclose or render obvious every element of independent claims 1 and 23 and most of the dependent claims. I rely on Dang as additional support for certain elements of claims 1 and 23; to the extent the Board finds that Goldberg does not disclose those elements, Dang supplies them.

1. Motivation to Combine

127. All three references address the same problem: using a medical knowledge graph to support clinical diagnosis. Goldberg provides a hybrid knowledge graph integrating clinical data with SNOMED and LOINC and a diagnostic algorithm that selects pathways to generate differential diagnoses. Ex. 1006, Abstract; 5:51–62; 6:10–34; 7:14–35. Gnanasambandam publishes APIs that “perform analytics based on a knowledge graph pertaining to ontological data of a

medical condition” and provides them to client applications. Ex. 1007 at 100:24–34. Dang generates “a symptom-disease knowledge graph” from medical guidelines and uses it to calculate disease likelihoods and generate inquiries “for a more targeted medical diagnosis.” Ex. 1008 at ¶ [0025]. A POSITA working on any one of these systems would naturally consult the others.

128. A POSITA would have been motivated to deploy Goldberg’s system through Gnanasambandam’s APIs. Goldberg already receives clinical inputs from a separate application and returns diagnostic results. Ex. 1006 at 7:14–35; FIG. 2. Gnanasambandam’s API layer formalizes that input using a well-established healthcare-IT design pattern. Ex. 1007 at 100:24–34; 101:12–37. This requires no change to Goldberg’s knowledge graph.

129. A POSITA would have been motivated to implement the combined system using Dang’s disclosed techniques. For example, Goldberg describes laboratory tests as a distinct node set but does not expressly illustrate a knowledge graph containing these nodes in a figure. Ex. 1006 at 16:60–17:3, 17:31–34; FIGS. 1, 4. Dang expressly illustrates test nodes as a separate node type connected through feature nodes to disease nodes, with condition-specific probabilities on the connections (e.g., $P(f_4|d_2)$). Ex. 1008 at ¶¶ [0028]–[0031]; FIG. 1B. That is a natural way to implement Goldberg’s node sets in a graph database. As another example, Goldberg states that “there are several alternative ways to implement such an

inference engine.” Ex. 1006 at 7:44–47. Dang’s inference method reasons “through hyponymy” between generic and specific nodes using the SNOMED type-of and sub-type relations Goldberg’s graph already contains. Ex. 1008 at ¶¶ [0010], [0047]–[0050]. As another example, Goldberg recommends a “next best step” without specifying how candidates are ranked. Ex. 1006 at 7:25–36. Dang selects the candidate with the maximal “degree of distinguish” and reports higher diagnostic accuracy with fewer questions than benchmark methods. Ex. 1008 at ¶¶ [0058]–[0059].

130. Each reference describes an implemented system using knowledge graphs, with medical concepts as nodes and weighted or typed relationships as edges. Combining them applies known graph-design and inference techniques to data of the same type, so a POSITA would have reasonably expected success.

2. Independent Claim 1

131. Each element of independent claim 1 of the ’883 Patent is disclosed or rendered obvious by the combination of Goldberg and Gnanasambandam. Where noted below, Dang provides additional support for claim 1 to the extent certain elements are not disclosed or rendered obvious by Goldberg. I address each element below.

a. Element [1pre]

132. Element [1pre] recites: “A method of providing reasoning capability to integrated medical evidence from medical evidence sources, said method comprising.”

133. To the extent the preamble is limiting, Goldberg discloses a healthcare platform that provides reasoning capability to integrated medical evidence from medical evidence sources. Goldberg’s system includes “a healthcare knowledge graph database for storing a medical hybrid knowledge graph, the hybrid knowledge graph including healthcare data from one or more ontology-based knowledge graphs and clinical data from one on [sic] more sources; a data processing engine running a diagnostic algorithm, designed to process a patient condition data and to process a hybrid knowledge graph database, and select one or more pathways on the knowledge graph to generates [sic] one or more differential diagnoses for suspected disorders.” Ex. 1006, Abstract. Goldberg explains that the hybrid knowledge graph translates medical knowledge from books, articles, publications, and personal knowledge into a graph integrating clinical knowledge with ontology-based graphs such as SNOMED and LOINC. Ex. 1006 at 5:51–62; 6:20–40.

134. Gnanasambandam discloses a platform for providing reasoning capability to integrated medical evidence. Gnanasambandam’s system uses “artificial intelligence to perform analytics based on a knowledge graph pertaining

to ontological data of a medical condition.” Ex. 1007 at 1:32–42; claim 1. The platform publishes APIs through which software applications access knowledge-graph-based medical analytics. Ex. 1007 at 100:24–34, 101:3–24, 102:22–46.

135. A POSITA would have been motivated to deploy Goldberg’s hybrid knowledge graph diagnostic system using Gnanasambandam’s API platform because both are directed to healthcare knowledge-graph systems for clinical decision support, and exposing clinical reasoning engines through APIs was a well-established design pattern in healthcare IT. Ex. 1007 at 100:24–34, 101:12–25. Goldberg’s system already accepts patient inputs and outputs diagnostic results. Gnanasambandam’s API architecture formalizes this for programmatic access by external software applications. The resulting knowledge-graph-based diagnostic engine accessible via published APIs would have been predictable.

b. Element [1a]

136. Element [1a] recites: “receiving, via an application programming interface (API), a medical query from a software application for querying said knowledge graph.”

137. Goldberg provides knowledge-graph-based clinical reasoning that processes medical queries. Goldberg’s system receives patient condition data, including history, symptoms, signs, and other clinical features, through user interfaces and processes that data against its hybrid knowledge graph using a


diagnostic algorithm that selects pathways to generate differential diagnoses and recommend next steps. Ex. 1006 at 7:14–40, FIG. 2, Abstract. Goldberg defines pathways by selecting connectors that connect pairs of nodes, where the connectors are derived from ontological and clinical relationships between medical terms, and Goldberg’s system outputs diagnostic results based on those pathway selections. Ex. 1006 at 14:35–46. Goldberg thus discloses a knowledge graph that can be queried with patient-specific medical data to produce diagnostic outputs.

138. Goldberg does not expressly disclose that this querying occurs via an application programming interface (API) through which an external software application submits queries and receives output. However, as discussed above, using APIs for accessing medical knowledge graphs was known in the art and widely practiced prior to the ’883 patent. *Supra* § VI.E. In the combination, Gnanasambandam supplies this.

139. Gnanasambandam’s API architecture publishes APIs that use artificial intelligence to perform analytics based on a medical knowledge graph and makes them available to client devices running software applications. Ex. 1007 at 100:24–34, 101:3–24, 102:22–46. Gnanasambandam expressly discloses “publishing, for implementation in the autonomous multipurpose application, at least one API that uses artificial intelligence to perform analytics based on a knowledge graph

pertaining to ontological data of a medical condition, wherein the at least one API is included in a set of APIs.” Ex. 1007 at 1:32–42; claim 1.

140. Gnanasambandam discloses that “[a]pplication back-end component 9104 is configured to publish one or more APIs 9103. As described previously, the one or more APIs 9103 may use artificial intelligence to perform analytics based on a knowledge graph pertaining to ontological data of a medical condition.” Ex. 1007 at 100:24–34; FIG. 91.



Ex. 1007 at FIG. 91 (annotated)

141. The APIs are provided to a client device (a computing device running a software application), which subscribes to and invokes them. Ex. 1007 at 100:28–34, 101:3–24, 102:37–46; FIG. 92, step 9206.




FIG. 92

Ex. 1007 at FIG. 92 (annotated)

142. In the combined system, Gnanasambandam's API layer formalizes the input-output pathway for Goldberg's diagnostic engine, enabling external software applications to submit medical queries and receive diagnostic outputs programmatically. Together, Goldberg and Gnanasambandam disclose the specific flow the claim requires: a software application sends a medical query, the query is run against the knowledge graph, and the output is returned to the application.

Goldberg supplies the application and the separate server-side reasoning engine. Gnanasambandam supplies the API layer between them.

143. Goldberg already separates the querying application from the knowledge-graph engine. For example, Goldberg’s platform 200 “includes a cloud-based inference engine 210, communicatively coupled to a knowledge base 220, and a patient management App. 230.” Ex. 1006 at 7:14–18; FIG. 2. The App takes “a list of given history, symptoms, signs and other clinical features” and “generates a list of suspected diseases and complications.” *Id.* at 7:25–30. Goldberg explains that its modules “may be implemented using one or more front-end and server-side techniques.” *Id.* at 8:9–11. Goldberg also describes its application as “an interface for query-based interaction with the knowledge base.” *Id.* at 8:59–62. Goldberg gives concrete examples of a “data query from the Associative Medical Knowledge Base which provides an answer for questions such as ‘How do CROHN’s and CELLIAC [sic] compare,’” and a query answering “what are all the diseases which could present with a ‘Failure to Thrive.’” *Id.* at 9:1–20; FIGS. 4–5. Goldberg therefore discloses a client-side application that sends medical queries across a network to a separate server-side engine and knowledge graph and receives answers in return. The only thing Goldberg does not do is expressly name the interface between App 230 and engine 210 as an “API.”

144. Gnanasambandam’s APIs are knowledge-graph analytics services that client applications invoke. Gnanasambandam’s “application back-end component 9104 is configured to publish one or more APIs 9103,” and “the one or more APIs 9103 may use artificial intelligence to perform analytics based on a knowledge graph pertaining to ontological data of a medical condition.” Ex. 1007 at 100:24–34; FIG. 91. The back-end component then provides “computer instructions 9112 of the one or more APIs 9103” to client device 9106 as part of front-end component 9110. *Id.* at 100:29–34. Front-end component 9110 may be a web page or “an Internet-enabled application executing on client device 9106.” *Id.* at 100:2–4. The client and the platform communicate by message: the web browser “may interpret a message (e.g., via HTTP messages) received from application back-end component 9104.” *Id.* at 100:50–54. Gnanasambandam’s APIs include a “symptomology API” and a “health history API,” and Gnanasambandam calls this collection of APIs “an interoperability layer of a cognitive intelligence platform in healthcare.” *Id.* at 101:55–67, 102:19–21.

145. A POSITA would have understood that a client application invokes an API “that use[s] artificial intelligence to perform analytics based on a knowledge graph” by sending a request with the input data to be analyzed. The analytics then run against the knowledge graph on the platform side, and the result comes back to

the client. That is what it means for an API to “perform analytics” for an application running elsewhere. In the combined system, the query would proceed as follows:

- i. Goldberg’s Patient Management App 230 (the “software application”) collects a patient’s history, symptoms, signs and lab results, or a comparison question such as those shown in Goldberg's FIGS. 4–5;
- ii. The App sends that input as a request to a published API like Gnanasambandam’s, which “perform[s] analytics based on a knowledge graph.” This is receiving, via an API, a medical query from a software application for querying said knowledge graph;
- iii. Behind the API, Goldberg’s cloud-based inference engine 210 runs its diagnostic algorithm over the hybrid knowledge graph in knowledge base 220; and
- iv. The API returns the result, such as the differential diagnosis list, the next-step recommendation or the feature comparison, to the App.

In this arrangement, Gnanasambandam’s API takes the place of the unnamed network interface between Goldberg’s App 230 and inference engine 210.

146. I discuss motivation to combine above and incorporate that discussion here. *Supra* § X.A.1. A POSITA would have been motivated to add Gnanasambandam’s API architecture to Goldberg’s diagnostic engine because Goldberg already processes patient queries and returns diagnostic results.

Gnanasambandam's API layer formalizes that input and output pathway for programmatic access by external software applications, without changing Goldberg's underlying diagnostic logic. The result would have been predictable because the API provides a routine interface for a known query-response system.

c. Element [1b]

147. Element [1b] recites: "wherein said knowledge graph includes a plurality of nodes, which include a plurality of medical condition nodes, a plurality of finding nodes, and a plurality of investigation nodes."

148. Goldberg's hybrid knowledge graph provides the claimed medical condition nodes, finding nodes, and investigation nodes within a unified knowledge graph. Goldberg describes a process of integrating multiple source data sets, including a healthcare ontology such as SNOMED, an ontology dedicated to laboratory tests and procedures such as LOINC, and clinical probability data derived from medical publications, into "a new hybrid knowledge graph." Ex. 1006 at 6:20–40. The result is a single knowledge graph. Goldberg describes "a database for storage of medical hybrid knowledge graph" (Ex. 1006 at 12:44–54), and the system queries this single hybrid graph to return results spanning all node types. For example, "multiple pathways may be searched for when looking up Pneumonia" and "[t]he system may subsequently show all associations with findings, lab tests, etc." Ex. 1006 at 11:55–67.

149. Goldberg’s claims confirm that the hybrid knowledge graph is a single graph containing all claimed node types. Claim 9 recites “a medical knowledge graph” including “a plurality of nodes” comprising six sets: (1) a first set of nodes representing diseases, (2) a second set of nodes representing associated conditions, (3) a third set of nodes representing symptoms, (4) a fourth set of nodes representing complications, (5) a fifth set of nodes representing laboratory tests, and (6) a sixth set of nodes representing treatments. Ex. 1006 at 15:60–16:21. These nodes are connected by “a plurality of connectors” derived from both ontological and clinical relationships. Ex. 1006 at 16:22–31.

150. Goldberg’s claim 11, which depends indirectly from claim 9 via claim 10, further confirms the unified nature of the graph by reciting that the pathways include “at least one node from the fifth set of nodes” (laboratory tests) alongside nodes from other sets, and the method determines “a specific laboratory test in the plurality of laboratory tests to be administered to the patient” from those pathways. Ex. 1006 at 16:55–17:5. This cross-type pathway traversal (selecting laboratory tests from the same pathways used to traverse disease and symptom nodes) is within a single, unified knowledge graph. Goldberg’s claim 13 independently confirms the same six node set. Ex. 1006 at 17:11–37.

151. FIG. 1 of Goldberg visually confirms that its knowledge graph stores multiple node types within a single, unified knowledge base structure. FIG. 1 is a

table depicting “a sample of knowledge base rules for Celiac disease.” Ex. 1006 at 6:65–7:2.

FIG. 1

Source Disorder	Type of Rule	Rule Target	Probability	Strength of source	Link to source
Celiac disease (disorder)	ASSOCIATED CONDITION	Addison's disease (disorder)	1-2	1	https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3496881/
Celiac disease (disorder)	ASSOCIATED CONDITION	Alopecia areata (disorder)	1-2	2	https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3496881/
Celiac disease (disorder)	ASSOCIATED CONDITION	Diabetes mellitus type 1 (disorder)	1-2	3	https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3496881/
Celiac disease (disorder)	ASSOCIATED CONDITION	Migraine (disorder)	2-5	4	https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3496881/
Celiac disease (disorder)	ASSOCIATED CONDITION	Systemic lupus erythematosus (disorder)	1-2	5	https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3496881/
Celiac disease (disorder)	COMPLICATION	Bile acid malabsorption syndrome (disorder)	2-5	5	https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3496881/
Celiac disease (disorder)	COMPLICATION	Chronic diarrhea (disorder)	10-20	1	https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3496881/
Celiac disease (disorder)	COMPLICATION	Chronic liver disease (disorder)	2-5	2	https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3496881/
Celiac disease (disorder)	EXPOSURE	Gluten (substance)	99-100	3	https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3496881/
Celiac disease (disorder)	GENETIC MUTATIONS	HLA-DQ2	90-100	4	https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3496881/
Celiac disease (disorder)	LABORATORY TESTS	Amylase [Enzymatic activity/volume] in Serum, Plasma or Blood:HIGH	5-10	1	https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3496881/

Ex. 1006 at FIG. 1

152. The table is organized with columns for “Source Disorder,” “Type of Rule,” “Rule Target,” “Probability,” “Strength of source,” and “Link to source.” *Id.* at FIG. 1. All entries in FIG. 1 share a single Source Disorder (“Celiac disease (disorder)”) and the “Type of Rule” column categorizes each rule target by its clinical role within the graph. The rule types shown include ASSOCIATED CONDITION (e.g., Addison’s disease, Diabetes mellitus type 1), COMPLICATION

(e.g., Chronic diarrhea, Chronic liver disease), EXPOSURE (Gluten), GENETIC MUTATIONS (HLA-DQ2), and LABORATORY TESTS (Amylase in Serum). *Id.* at FIG. 1.

153. The extended version of this same knowledge base structure is further illustrated in FIG. 4, which provides a feature-by-feature comparison for Celiac disease and Crohn's disease.

FIG. 4

Type of Feature	Feature	Likelihood with Celiac disease	Likelihood with Crohn's disease
SYMPTOM/COMPLAINT	Abdominal pain (finding)	5-10	
SYMPTOM/COMPLAINT	Abdominal pain (finding)# Chronic (qualifier value)	90-100	
COMPLICATION	Acute hepatitis (disorder)	2-5	
COMPLICATION	Acute intestinal obstruction (disorder)	5-10	
COMPLICATION	Acute upper gastrointestinal hemorrhage (disorder)	2-5	
ASSOCIATED CONDITION	Addison's disease (disorder)	1-2	
LABORATORY TESTS	Amylase [Enzymatic activity/volume] in Serum, Plasma or Blood:HIGH	5-10	
COMPLICATION	Anemia of chronic disorder (disorder)		5-10
SIGNS	Ataxia (finding)	1-2	
ASSOCIATED CONDITION	Autoimmune hepatitis (disorder)	1-2	5-10
LABORATORY TESTS	Ceruloplasmin [Mass/volume] in Serum or Plasma:LOW	2-5	
COMPLICATION	Chronic diarrhea (disorder)	10-20	5-10
COMPLICATION	Chronic liver disease (disorder)	2-5	2-5
COMPLICATION	Constipation (disorder)	5-10	
COMPLICATION	Delayed puberty (disorder)	2-5	5-10
COMPLICATION	Depressive disorder (disorder)	2-5	
ASSOCIATED CONDITION	Diabetes mellitus type 1 (disorder)	1-2	
LABORATORY TESTS	Endomysium Ab [Units/volume] in Serum:HIGH	90-95	
COMPLICATION	Exacerbation of Crohn's disease of large intestine (disorder)		90-100
COMPLICATION	Exocrine pancreatic insufficiency (disorder)		2-5
SYMPTOM/COMPLAINT	Failure to gain weight (finding)	5-10	
COMPLICATION	Failure to thrive (disorder)	60-70	30-40
COMPLICATION	Iron deficiency anemia (disorder)	5-10	5-10
ASSOCIATED CONDITION	Migraine (disorder)	2-5	2-5
COMPLICATION	Osteopenia (disorder)	2-5	5-10
RISK FACTORS	Periods of life (qualifier value):ADULTHOOD	30-40	40-50
COMPLICATION	Vitamin D deficiency (disorder)	2-5	10-20


Ex. 1006 at FIG. 4

154. FIG. 4 adds the additional rule types SYMPTOM/COMPLAINT (e.g., “Abdominal pain (finding)”), SIGNS (e.g., “Ataxia (finding)”), and RISK

FACTORS (e.g., “Periods of life: ADULTHOOD”) alongside the ASSOCIATED CONDITION, COMPLICATION, and LABORATORY TESTS categories already shown in FIG. 1. *Id.*, FIG. 4. A POSITA would understand from FIGS. 1 and 4 together that Goldberg’s knowledge graph contains, at minimum, disease nodes, associated condition nodes, complication nodes, symptom/complaint nodes, sign nodes, laboratory test nodes, and risk factor nodes which map to and exceed the claimed plurality of “medical condition nodes,” “finding nodes,” and “investigation nodes.”

155. FIG. 9 of Goldberg provides a graphical depiction of a portion of the hybrid knowledge graph, illustrating how medical condition nodes and finding nodes are connected through both ontological and clinical relationships. Ex. 1006, FIG. 9, 5:5–7, 10:4–12.


FIG. 9



Ex. 1006 at FIG. 9

156. FIG. 9 shows “an example of a hybrid graph showing clinical relations connecting Pneumonia and Fever.” Ex. 1006 at 10:3–5. The graph contains seven nodes: Pneumonia, Septic shock, Systemic inflammatory response syndrome, Sepsis, Pneumonitis, Legionella pneumonia, and Fever. *Id.*, FIG. 9. Goldberg explains that in this figure, “darker arrows represent SNOMED 'type-of' relations” while “[l]ighter arrows reflect clinical relations taken from medical literature,” with single-direction arrows representing causal relationships and double-sided arrows representing clinical associations. *Id.* at 10:5–12.

157. FIG. 10 of Goldberg provides a detailed illustration of how its hybrid knowledge graph enables traversal across multiple node types through linked pathways.




158. FIG. 10 shows “an example of a hybrid graph showing connections between Myocardial infarction and abdominal pain.” Ex. 1006 at 10:13–23. The graph contains approximately nineteen nodes, including disease nodes (e.g., Myocardial infarction, Ischemic heart disease), complication/condition nodes (e.g., Hepatic Congestions, Pulmonary hypertension, Left ventricular hypertrophy), and finding nodes (e.g., Abdominal pain RUQ, Hepatomegaly, Cardiomegaly, Left ventricular hypertrophy, Hypokalemia). *Id.*, FIG. 10. The figure includes an express legend identifying three types of connectors: solid arrows for “Cause of”

relationships, dashed arrows for “Type of” (ontological) relationships, and dotted arrows for “Associated with” (clinical association) relationships. *Id.*

159. FIGS. 1, 9, and 10 in Goldberg confirm that its hybrid knowledge graph (1) stores multiple distinct node types, including a plurality of medical condition nodes, finding nodes, and investigation nodes, within a single unified structure (FIG. 1); (2) connects those nodes through both ontological relationships derived from SNOMED and clinical relationships derived from medical literature (FIG. 9); and (3) enables traversal across linked pathways that cross node types and relationship types to support clinical reasoning from a disease to its manifestations (FIG. 10).

160. To the extent Goldberg is found not to disclose element [1b], Dang provides additional structural support for this element, particularly with respect to investigation nodes. Dang’s FIG. 1A is a block diagram of a symptom-disease meta-model knowledge graph that expressly includes “Lab Test & Check” as a dedicated node type alongside Disease, Symptom, Sign, Observation Item, and other node types. Dang (Ex. 1008) at ¶ [0027]; FIG. 1A. Dang’s FIG. 1A is a graphical depiction of the knowledge graph topology that shows investigation nodes (“Lab Test &

Check”) as structurally distinct node types within the graph, with express directional edges connecting them to other node types.



Dang at FIG. 1A (annotated)

161. Dang’s FIG. 1B provides further architectural detail. FIG. 1B depicts a three-tier operational knowledge graph in which “the factual nodes include symptom nodes $\{s_1, s_2, s_3 \dots s_m\}$; context nodes $\{c_1 \dots c_p\}$, and test nodes $\{t_1, t_2 \dots t_j\}$.” Dang (Ex. 1008) at ¶ [0029]. The test nodes are structurally present in the same unified graph as the feature nodes $\{f_1, f_2 \dots f_x\}$ and disease nodes $\{d_1, d_2, d_3 \dots d_n\}$, connected by probabilistic edges. Dang (Ex. 1008) at ¶¶ [0028]–[0031]; FIG.

1B. This three-tier graph expressly depicts investigation nodes (test nodes) as structurally distinct nodes within the knowledge graph.

100B




FIG. 1B

Dang at FIG. 1B (annotated)

162. A POSITA would have been motivated to incorporate investigation nodes as structurally distinct nodes within a knowledge graph, as expressly shown in Dang, because this is consistent with Goldberg’s description. Goldberg discloses “a fifth set of nodes representing the plurality of laboratory tests” as one of six node sets within a unified knowledge graph. Ex. 1006 at 16:8–24. Dang’s FIG. 1A and

FIG. 1B expressly illustrate the structural implementation of this concept, depicting test nodes as a dedicated node type within the graph topology, with defined edges connecting them through intermediate feature nodes to disease nodes. Ex. 1008 at ¶¶ [0027]–[0029]; FIGS. 1A, 1B. A POSITA designing a clinical knowledge graph with the node categories disclosed in Goldberg would have recognized the benefit of implementing those nodes using Dang’s topology because this enables more reliable automated traversal and reduces ambiguity in how the diagnostic algorithm identifies and processes investigation nodes. The result would have been predictable because both references use knowledge graphs for diagnostic reasoning, and Dang’s graph topology merely provides an express illustration of the structural implementation of the node categories that Goldberg already recites.


d. Element [1c]

163. Element [1c] recites: “wherein at least one of said finding nodes in said knowledge graph is traversable to a first medical condition node to indicate that a finding represented by said at least one of said finding nodes is associated with a first medical condition represented by said first medical condition node.”

164. Goldberg’s knowledge graph includes express traversal paths from finding nodes (symptoms) to condition nodes (diseases). Goldberg discloses that the application connects clinical manifestations to originating diseases through paths that include problem lists and complications. For example, the system tracks

jaundice back to a complication and then to the originating liver disorder. Ex. 1006 at 8:32–40.

165. FIG. 9 of Goldberg shows a hybrid graph with clinical relations connecting Pneumonia (a medical condition) and Fever (a symptom).



Ex. 1006 at FIG. 9 (annotated)

166. FIG. 10 of Goldberg shows connections between Myocardial infarction and abdominal pain through multiple paths that include type-of and causal relationships. Ex. 1006 at 10:13–23.

167. Goldberg’s claim 1 recites defining pathways by selecting connectors that connect pairs of nodes, with connectors derived from ontological and clinical relationships between medical terms. Ex. 1006 at 14:35–46.

168. To the extent Goldberg does not expressly disclose the required traversability, Dang clearly illustrates it. For example, in FIG. 1A of Dang, the

“Symptom” node (finding node) is traversable to the “Disease” node (medical condition node). In FIG. 1B of Dang, the symptom nodes are traversable to disease nodes (e.g., node s_1 is traversable to d_1 ; node s_2 is traversable to d_1 and d_2 , etc.). I discussed the reasons for combining Goldberg and Dang (to the extent necessary) above and incorporate that analysis here. *Supra* § X.A.1.

e. Element [1d]

169. Element [1d] recites: “wherein at least one of said investigation nodes in said knowledge graph is traversable to a second medical condition node to indicate that an investigation represented by said at least one of said investigation nodes is associated with a second medical condition represented by said second medical condition node.”

170. This element requires that at least one investigation node in the knowledge graph is “traversable to” a condition node, indicating an association between the investigation and the condition. Under its plain and ordinary meaning, “traversable to” means that a connected path exists in the knowledge graph from the investigation node to the condition node, including through intermediate nodes. Goldberg discloses this element.

171. For example, Goldberg’s claim 9 expressly recites “a fifth set of nodes representing the plurality of laboratory tests” as part of the medical knowledge graph. Ex. 1006 at 16:8–24. Goldberg’s claim 11, which depends from claim 9,

recites that “the one or more pathways comprises at least one node from the fifth set of nodes [laboratory tests] and at least one node from the sixth set of nodes [treatments],” and the method further comprises “determining, by the at least one processor and from the one or more pathways, a specific laboratory test in the plurality of laboratory tests to be administered to the patient.” Ex. 1006 at 16:55–17:5. Goldberg’s claims 13–14 contain a parallel recitation. Ex. 1006 at 17:10–18:38 (reciting “a fifth set of nodes representing the plurality of laboratory tests” and determining a specific laboratory test to be administered to the patient). This expressly describes traversable paths between investigation nodes (laboratory tests) and condition nodes (diseases) through the pathway mechanism. Goldberg’s system determines which laboratory test to administer based on pathways that connect disease nodes to laboratory-test nodes.

172. As another example, Goldberg discloses that “multiple pathways may be searched for when looking up Pneumonia” and that “[t]he system may subsequently show all associations with findings, lab tests, etc.” Ex. 1006 at 11:55–67. This is express runtime retrieval of investigation nodes (lab tests) in response to a disease query (i.e., traversal from condition nodes to investigation nodes and vice versa).

173. Goldberg’s FIG. 1 further illustrates investigation-to-condition traversability in the unified knowledge base. FIG. 1 shows ‘LABORATORY

TESTS' as a rule type linking Celiac disease (a medical condition node) to specific laboratory tests such as 'Amylase [Enzymatic activity/volume] in Serum, Plasma or Blood: HIGH,' with a probability value of 5–10 on the connection. Ex. 1006, FIG. 1; 6:65–7:2.

FIG. 1

Source Disorder	Type of Rule	Rule Target	Probability	Strength of source	Link to source
Celiac disease (disorder)	ASSOCIATED CONDITION	Addison's disease (disorder)	1-2	1	https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3496881/
Celiac disease (disorder)	ASSOCIATED CONDITION	Alopecia areata (disorder)	1-2	2	https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3496881/
Celiac disease (disorder)	ASSOCIATED CONDITION	Diabetes mellitus type 1 (disorder)	1-2	3	https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3496881/
Celiac disease (disorder)	ASSOCIATED CONDITION	Migraine (disorder)	2-5	4	https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3496881/
Celiac disease (disorder)	ASSOCIATED CONDITION	Systemic lupus erythematosus (disorder)	1-2	5	https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3496881/
Celiac disease (disorder)	COMPLICATION	Bile acid malabsorption syndrome (disorder)	2-5	5	https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3496881/
Celiac disease (disorder)	COMPLICATION	Chronic diarrhea (disorder)	10-20	1	https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3496881/
Celiac disease (disorder)	COMPLICATION	Chronic liver disease (disorder)	2-5	2	https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3496881/
Celiac disease (disorder)	EXPOSURE	Gluten (substance)	99-100	3	https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3496881/
Celiac disease (disorder)	GENETIC MUTATIONS	HLA-DQ2	90-100	4	https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3496881/
Celiac disease (disorder)	LABORATORY TESTS	Amylase [Enzymatic activity/volume] in Serum, Plasma or Blood:HIGH	5-10	1	https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3496881/

Ex. 1006 at FIG. 1

174. FIG. 4 shows a feature-by-feature comparison of Celiac disease and Crohn's disease drawn from the same knowledge base. It lists LABORATORY TESTS entries (e.g., Amylase: HIGH; Ceruloplasmin: LOW; Endomysium Ab: HIGH) with likelihood values alongside SYMPTOM/COMPLAINT, SIGNS,

COMPLICATION, and ASSOCIATED CONDITION entries in a single table. Ex. 1006, FIG. 4; 9:1–10. These entries show that investigation nodes sit in the same knowledge base as condition nodes and are linked to them by typed rules with probability values.

FIG. 4


Type of Feature	Feature	Likelihood with Celiac disease	Likelihood with Crohn's disease
SYMPTOM/COMPLAINT	Abdominal pain (finding)	5-10	
SYMPTOM/COMPLAINT	Abdominal pain (finding)# Chronic (qualifier value)	90-100	
COMPLICATION	Acute hepatitis (disorder)	2-5	
COMPLICATION	Acute intestinal obstruction (disorder)	5-10	
COMPLICATION	Acute upper gastrointestinal hemorrhage (disorder)	2-5	
ASSOCIATED CONDITION	Addison's disease (disorder)	1-2	
LABORATORY TESTS	Amylase [Enzymatic activity/volume] in Serum, Plasma or Blood:HIGH	5-10	
COMPLICATION	Anemia of chronic disorder (disorder)		5-10
SIGNS	Ataxia (finding)	1-2	
ASSOCIATED CONDITION	Autoimmune hepatitis (disorder)	1-2	5-10
LABORATORY TESTS	Ceruloplasmin [Mass/volume] in Serum or Plasma:LOW	2-5	
COMPLICATION	Chronic diarrhea (disorder)	10-20	5-10
COMPLICATION	Chronic liver disease (disorder)	2-5	2-5
COMPLICATION	Constipation (disorder)	5-10	
COMPLICATION	Delayed puberty (disorder)	2-5	5-10
COMPLICATION	Depressive disorder (disorder)	2-5	
ASSOCIATED CONDITION	Diabetes mellitus type 1 (disorder)	1-2	
LABORATORY TESTS	Endomysium Ab [Units/volume] in Serum:HIGH	90-95	
COMPLICATION	Exacerbation of Crohn's disease of large intestine (disorder)		90-100
COMPLICATION	Exocrine pancreatic insufficiency (disorder)		2-5
SYMPTOM/COMPLAINT	Failure to gain weight (finding)	5-10	
COMPLICATION	Failure to thrive (disorder)	60-70	30-40
COMPLICATION	Iron deficiency anemia (disorder)	5-10	5-10
ASSOCIATED CONDITION	Migraine (disorder)	2-5	2-5
COMPLICATION	Osteopenia (disorder)	2-5	5-10
RISK FACTORS	Periods of life (qualifier value):ADULTHOOD	30-40	40-50
COMPLICATION	Vitamin D deficiency (disorder)	2-5	10-20

Ex. 1006 at FIG. 4

175. Because Goldberg's unified hybrid knowledge graph includes a plurality of diseases (e.g., as confirmed by claim 9) the finding node that is


traversable to one condition (e.g., Fever traversable to Pneumonia, per FIG. 9 and 10:5–20) and the investigation node that is traversable to another condition (e.g., Amylase traversable to Celiac disease, per FIG. 1 and 6:65–7:2) satisfy the claim requirement that the “second medical condition node” be different from the “first medical condition node.” The unified graph architecture ensures these are nodes in a single hybrid knowledge graph.

176. To the extent element [1d] is not disclosed by Goldberg, this element is supplied by Dang. For example, in FIG. 1A of Dang, an investigation node (Lab Test & Check) is traversable to a medical condition node (Disease) to indicate that the investigation represented by the investigation node is associated with the medical condition represented by the medical condition node.

100A**FIG. 1A**

Dang at FIG. 1A (annotated)

177. As another example, Dang's FIG. 1B shows how multiple investigation nodes (shown as test nodes $t_1, t_2 \dots t_j$) are traversable to multiple medical condition nodes ($d_1, d_2, d_3 \dots d_n$) in the knowledge graph.

100B

Dang proceeds from step 502 to infer evidential nodes (step 504), disease feature nodes (step 506), and possible disease nodes (step 508), it traverses from factual nodes, including test nodes, through intermediate nodes to disease nodes. *Id.* at ¶¶ [0047]–[0050]; FIG. 5. In FIG. 1B, for example, test node t_2 is connected to feature nodes f_k , which in turn are connected to disease node d_n (via $P(f_k|d_n)$). Ex. 1008, FIG. 1B (annotated); ¶ [0031]. Although Dang’s text describes the connections of symptom nodes to feature nodes (*id.* at ¶ [0030]), a POSITA viewing FIG. 1B would understand that test nodes, which Dang identifies as factual nodes (*id.* at ¶¶ [0028]–[0029]), are connected to feature nodes.

179. I provided reasons above why a POSITA would have been motivated to implement Dang’s architecture in Goldberg and incorporate that analysis here. *Supra* § X.A.1.

f. Element [1e]

180. Element [1e] recites: “wherein one or more of said nodes in said knowledge graph are extracted from a plurality of medical evidence sources and are linked to a medical terminology ontology.”

181. As discussed above, extracting information from medical evidence sources to construct nodes in a knowledge graph was well known in the art and routine prior to the ’883 patent. *Supra* § VI.C. Similarly, medical terminology

ontologies, such as SNOMED and LOINC, were well known in the art prior to the '883 patent. *Supra* § VI.A.

182. Goldberg discloses that one or more of the nodes in the knowledge graph are extracted from a plurality of medical evidence sources. For example, Goldberg “enable[s] the translation of medical knowledge from multiple sources (e.g., books, articles, publications, personal knowledge etc.), into a hybrid knowledge graph which integrates clinical knowledge along with other ontology based knowledge graphs, such as SNOMED, LOINC or other types of medical databases (e.g., including costs and risks of lab and imaging tests, medicine’s side-effects, disease prevalence and incidence rates, etc.).” Ex. 1006 at 5:51–62.

183. Goldberg also describes that “[d]ata is derived from evidence-based or other sources” and that “[e]ach such association within the knowledge base has a source such as textbook, article etc.” with associated rankings. Ex. 1006 at 6:36–44.

184. Goldberg explains that its hybrid knowledge graph can integrate other knowledge graphs constructed from different medical evidence sources. For example, “a first knowledge graph is derived from a healthcare Ontology which includes a large vocabulary of medical concepts. (e.g., SNOMED).” Ex. 1006 at 6:10-15. As another example, a “second knowledge graph is derived from a different source such as an ontology dedicated to laboratory tests and procedures (e.g., LOINC).” *Id.* at 6:17-20. As a further example, a “third knowledge graph is created

or generated by medical experts or healthcare professionals, who research medical publications to build a clinical probability graph that describes clinical statistical information that links medical concepts that are described in the Ontological graphs.” *Id.* at 6:20-30; *see also id.* at 6:5-28.

185. Goldberg further discloses that “[d]ata is derived from evidence-based or other sources” and “[e]ach such association within the knowledge base has a source such as textbook, article etc.” Ex. 1006 at 6:35-44. For example, FIGS. 8A and 8B show publication sources including Harrison’s Principles of Internal Medicine and journal citations. Ex. 1006 at 5:51–62; 6:20–38; 9:61–10:3, FIGS. 8A–8B.

FIG. 8A

EPIDEMIOLOGY		
Given	Distribution	Breakdown By Gender
Ischemic heart disease (disorder)	Population	Male (qualifier.value)
		65%
		Female (qualifier.value)
		35%
Publication Name		
Harrison's Principles of Internal Medicine, 20e		
Chapter		
267: Ischemic Heart Disease		
Location (URL)		
https://accessmedicine.mhmedical.com/content.aspx?sectionid=192029547&bookid=21294&resultclick=2#1160046746		
Given	Incidence of	Myocardial infarction (disorder)
Diabetes mellitus (disorder)		45.5/10000/yr
Location (URL)		
http://www.ncbi.nlm.nih.gov/pmc/articles/PMC2473666/abstract		
Given	Risk of	Myocardial infarction (disorder)
Atrial fibrillation (disorder)		1.7x lift
Publication Name		
JAMA Intern Med 2014 Jan 1;174(1):107		
Location (URL)		
https://www.ncbi.nlm.nih.gov/pubmed/24190640?from=Abstract		

Ex. 1006 at FIG. 8A (annotated)

186. Goldberg further discloses that “at least some of the entering of healthcare publications content is integrated into the associative medical database by automated tools, optionally in cooperation with human editors.” Ex. 1006 at 10:60–63.

187. These disclosures confirm that one or more nodes in Goldberg’s hybrid knowledge graphs are linked to a medical terminology ontology. Goldberg expressly integrates SNOMED and LOINC. Ex. 1006 at 1:54–56; 6:5–28. SNOMED is a well-known reference ontology covering hundreds of thousands of medical terms. *Id.* at 1:54–56. Goldberg’s hybrid knowledge graph includes a knowledge graph derived from a healthcare ontology such as SNOMED and a knowledge graph derived from an ontology dedicated to laboratory tests and procedures such as LOINC. *Id.* at 6:5–28. Thus, one or more nodes are linked to a medical terminology ontology.


g. Element [1f]

188. Element [1f] recites: “wherein said knowledge graph is configured to support reasoning capability by allowing traversal of ancestor–descendant relationships among one or more of said medical condition nodes, said finding nodes, or said investigation nodes.”

189. Goldberg’s knowledge graph is configured to support reasoning capability by allowing traversal of ancestor-descendant relationships. Goldberg

integrates SNOMED, whose foundational organization uses an “is-a” hierarchy. FIG. 6 shows express “is-a” edges, including Bacterial Pneumonia is-a Infective Pneumonia is-a Pneumonia.


FIG. 6



Ex. 1006 at FIG. 6

190. FIG. 7 shows “is-a” edges among disease concepts, including Myocardial infarction, Ischemic heart disease, and Myocardial disease. Ex. 1006 at FIGS. 6–7; 9:50–60.

FIG. 7



Ex. 1006 at FIG. 7

191. Goldberg’s model expressly includes “Is a” or “Sub-type” as a relation type. The specification confirms that the model may consider subtype relations in the ontology, including subtypes of diseases such as iron-deficiency anemia being a subtype of anemia. Ex. 1006 at 10:35–50, FIG. 11. Goldberg’s diagnostic algorithm uses these ontology and clinical relationships to select pathways and to generate differential diagnoses. The SNOMED hierarchy therefore supplies the ancestor–descendant structure through which the hybrid graph supports clinical reasoning. Goldberg’s SNOMED concept-tree examples illustrate the ontology’s ancestor–

descendant organization, including the Pneumonia hierarchy and related disease concepts. Ex. 1006, FIGS. 6–7.

192. A POSITA would have understood SNOMED’s “is-a” hierarchy, as integrated in Goldberg’s knowledge graph, to be the ancestor–descendant relationships recited in element [1f]. While I understand that Xia is not part of Ground 1, it corroborates a POSITA’s understanding of how “is_a” relationships were used in medical knowledge graphs. Xia built a biomedical knowledge graph from concepts in the Human Disease Ontology and the Symptom Ontology. Ex. 1010 at 1118-23. Xia explains that “[r]elationships are built between entities with the same attribute to reflect the hierarchy information. The ‘is_a’ relationship is a reflection that the former concept is a child of the latter.” *Id.* at 1119. Xia also describes that “[e]ach mentioned concept was then traced all the way up to all its ancestors, intending to reflect that mentioning of a concept is also mentioning of its ancestors.” *Id.* Xia thus treated “is_a” links as parent–child links and followed them upward to a concept’s ancestors. That is the same understanding a POSITA would have applied to the “is-a” edges in Goldberg’s SNOMED-integrated knowledge graph.

193. Goldberg’s disclosure is also consistent with the ’883 Patent’s description that SNOMED and LOINC provide reasoning for clinical concepts and laboratory tests. The ’883 Patent describes the use of SNOMED-CT and LOINC for

reasoning. *See, e.g.*, Ex. 1001 at 5:13–19, 8:38–58. Under that description, Goldberg’s ontology-integrated graph with SNOMED “is-a” relationships supports the claimed reasoning capability.

194. To the extent Goldberg is found not to disclose element [1f], Dang discloses a knowledge graph configured to support reasoning over ancestor–descendant relationships among its nodes. Dang’s diagnosis method infers evidential nodes by “identifying one or more candidate evidential nodes among the set of factual nodes, wherein the identifying is based on hyponymy relationships between at least two factual nodes in the set of factual nodes.” Ex. 1008 at ¶ [0010]; claim 9. Dang explains that hyponymy is a relationship “between a generic term (a hypernym) and an instance (hyponym) of the generic term,” and that the semantic field of a hypernym, “also known as a superordinate, is broader than the semantic field of a hyponym.” *Id.* at ¶ [0048]. Dang’s factual nodes include “symptom nodes, test nodes, context nodes.” *Id.* at ¶ [0047].

195. A POSITA would have understood a hypernym–hyponym relationship between two factual nodes to be a parent–child (ancestor–descendant) relationship between those nodes: the hypernym node is the superordinate, more general concept, and the hyponym node is a more specific instance of it. That is the same structure as a SNOMED “is-a” relationship. Ex. 1006, FIGS. 6–7. Dang ties these relationships to nodes of the knowledge graph (“hyponymy relationships between at least two

factual nodes”), not merely to words in free text. Ex. 1008 at ¶ [0010]. Although Dang describes a single generic-to-specific step, a POSITA would have understood that applying that step repeatedly across a multi-level hierarchy such as SNOMED’s reaches every ancestor or descendant of a concept. Dang’s knowledge graph is therefore “configured to support reasoning capability by allowing traversal of ancestor–descendant relationships” among its finding (symptom) and investigation (test) nodes.

196. A POSITA would have been motivated to configure Goldberg’s hybrid knowledge graph to support Dang’s hyponymy-based reasoning for the reasons discussed above. *Supra* § X.A.1. Goldberg’s graph already contains SNOMED “type-of” relations on its diagnostic paths and sub-type relations between diseases. Ex. 1006 at 10:3–23; FIGS. 10–11. Dang shows how such generic-to-specific relationships between nodes are used during diagnostic inference. The result would have been predictable because it applies Dang’s reasoning to relationships Goldberg’s graph already stores.

h. Element [1g]

197. Element [1g] recites: “wherein said medical terminology ontology includes a Systemized Nomenclature of Medicine—Clinical Terms (SNOMED-CT) database, or a Logical Observation Identifiers, Names, and Codes (LOINC) database.”

198. Goldberg expressly names both SNOMED and LOINC as components of its hybrid knowledge graph. Goldberg describes SNOMED as an encompassing ontology covering hundreds of thousands of medical terms and identifies LOINC as an ontology dedicated to laboratory tests and procedures. Ex. 1006 at 1:54–2:1; 6:5–28. Goldberg’s express integration of SNOMED and LOINC supplies the claimed ontology limitation. Goldberg’s first graph is derived from a healthcare ontology such as SNOMED, and its second graph is derived from an ontology dedicated to laboratory tests and procedures such as LOINC. Ex. 1006 at 6:5–28.

199. The ’883 Patent specification itself confirms that SNOMED-CT and LOINC were known, pre-existing ontologies. Ex. 1001 at 4:12–37, 5:13–19, 7:4–13. As I discussed in the Technology Background section, SNOMED-CT and LOINC are established controlled vocabularies for clinical diagnoses and laboratory tests. *Supra* § VI.A.

i. Element [1h]

200. Element [1h] recites: “identifying, using said reasoning capability, a set of ancestor–descendant relationships of said knowledge graph that are relevant to said medical query, wherein said reasoning capability including traversing said knowledge graph to identify one or more of said medical condition nodes, said finding nodes, or said investigation nodes that are relevant to said medical query.”

201. Goldberg’s diagnostic algorithm provides the claimed reasoning capability. Goldberg’s diagnostic algorithm expressly traverses the knowledge graph at runtime. Claim 1 recites defining pathways by selecting connectors that connect pairs of nodes, where the connectors are derived from ontological and clinical relationships between medical terms. Claim 9 further recites selecting pathways between pairs of nodes by selecting connectors that represent aspects of the patient’s health data. Ex. 1006 at 14:35–46; 16:35–45.

202. Because Goldberg’s ontological relationships include the SNOMED “is-a” hierarchy, a POSITA would understand that pathway selection operates over ancestor–descendant relationships. The pathways are selected in response to the specific patient’s health data and are query-specific. Goldberg provides a “path to decision”—the clinical reasoning path of associations used to make a diagnosis or recommend a next step, including links to original medical textbook or academic paper sources. Ex. 1006 at 8:41–47. Goldberg also discloses multi-hop traversal, tracking jaundice back through a complication to the originating liver disorder. *Id.* at 8:32–40. Goldberg further discloses that multiple pathways may be searched when looking up Pneumonia and that the system may show associations with findings and laboratory tests. Ex. 1006 at 11:55–67. This runtime, multi-pathway search identifies graph nodes relevant to the patient query.

203. Goldberg further discloses that the platform enables “patient-tailored reference lookup incorporated with path-to-evidence citations, findings and suspected diseases.” Ex. 1006 at 12:1–10.

204. Xia corroborates and confirms a POSITA’s understanding in this regard. Xia illustrates how researchers used a concept hierarchy in both building and using a medical knowledge graph. When processing literature records, Xia “traced [each mentioned concept] all the way up to all its ancestors, intending to reflect that mentioning of a concept is also mentioning of its ancestors,” before counting concept occurrences and co-occurrences. Ex. 1010 at 1119. When applying its knowledge graph to diagnosis, Xia took the hierarchy into account: “where we have hierarchy of concepts, when talking about the rank of a diagnosis, we compare with disease concepts of the same level for fair comparison.” *Id.* at 1121. Xia also illustrates “how hierarchical information was used and displayed” in an example diagnosis output, explaining that “[w]hen deciding whether the diagnosis is accurate, we compare weights of other diseases on the same level.” *Id.* at 1123; Fig. 3. These passages are consistent with my opinion that a POSITA would have understood that the “is-a” relationships in a medical knowledge graph such as Goldberg’s can be used to relate a query concept to its ancestors and descendants.

205. The ’883 Patent specification explains that “traversing said knowledge graph to identify one or more of said medical condition nodes, said finding nodes,

or said investigation nodes that are relevant to said medical query” encompasses using a stored ontological hierarchy to expand a query’s reach to related concepts. The specification’s reasoning service 72 performs this function by “us[ing] an internal or external taxonomy or any other suitable medical hierarchy in order to perform this inferential reasoning.” Ex. 1001 at 23:64–67. The specification gives the example that, if a diagnostic factor is “double quotidian fever,” the service finds its parent “fever” and traverses upward in the taxonomy to add relevant diagnostic factors. *Id.* at 23:66–24:4. Goldberg’s SNOMED hierarchy and runtime pathway selection are consistent with this description.

206. Goldberg’s diagnostic algorithm already selects query-specific pathways using ontological and clinical relationships. Deploying that runtime pathway selection through Gnanasambandam’s API architecture would have been predictable because the references address complementary aspects of the same healthcare knowledge-graph system.

207. To the extent Goldberg is found not to disclose element [1h], Dang discloses identifying, using a reasoning capability, ancestor–descendant relationships of the knowledge graph that are relevant to a medical query, and traversing the graph to identify relevant nodes. Dang’s graph-based diagnosis process begins with “data from a set of patient complaint records” and “detects factual nodes (for example, symptom nodes, test nodes, context nodes)” that match

the input, using “semantic analysis and matching” with a confidence value (“Confidence of fact=similarity(input->node)”). Ex. 1008 at ¶ [0047]; *see also id.* at ¶ [0009] (“measuring, using semantic matching, a similarity between an input text [and] one or more nodes in the knowledge graph data structure, the input text comprising one or more observed symptoms”). The patient’s complaint is the medical query.

208. Dang then “infer[s] (step 504) evidential nodes from the nodes detected (step 502) in the input data,” and “the inferring may be performed using reasoning through hyponymy.” Ex. 1008 at ¶ [0048]. The hyponymy relationships used are “between at least two factual nodes” of the knowledge graph. *Id.* at ¶ [0010]. Because this reasoning runs on the nodes matched from the patient’s input, the hypernym–hyponym relationships it identifies are the ancestor–descendant relationships “relevant to said medical query.” Dang then continues through the graph and infers disease feature nodes from the evidential nodes (step 506) and possible diseases from the feature nodes (step 508), with a confidence value computed at each step. *Id.* at ¶¶ [0049]–[0050]. Dang adds that “disease information may be aggregated through hyponymy analysis.” *Id.* at ¶ [0051]. This is “traversing said knowledge graph to identify one or more of said medical condition nodes, said finding nodes, or said investigation nodes that are relevant to said medical query.” Dang also claims this

sequence of detecting factual nodes, inferring evidential nodes, inferring feature nodes, and inferring possible disease nodes. *Id.*, claims 7, 9.

209. For the reasons given above, a POSITA would have been motivated to use Dang’s query-time hyponymy reasoning within Goldberg’s pathway selection, which already operates over SNOMED “type-of” relations. Doing so would allow a patient’s specific finding to be matched to broader findings stored in the graph, and the reverse. The result would have been predictable because both references reason over the same kind of graph for the same purpose: identifying candidate diseases from a patient’s findings.

j. Element [1i]

210. Element [1i] recites: “outputting, responsive to said medical query, an output to said software application, wherein said output includes data related to one or more of said medical condition nodes, said finding nodes, or said investigation nodes that are relevant to said medical query in order for said software application to generate a medical determination.”

211. Gnanasambandam discloses generating a medical determination via API-connected software applications. Gnanasambandam’s APIs provide analytics output to consuming software applications, including data that may be used to generate a diagnosis of a medical condition for a patient. The API platform delivers this output to client devices running the autonomous multipurpose application,

which then presents the medical determination. Ex. 1007 at 101:12–37, 102:22–24, 100:24–34.

212. Goldberg discloses outputting diagnostic results. Its diagnostic algorithm generates differential diagnoses for suspected disorders, and the application generates a list of suspected diseases and complications and suggests the next best step to narrow the list. Ex. 1006, Abstract; 7:25–50.

213. A POSITA would have been motivated to deliver Goldberg’s diagnostic output through Gnanasambandam’s API framework because Goldberg already generates differential diagnoses and next-step recommendations, while Gnanasambandam’s APIs deliver knowledge-graph-based analytics to consuming software applications. The API layer would serve as a predictable transport mechanism for Goldberg’s existing diagnostic output, allowing the software application to generate a medical determination.

3. Dependent Claim 2

214. Claim 2 depends from claim 1 and recites that “each of said investigation nodes is linked to a corresponding result node holding a result of said each investigation node, each of said corresponding result nodes being linked to said at least one of said medical condition nodes to which said each investigation node is linked.” I incorporate by reference my discussion of claim 1 above. *Supra* § X.A.2.

215. The ‘883 Patent specification provides context for understanding the plain meaning of the phrase “a corresponding result node holding a result of said each investigation node.” The specification describes a graph database where each elemental diagnostic investigation (“EDI”) node representing a laboratory test or examination is associated with a separate “EDI Result” node that holds the condition-specific result of that test. The specification explains that “the single laboratory test EDI urine dipstick for white cells 332, is at the end of the linkage, with its possible result, EDI Result low 326, being between nodes 320 and 332.” Ex. 1001 at 19:17–20. Further, “an EDI Result node provides a result specific to its indicated condition, and is indicative of the condition only when the result has a particular value, such as ‘low’ in node 326.” *Id.* at 19:20–24. The architecture places the EDI Result node between the investigation group (which links to the condition node) and the EDI node (which holds the elemental test description), forming the chain: Condition → Investigation Group → EDI Result → EDI.

216. The specification further describes the content of EDI Result nodes. Each EDI Result node “includes the actual result from the laboratory test or examination (e.g., ‘low,’ 32, greater than 100, in-between 50 and 100, yellowish, etc.) to indicate the condition,” and the type of the result may be “a value, range, a binary, and ordinal, nominal, etc.” Ex. 1001 at 21:56–61. The specification also states that “EDI Result nodes include the result of the test (described using terms

such as ‘high’ or ‘low’, using ranges, using binary values, using ordinal types, etc.).” *Id.* at 22:12–16. Thus, the ‘883 Patent’s own description of a “result” associated with an investigation is a qualifier—such as “high” or “low”—indicating the outcome of a laboratory test in the context of a particular medical condition.

217. Goldberg discloses this same separation between test identity and result qualifier. Goldberg's FIG. 1 is a table depicting “a sample of knowledge base rules for Celiac disease.” Ex. 1006 at 6:65–7:2. The table includes a row with the Type of Rule “LABORATORY TESTS” and the Rule Target “Amylase [Enzymatic activity/volume] in Serum, Plasma or Blood: HIGH,” linked to Celiac disease with a probability of 5–10. Ex. 1006, FIG. 1. The Rule Target entry contains two distinct components separated by a colon delimiter: (1) the test identity “Amylase [Enzymatic activity/volume] in Serum, Plasma or Blood,” which is a LOINC-coded laboratory test name; and (2) the result qualifier “HIGH” which indicates the condition-specific outcome of that test that is associated with Celiac disease.

Celiac disease (disorder)	LABORATORY TESTS	Amylase [Enzymatic activity/volume] in Serum, Plasma or Blood:HIGH	5-10	1	https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3462881/
------------------------------	---------------------	---	------	---	---

Ex. 1006 at FIG. 1 (excerpted and annotated)

218. The result qualifier “HIGH” in Goldberg is the same type of information that the ‘883 Patent describes as the content of a result node. As discussed above, the ‘883 Patent specification states that EDI Result nodes hold

results “described using terms such as ‘high’ or ‘low.’” Ex. 1001 at 22:12–16. Goldberg’s “HIGH” qualifier performs the identical function: it specifies the condition-specific result value (high amylase) that is associated with the medical condition (Celiac disease). Just as the ‘883 Patent’s EDI Result node 326 holds “low” as the result of the urine dipstick test for a particular condition, Goldberg’s “:HIGH” holds “HIGH” as the result of the Amylase test for Celiac disease.

219. FIG. 4 of Goldberg provides an expanded comparison for Celiac disease and Crohn’s disease, which confirms that laboratory test entries follow the same pattern. FIG. 4 includes the entry “Amylase [Enzymatic activity/volume] in Serum, Plasma or Blood: HIGH” under the LABORATORY TESTS category for Celiac disease with a likelihood of 5–10. Ex. 1006, FIG. 4. This further demonstrates that Goldberg’s knowledge base consistently associates result qualifiers with laboratory tests in a manner that links the test result to a specific condition.

LABORATORY TESTS	Amylase [Enzymatic activity/volume] in Serum, Plasma or Blood: HIGH	5-10	
COMPLICATION	Anemia of chronic disorder (disorder)		5-10
SIGNS	Ataxia (finding)	1-2	
ASSOCIATED CONDITION	Autoimmune hepatitis (disorder)	1-2	5-10
LABORATORY TESTS	Ceruloplasmin [Mass/volume] in Serum or Plasma: LOW	2-5	

Ex. 1006 at FIG. 4 (excerpted and annotated)

220. In Goldberg’s knowledge base, investigation nodes remain condition-independent. For example, Goldberg’s claim 9 recites a “fifth set of nodes representing the plurality of laboratory tests” as part of its medical knowledge graph.

Ex. 1006 at 16:8–24. Goldberg’s claim 11 recites that the system determines “a specific laboratory test in the plurality of laboratory tests to be administered to the patient” from pathways that connect laboratory tests to disease nodes. Ex. 1006 at 16:55–17:3. For the same laboratory test (e.g., Amylase) to be usable across multiple disease pathways—Celiac disease, pancreatitis, and others—the test node must remain condition-independent, while the condition-specific result value (e.g., HIGH for Celiac disease, or a different threshold for pancreatitis) resides in a separate data element linked to both the test and the condition. This is precisely the architecture shown in FIG. 1, where the colon-delimited “:HIGH” separates the condition-specific result from the condition-independent test identity.

221. To the extent Goldberg does not expressly implement the result qualifier as a separate, structurally distinct node in a graph database, it would have been obvious to a POSITA to implement the result qualifier as a separate result node linked to the investigation node and the condition node. This is because the fundamental design principle of maintaining condition-independence for reusable data elements was well understood in the art of knowledge graph construction before the ‘883 Patent’s priority date. Knowledge graphs in healthcare and other domains routinely separate entity identity from context-dependent attributes to enable the same entity to participate in multiple relational contexts.

222. In the case of laboratory tests, the separation of test identity from result interpretation is inherent in the structure of laboratory medicine itself: the same test (e.g., Amylase) can yield different results (e.g., HIGH, LOW, normal), and the clinical significance of each result depends on the suspected condition. A POSITA designing a knowledge graph that includes laboratory tests linked to multiple conditions would have recognized that storing the result qualifier in a separate linked node is the natural and predictable implementation for maintaining the test node's reusability. This is a straightforward application of well-known graph database design principles to the clinical data that Goldberg organizes.

223. To the extent Goldberg's result qualifier is found not to be a "result node," Dang's knowledge graph includes "an observation item node" and "a lab test and check node" as distinct node types, connected by "[d]irectional edges" that "indicate a flow of information from one node to another." Ex. 1008 at ¶¶ [0027], [0036] (a body temperature of 38° C "may be indicated ... via the checks and tests nodes."); FIG. 1A. As shown in annotated FIG. 1A, the Observation Item node is positioned on the path between the Lab Test & Check node and the Disease node. A POSITA would have been motivated to implement Goldberg's ":HIGH" qualifier as a separate node, following Dang's structure. Doing so keeps the laboratory-test node reusable across diseases while the result is held in its own node.

4. Dependent Claim 3

224. Claim 3 recites “A method as recited in claim 2 wherein each of said investigation nodes represents a laboratory test or an examination, and wherein said result node corresponding to said each of said investigation nodes represents a result of said laboratory test or said examination.” I incorporate by reference my discussion of dependent claim 2 above. *Supra* § X.A.3.

225. Goldberg discloses or renders obvious this element for the reasons already discussed above in connection with claim 2. Goldberg’s investigation nodes represent laboratory tests—for example, the Amylase test shown in FIG. 1 is expressly categorized under “LABORATORY TESTS” and identified by a LOINC-coded test name: “Amylase [Enzymatic activity/volume] in Serum, Plasma or Blood.” Ex. 1006, FIG. 1; *see also id.* at 6:17–28 (describing integration of “an ontology dedicated to laboratory tests and procedures (e.g., LOINC).”).

226. The corresponding result qualifier “HIGH” represents a result of that laboratory test—specifically, the condition-specific outcome indicating elevated amylase associated with Celiac disease. As discussed in connection with claim 2, this result qualifier is the same type of information that the ‘883 Patent describes as the content of a result node.

5. Dependent Claim 4

227. Claim 4 recites “A method as recited in claim 1 wherein said finding nodes are diagnostic factor nodes.” I incorporate by reference my discussion of independent claim 1 above. *Supra* § X.A.2.

228. The '883 Patent specification explains that a diagnostic factor is “a factor that the physician can observe, detect or measure quickly or have the patient describe to him during an office visit (traditional signs and symptoms),” whereas an investigation is “a laboratory test or an examination performed by a specialist.” Ex. 1001 at 9:21–28. Examples of diagnostic factors include “patient has a headache,” “patient skin appears yellow,” and “patient reflexes appear normal.” *Id.* Claim 4 requires that the “finding nodes” disclosed in element [1b] function as diagnostic factor nodes, e.g., nodes representing symptoms that a physician can detect during a patient encounter.

229. Goldberg’s hybrid knowledge graph includes symptom nodes that function as diagnostic factor nodes within the meaning of the '883 Patent. Goldberg’s claim 9 recites “a third set of nodes representing the plurality of symptoms” as one of the six node sets in its medical knowledge graph. Ex. 1006 at 16:8–24. These symptom nodes represent clinical observations: signs and symptoms that a physician detects during a patient encounter.


230. Goldberg's figures provide specific examples. FIG. 4 illustrates the knowledge base with express "SYMPTOM/COMPLAINT" entries (e.g., "Abdominal pain (finding)") and "SIGNS" entries (e.g., "Ataxia (finding)")—both of which are diagnostic factors that a physician can observe or that a patient can report during an office visit. Ex. 1006, FIG. 4.

FIG. 4

Type of Feature	Feature	Likelihood with Celiac disease	Likelihood with Crohn's disease
SYMPTOM/COMPLAINT	Abdominal pain (finding)	5-10	
SYMPTOM/COMPLAINT	Abdominal pain (finding)# Chronic (qualifier value)	90-100	
COMPLICATION	Acute hepatitis (disorder)	2-5	
COMPLICATION	Acute intestinal obstruction (disorder)	5-10	
COMPLICATION	Acute upper gastrointestinal hemorrhage (disorder)	2-5	
ASSOCIATED CONDITION	Addison's disease (disorder)	1-2	
LABORATORY TESTS	Amylase [Enzymatic activity/volume] in Serum, Plasma or Blood:HIGH	5-10	
COMPLICATION	Anemia of chronic disorder (disorder)		5-10
SIGNS	Ataxia (finding)	1-2	
ASSOCIATED CONDITION	Autoimmune hepatitis (disorder)	1-2	5-10
LABORATORY TESTS	Ceruloplasmin [Mass/volume] in Serum or Plasma:LOW	2-5	
COMPLICATION	Chronic diarrhea (disorder)	10-20	5-10
COMPLICATION	Chronic liver disease (disorder)	2-5	2-5
COMPLICATION	Constipation (disorder)	5-10	
COMPLICATION	Delayed puberty (disorder)	2-5	5-10
COMPLICATION	Depressive disorder (disorder)	2-5	
ASSOCIATED CONDITION	Diabetes mellitus type 1 (disorder)	1-2	
LABORATORY TESTS	Endomysium Ab [Units/volume] in Serum:HIGH	90-95	
COMPLICATION	Exacerbation of Crohn's disease of large intestine (disorder)		90-100
COMPLICATION	Exocrine pancreatic insufficiency (disorder)		2-5
SYMPTOM/COMPLAINT	Failure to gain weight (finding)	5-10	
COMPLICATION	Failure to thrive (disorder)	60-70	30-40
COMPLICATION	Iron deficiency anemia (disorder)	5-10	5-10
ASSOCIATED CONDITION	Migraine (disorder)	2-5	2-5
COMPLICATION	Osteopenia (disorder)	2-5	5-10
RISK FACTORS	Periods of life (qualifier value):ADULTHOOD	30-40	40-50
COMPLICATION	Vitamin D deficiency (disorder)	2-5	10-20

Ex. 1006 at FIG. 4

231. FIG. 10 shows finding nodes including Abdominal pain RUQ, Hepatomegaly, Cardiomegaly, and Left ventricular hypertrophy—all of which are clinical observations or physical findings that a physician can detect during examination. Ex. 1006, FIG. 10; 10:13–15. Each of these finding nodes represents information that can be observed, detected, or measured during a patient encounter.



232. Goldberg confirms that the symptom and finding nodes function as diagnostic factors in the system’s diagnostic process. Goldberg discloses that “[u]pon a list of given history, symptoms, signs and other clinical features the app generates a list of suspected diseases and complications.” Ex. 1006 at 7:27–30. The

system uses these symptoms, signs, and clinical features for running its “mathematical inference algorithm that finds the most likely explanation for the patient’s condition.” *Id.* at 8:15–20. In other words, Goldberg’s symptom and finding nodes are the clinical observations that the system uses to generate differential diagnoses and recommend next steps. A POSITA would have understood Goldberg’s symptom nodes, sign nodes, and other clinical-feature nodes as diagnostic factor nodes within the meaning of claim 4.

6. Dependent Claim 5

233. Claim 5 recites that “said software application is an artificial intelligence (AI) application.” I incorporate by reference my discussion of independent claim 1 above. *Supra* § X.A.2.

234. Goldberg’s platform uses artificial intelligence techniques both to reason over its knowledge graph and to build it. Goldberg’s application uses patient inputs (e.g., demographics, history, symptoms, signs, lab results) and “their corresponding associations within the knowledge base to run a mathematical inference algorithm that finds the most likely explanation for the patient’s condition.” Ex. 1006 at 8:14–21. Goldberg explains that “App algorithms are used that are based on mathematical inference computations adapted to handle the large number of associations between the various clinical features,” and that the algorithm “takes into consideration the strength of the various associations as indicated in the

knowledge base as well as data such as the cost or pain and time it takes to get required lab tests, imaging or other diagnostic procedures.” *Id.* at 7:29–42. Goldberg further states that “AI and machine learning algorithms can be applied to large sets of individual patient records to look for patterns and create ‘clinical rules derived from big data’” that are added to the hybrid knowledge graph. *Id.* at 14:7–18. This inference-based diagnostic reasoning, together with machine-learned rule generation, is a form of artificial intelligence.

235. Goldberg’s user-facing application is described primarily as an interface for entering patient data and viewing diagnostic results. Ex. 1006 at 7:25–50. To the extent Goldberg does not expressly characterize the overall software application in AI terms, Gnanasambandam supplies this express disclosure.

236. Gnanasambandam’s claim 1 recites “publishing, for implementation in the autonomous multipurpose application, at least one API that uses artificial intelligence to perform analytics based on a knowledge graph pertaining to ontological data of a medical condition.” Ex. 1007, claim 1. The autonomous multipurpose application is the software application that subscribes to and invokes these AI-powered APIs. *Id.* at 100:24–34; 101:12–37; 102:37–46. Those analytics are performed by the cognitive intelligence platform’s “artificial intelligence engine 109 (‘AI Engine’ in FIG. 1) that uses one or more machine learning models.” *Id.* at 18:30–36; FIG. 1.




FIG. 1

Ex. 1007 at FIG. 1 (annotated)

237. These machine learning models are trained to generate knowledge graphs, transform unstructured data using knowledge graphs and logical structures, and identify possible health-related information by comparing patient data to the knowledge graph. Ex. 1007 at 18:55–19:15. The AI engine performs the analytics underlying the API outputs, including diagnoses of medical conditions, and “may continuously learn based on input data (e.g., evidence-based guidelines, clinical trials, physician research, electronic medical records, etc.) and modify the individual elements and predicates.” *Id.* at 8:24–39. A software application that invokes these

APIs to obtain knowledge-graph analytics is an AI application within the ordinary meaning of that term.

238. In the combined system, Gnanasambandam's software application running on a client device queries the knowledge graph via an API and receives the output for generating a medical determination. A POSITA would have been motivated to deploy Goldberg's knowledge-graph-based diagnostic system through Gnanasambandam's AI-powered API platform because both references operate in the same field of AI-driven healthcare knowledge systems, and the combination would enable Goldberg's diagnostic reasoning to be delivered programmatically to external software applications through APIs that expressly use artificial intelligence. Goldberg's system already runs an AI-based inference algorithm that selects pathways and generates differential diagnoses. Ex. 1006, Abstract. Gnanasambandam's platform formalizes this type of AI-based analytics for deployment through published APIs accessible by client devices. Ex. 1007 at 100:24–34; 101:12–25. The resulting AI application that queries a medical knowledge graph via an API and receives diagnostic outputs would have been entirely predictable because each reference independently employs AI in the context of medical knowledge graph analytics, and the combination merely applies Gnanasambandam's API delivery mechanism to Goldberg's existing AI-based diagnostic engine.

7. Dependent Claim 6

239. Claim 6 recites that “said medical query includes at least one patient diagnostic factor extracted from an electronic health record (EHR) of a patient,” and further requires “matching said patient diagnostic factor to one of said finding nodes, wherein said output is a medical condition associated with one of said medical condition nodes to which said one of said finding nodes is linked.” I incorporate by reference my discussion of independent claim 1 above. *Supra* § X.A.2.

240. Goldberg’s patient management application receives patient condition data including “history, symptoms, signs and other clinical features” and uses them as inputs to its diagnostic algorithm. Ex. 1006 at 7:25–40. The application “uses the inputs and their corresponding associations within the knowledge base to run a mathematical inference algorithm that finds the most likely explanation for the patient's condition.” Ex. 1006 at 8:15–25. This process matches a patient diagnostic factor to a finding node: the system takes the patient’s reported symptom (a diagnostic factor) and matches it against symptom nodes (finding nodes) in the knowledge graph to identify associated diseases.

241. Goldberg confirms that the output of this matching process is a medical condition. The application “generates a list of suspected diseases and complications” and “suggests the next best step to narrow down that list by ruling in some diseases or ruling out other high-risk diseases.” Ex. 1006 at 7:27–32; *see also id.*, Abstract.

The differential diagnoses (suspected diseases) are medical conditions associated with the condition nodes in the knowledge graph to which the matched finding nodes (symptoms) are linked.

242. FIG. 5 of Goldberg further illustrates this matching-and-output workflow. FIG. 5 shows a table listing diseases alongside their likelihood of presenting with given features (Failure to thrive & Ataxia). Ex. 1006, FIG. 5; 9:13–20. The table demonstrates that, when specific findings (patient diagnostic factors) are entered into the system, the output is a ranked list of medical conditions associated with those findings (e.g., Celiac disease has a 60-70% likelihood of presenting with Failure to thrive, and a 1-2% likelihood of presenting with Ataxia).

FIG. 5


Feature	Disease	Likelihood to present with disease
Failure to thrive (disorder)	Adrenal cortical hypofunction (disorder)	20-30
Failure to thrive (disorder)	Adrenocorticotrophic hormone deficiency (disorder)	30-40
Failure to thrive (disorder)	Anemia (disorder)	40-50
Failure to thrive (disorder)	Atrial septal defect (disorder)	5-10
Ataxia (finding)	Celiac disease (disorder)	1-2
Failure to thrive (disorder)	Celiac disease (disorder)	60-70
Failure to thrive (disorder)	Crohn's disease (disorder)	30-40
Failure to thrive (disorder)	Chronic diarrhea (disorder)	40-50
Failure to thrive (disorder)	Chronic liver disease (disorder)	20-30
Failure to thrive (disorder)	Cystic fibrosis (disorder)	30-40
Ataxia (finding)	Degenerative disease of the central nervous system (disorder)	70-80
Failure to thrive (disorder)	Disorder of the urea cycle metabolism (disorder)	60-70
Failure to thrive (disorder)	Glucosylceramide beta-glucosidase deficiency (disorder)	60-70
Failure to thrive (disorder)	Glycogen storage disease type Ia (disorder)	60-70
Failure to thrive (disorder)	Malabsorption syndrome (disorder)	20-30
Failure to thrive (disorder)	Milk-induced pulmonary disease in infant (disorder)	5-10
Failure to thrive (disorder)	Mitochondrial cytopathy (disorder)	60-70
Failure to thrive (disorder)	Tyrosinemia (disorder)	20-30
Failure to thrive (disorder)	Tyrosinemia (disorder)	60-70
Ataxia (finding)	Tyrosinemia (disorder)	5-10
Failure to thrive (disorder)	Ventricular septal defect (disorder)	10-20
Ataxia (finding)	Wilson's disease (disorder)	5-10

Ex. 1006 at FIG. 5

243. Goldberg's specification describes patient data being entered through the patient management application's user interface. Ex. 1006 at 7:25–40. While Goldberg references the development of EHR systems and their role in healthcare (Ex. 1006 at 1:46–58), and while the system receives the same types of clinical data

that are stored in EHRs (history, symptoms, signs, lab results), Goldberg does not expressly state that its system extracts patient diagnostic factors from an EHR.

244. Gnanasambandam expressly discloses extracting patient diagnostic information from electronic medical records. Gnanasambandam discloses a method in which the platform receives “an electronic medical record (EMR) including notes pertaining to a patient,” processes those notes to “obtain indicia including a subject, an object, a word, a cardinal, a phrase, a concept, a sentence, a predicate, or some combination thereof,” and then identifies “a possible medical condition of the patient by identifying a similarity between the indicia and a knowledge graph representing knowledge pertaining to the possible medical condition.” Ex. 1007 at 56:50–57:37, FIG. 25; *see also id.* at 125:10–29 (Clause 122), 126:28–47 (Clause 132).



Ex. 1007 at FIG. 25

245. Gnanasambandam further discloses that “[t]he EMRs may be provided directly to the cognitive intelligence engine 102 and/or retrieved using an application programming interface (API) from an EMR system used by the physician. The patient notes may be extracted from the EMRs.” Ex. 1007 at 53:4–12.

246. Gnanasambandam also discloses that the knowledge graphs used by its platform are generated and updated using data extracted from electronic medical records. The knowledge graphs and logical structures “may be generated and updated continuously or on a periodic basis by an artificial intelligence engine with evidence-based guidelines, physician research, patient notes in EMRs, physician feedback, and so forth.” Ex. 1007 at 8:24–39. The AI engine “may continuously learn based on input data (e.g., evidence-based guidelines, clinical trials, physician research, electronic medical records, etc.) and modify the individual elements and predicates.” Ex. 1007 at 8:24–39.

247. A POSITA would have been motivated to configure the combined Goldberg/Gnanasambandam system to extract patient diagnostic factors from an EHR because Goldberg’s diagnostic algorithm already processes the same types of clinical data that EHRs contain—history, symptoms, signs, and lab results—and Gnanasambandam expressly provides the architecture for receiving and processing EHR data through its AI-powered platform. Ex. 1007 at 53:4–12; 56:62–57:1; 57:1–2. Integrating EHR data as an input source for a clinical decision support system was

a well-established practice in healthcare IT before the ‘883 Patent’s priority date, and both references operate in this field. A diagnostic system that extracts patient factors from EHR records, matches them to finding nodes in the knowledge graph, and outputs associated medical conditions would have been predictable because it merely connects Goldberg’s existing diagnostic matching workflow to an EHR data source that Gnanasambandam discloses.

248. Dang provides additional support. Dang’s input is a patient’s “electronic medical record (EMR)” indicating, e.g., “a runny nose, a fever,” and a temperature reading. Ex. 1008 at ¶ [0036]. Using “patient complaint records,” Dang “detect[s] ... factual nodes using semantic analysis and matching” (step 502). *Id.* at ¶ [0047]. It then outputs “the top disease list as a suspected disease list.” *Id.* at ¶ [0056].

8. Dependent Claim 7

249. Claim 7 recites: “A method as recited in claim 1 wherein said medical query is a clinical state of a patient including at least one likely medical condition of said patient and at least one patient diagnostic factor, said method further comprising: matching said patient diagnostic factor to one of said finding nodes, and wherein said output is an unknown diagnostic factor associated with one of said finding nodes that is not yet known to be present or absent in said patient.” I

incorporate by reference my discussion of independent claim 1 above. *Supra* § X.A.2.

250. Claim 7 recites three additional limitations relative to independent claim 1: (1) the medical query is a clinical state of a patient that includes at least one likely medical condition and at least one patient diagnostic factor; (2) the system matches the patient diagnostic factor to a finding node; and (3) the output is an unknown diagnostic factor associated with one of the finding nodes that is not yet known to be present or absent in the patient.

251. As to the first limitation, Goldberg's application receives "a list of given history, symptoms, signs and other clinical features," generates "a list of suspected diseases and complications," and then "suggests the next best step to narrow down that list." Ex. 1006 at 7:25–35. It creates "a differential diagnosis list based on inputs entered for a specific patient" and can "rank likely differential diagnosis results by their probabilistic score." *Id.* at 8:10–24. At each stage, the patient's clinical state therefore includes both likely medical conditions (the ranked differential) and known diagnostic factors (the entered findings), and that state is the input from which Goldberg recommends a next step. *Id.* at 7:25–42; 8:10–32.

252. As to the second limitation, Goldberg's application "uses the inputs and their corresponding associations within the knowledge base" to reach its differential, which requires matching the patient's reported findings to the corresponding finding

nodes. Ex. 1006 at 8:14–21. I address this matching in connection with element [1h] and later for claim 20. *Supra* § X.A.2.i; *infra* § X.A.19.

253. As to the third limitation, Goldberg’s next step can be an item of patient information not yet known. For example, the application “may recommend test A over test B” because “[a] positive or negative answer for test A would increase the probability of suspected diagnosis.” Ex. 1006 at 8:20–32. Goldberg does not, however, describe the next step as a specific symptom or other finding to ask the patient about, or explain how candidate next steps are ranked. Dang supplies that disclosure.

254. Dang generates “a best next inquiry question, for use in a medical diagnosis process,” based on “likelihood values that a set of input symptoms exhibits one or more diseases.” Ex. 1008 at ¶ [0003]. After outputting a suspected disease list (FIG. 6), Dang’s computer system “selects (step 704) sub-symptoms or attributes as next inquiry candidates” and “selects (step 706) symptoms derived from related diseases as next inquiry candidates.” *Id.* at ¶¶ [0056], [0058]; FIGS. 6–7. For each candidate, it “calculates (step 710) its degree of distinguish (DoD),” and it then “selects (step 712) a candidate with a maximal DoD as the next inquiry suggestion.” *Id.* at ¶ [0058]; FIG. 7. Dang bases this selection on the Gini index, explaining that “the higher Gini index indicates greater inequality of symptoms and a more effective inquiry question.” *Id.* at ¶ [0040].

255. The selected inquiry is the “unknown diagnostic factor” that claim 7 recites because it is a symptom or attribute that the system proposes to ask the patient about, so a POSITA would have understood it to be a factor not yet known to be present or absent. Each candidate is a symptom represented in Dang’s knowledge graph and is drawn either from the patient’s input symptoms or from “related diseases,” i.e., disease nodes associated with the patient’s suspected conditions. Ex. 1008 at ¶¶ [0029]–[0031], [0058]; FIGS. 1B, 7.

256. A POSITA would have been motivated to implement Goldberg’s next-best-step recommendation using Dang’s next-inquiry method. Goldberg leaves the implementation open, stating that “there are several alternative ways to implement such an inference engine.” Ex. 1006 at 7:38–45. Dang supplies a specific method and reports that it works: on respiratory-disease test cases, Dang achieved “a top-1 diagnostic accuracy of 61.7% and top-3 diagnostic accuracy of 94.0%,” compared with 32.3% and 55.9% for one benchmark and 41.1% and 47.0% for the other, and needed an average of 13.5 questions compared with 27.7 and 16.6. Ex. 1008 at ¶ [0059]. Those results serve Goldberg’s stated goals of narrowing the differential “by ruling in some diseases or ruling out other high-risk diseases” and reducing “unnecessary tests and diagnostic errors.” Ex. 1006 at 7:25–35; 2:25–30.

257. The combination was predictable: both references reason over a graph of diseases and findings to recommend the next diagnostic step, and Dang’s method

operates on inputs Goldberg already produces, namely the patient’s findings and a ranked list of suspected diseases. Ex. 1006 at 7:25–35; 8:10–24; Ex. 1008 at ¶¶ [0056]–[0058]. Gnanasambandam’s API then delivers the selected inquiry to the software application, as discussed for element [1i]. *Supra* § X.A.2.j.

9. Dependent Claim 8

258. Claim 8 recites that “said each investigation node does not include information specific to a medical condition associated with one of said medical condition nodes to which said each investigation node is linked.” I incorporate by reference my discussion of independent claim 1 above. *Supra* § X.A.2.

259. Goldberg discloses this element. Goldberg’s hybrid knowledge graph derives its laboratory test nodes from LOINC, “an ontology dedicated to laboratory tests and procedures.” Ex. 1006 at 6:17–28. LOINC is a condition-independent taxonomy: it defines laboratory tests by what they measure (e.g., analyte, method, specimen type, and measurement property), not by what condition they are used to diagnose. Thus, laboratory test nodes in Goldberg’s knowledge graph inherit their identity from LOINC and do not include information specific to any medical condition.

260. Goldberg’s FIG. 1 confirms this condition-independence. The LABORATORY TESTS entry shows the LOINC-coded test name “Amylase [Enzymatic activity/volume] in Serum, Plasma or Blood,” i.e., a generic test

identifier that describes the analyte (Amylase), the measurement property (Enzymatic activity/volume), and the specimen type (Serum, Plasma or Blood). Ex. 1006, FIG. 1; 6:65–7:2. None of this information is specific to Celiac disease or any other condition. The condition-specific information (e.g., that the Amylase test is associated with Celiac disease with a probability of 5–10) does not reside in the test node.

261. FIG. 4 is consistent with this because it shows that Goldberg’s knowledge base stores clinical features as generic concepts whose likelihood values are specific to each condition. The same feature (e.g., “Chronic diarrhea (disorder)”) appears once, with a likelihood of 10–20 for Celiac disease and 5–10 for Crohn’s disease. Ex. 1006, FIG. 4. FIG. 4’s laboratory-test entries are populated for Celiac disease, and each is identified by its generic LOINC name (e.g., “Endomysium Ab [Units/volume] in Serum:HIGH”). Nothing in those names is specific to Celiac disease. *Id.* A POSITA would have understood that Goldberg’s laboratory-test concepts follow the same structure as its other features: the test is identified once, and the condition-specific likelihood sits in the rule. Ex. 1006 at 6:20–40; 16:8–24.

262. Goldberg’s claim 9 recites “a fifth set of nodes representing the plurality of laboratory tests” as a structurally distinct category within the medical knowledge graph. Ex. 1006 at 16:8–24. These laboratory test nodes are connected to disease nodes through “a plurality of connectors” derived from ontological and

clinical relationships, with condition-specific information encoded in the connectors (which include probability values per Goldberg claim 10), not in the test nodes. Ex. 1006 at 16:22–31. Goldberg’s claim 11 confirms that the same laboratory test can serve multiple disease pathways: it recites determining “a specific laboratory test in the plurality of laboratory tests to be administered to the patient” from pathways that connect laboratory test nodes to disease nodes. Ex. 1006 at 16:55–17:3. For the same test to participate in pathways leading to different diseases, the test node would be condition independent.

263. Dang confirms this design was well-known and routine. Dang’s test nodes are a separate set, “{t1, t2 ... tj}.” The condition-specific information sits on the connections: “each disease node is connected to one or more feature nodes ... each such connection may have a corresponding conditional probability P,” and $P(f_x|d_n)$ indicates “the likelihood of the feature f_x occurring when the disease d_n [is] observed.” Ex. 1008 at ¶¶ [0029], [0031]. Because disease-specific probabilities attach to the disease-to-feature connections rather than to the test nodes, a POSITA would have understood Dang’s test nodes to contain no information specific to any disease. This further confirms that a POSITA would understand or find obvious that Goldberg’s nodes contain no information specific to any disease.

10. Dependent Claim 9

264. Dependent claim 9 recites: “A method as recited in claim 1 wherein each of said investigation nodes is disambiguated across said plurality of medical condition nodes.” I incorporate by reference my discussion of independent claim 1 above. *Supra* § X.A.2.

265. Goldberg discloses this element through its integration of LOINC. LOINC assigns a unique identifier and standardized name to each laboratory test, which is the purpose of a controlled vocabulary: to provide standardized naming for laboratory observations across all clinical contexts. Goldberg derives its laboratory test nodes from LOINC, ensuring that each test node carries a single identity regardless of which disease it is linked to. Ex. 1006 at 6:17–28.

266. FIG. 1 illustrates this disambiguation. The test “Amylase [Enzymatic activity/volume] in Serum, Plasma or Blood” appears under LABORATORY TESTS for Celiac disease. Ex. 1006, FIG. 1. FIG. 4 lists the same test, identified by the same LOINC-coded name, in a query result generated from the same knowledge base. Ex. 1006, FIG. 4. More generally, Goldberg describes querying for “all the diseases that share a combination of one or more lab test results.” Ex. 1006 at 8:46–57. That query presupposes a single, unambiguous test concept linked to multiple diseases. The test node is not renamed or redefined when it appears in the context of a different disease; it retains the same unambiguous identity. This is disambiguation

across medical condition nodes: the Amylase test is the same even though it is linked to Celiac disease, Crohn's disease, or any other condition in the knowledge graph.

267. Goldberg's claim 9 further supports this reading. Claim 9 recites "a fifth set of nodes representing the plurality of laboratory tests" as a distinct set within the unified knowledge graph. Ex. 1006 at 16:8–24. By defining laboratory tests as a single, dedicated set of nodes—rather than creating separate test nodes for each disease—Goldberg's architecture disambiguates investigation nodes across medical conditions. Each laboratory test has one node in the graph, not multiple condition-specific variants.

11. Dependent Claim 10

268. Dependent claim 10 recites: "A method as recited in claim 1 wherein each of said investigation nodes is independent of said at least one of said medical condition nodes to which said each investigation node is traversable." I incorporate by reference my discussion of independent claim 1 above. *Supra* § X.A.2.

269. Goldberg discloses this element for the same reasons discussed above in connection with claims 8 and 9. Goldberg's laboratory test nodes are derived from LOINC and carry LOINC-coded names and identifiers that are defined independently of any disease. Ex. 1006 at 6:17–28. The test nodes exist as independent entities in the "fifth set of nodes" (claim 9) and are connected to disease nodes through typed connectors. Ex. 1006 at 16:8–24; FIG. 1.

270. The independence of investigation nodes from condition nodes is evident from the architecture of Goldberg’s hybrid knowledge graph. The graph integrates three separately derived components: (1) an ontology-based knowledge graph such as SNOMED, (2) an ontology for laboratory tests and procedures such as LOINC, and (3) clinical probability data linking concepts across these ontologies. Ex. 1006 at 6:5–28. Because the laboratory test nodes originate from LOINC they are independent of the disease nodes in the graph. The connections between test nodes and disease nodes are established through clinical probability data (the “rules” in FIGS. 1 and 4), not through any dependency of the test node on the disease node.

271. FIG. 1 illustrates this independence concretely. The Amylase test is identified by its LOINC-coded name and connected to Celiac disease through a probability value of 5–10. Ex. 1006, FIG. 1. If the Celiac disease node were removed from the graph, the Amylase test node would continue to exist with the same identity and content—it would simply no longer have a connection to that condition. This confirms independence: the investigation node’s existence and content do not depend on any condition node.

272. As discussed above for claim 2, Goldberg’s claims 9 and 11 confirm this independence. Goldberg’s claim 11 recites determining a specific laboratory test from pathways that connect test nodes to disease nodes—meaning the same test node participates in pathways to multiple diseases. Ex. 1006 at 16:60–17:2. A POSITA

would understand that the test node is independent of any condition node because the same test node is reusable across multiple disease pathways.

12. Dependent Claim 11

273. Dependent claim 11 recites: “A method as recited in claim 1 wherein said knowledge graph further includes for each of said investigation nodes, a data structure different from said each investigation node that is traversable to both said each investigation node and said at least one of said medical condition nodes to which said each investigation node is traversable, said data structure including attributes pertaining to said at least one of said medical condition nodes to which said each investigation node is traversable.” I incorporate by reference my discussion of independent claim 1 above. *Supra* § X.A.2.

274. Claim 11 recites “a data structure” rather than a “node.” Other claims in the ‘883 Patent use the term “node.” For example, claim 2 recites “a corresponding result node,” and independent claims 1 and 23 recite “medical condition nodes,” “finding nodes,” and “investigation nodes.” The use of “data structure” in claim 11 instead of “node” indicates that the claimed data structure is not limited to a graph node. A “data structure” is a broader term encompassing any organized container of data, including but not limited to graph nodes. This reading is confirmed by the ‘883 Patent’s specification, which states that “the conditions, factor groups, EDFs, investigation groups, and EDIs, need not necessarily be implemented in a particular

data structure shown or as nodes; each may be implemented using the representation of the underlying database.” Ex. 1001 at 18:5–9. The specification thus contemplates that the data structures in the knowledge graph may be implemented using representations other than nodes.

275. Under the plain meaning of “data structure,” Goldberg discloses this element. In FIG. 1 of Goldberg, each knowledge base rule is a structured data record containing multiple fields: Source Disorder, Type of Rule, Rule Target, Probability, Strength of source, and Link to source. Ex. 1006, FIG. 1; 6:65–7:2. For a LABORATORY TESTS entry, this rule record is a data structure that is different from the investigation node (the LOINC-coded test name) and that includes condition-specific attributes pertaining to the associated medical condition.

276. The rule record satisfies each element of claim 11. First, it is different from the investigation node: the investigation node is the LOINC-coded test identity (“Amylase [Enzymatic activity/volume] in Serum, Plasma or Blood”), while the rule record is the entire structured entry that associates the test with a specific disease through typed relationships and probabilistic values. Second, the rule record is traversable to the investigation node: the Rule Target field contains the laboratory test, connecting the record to the investigation node. Third, the rule record is traversable to the medical condition node: the Source Disorder field identifies “Celiac disease (disorder),” connecting the record to the condition node. Fourth, the

rule record includes attributes pertaining to the medical condition node: the Probability value (5–10) indicates the likelihood that the Amylase test result is associated with Celiac disease specifically; the Strength of source value indicates the reliability of the clinical evidence supporting this association; and the Link to source provides the citation underlying the condition-specific association. Ex. 1006, FIG. 1. Each of these attributes pertains to the relationship between the investigation and the specific medical condition, not to the investigation in the abstract.


277. FIG. 4 confirms that these attributes are condition-specific. The same feature carries different likelihood values for different conditions, e.g., “Failure to thrive (disorder)” at 60–70 for Celiac disease and 30–40 for Crohn’s disease. Ex. 1006, FIG. 4. Features that are not associated with a given condition have no value for that condition. For example, the Amylase entry has a value for Celiac disease and none for Crohn's disease. *Id.* The probability values in the rule record therefore belong to the condition, not to the test itself.

278. FIG. 11 of Goldberg further illustrates the hybrid model, showing qualified features (QF) connected to other concepts through different types of relations, including clinical relations entered by editors and semantic relations originating from ontological sources. Ex. 1006, FIG. 11; 10:30–31. Each rule connecting a medical concept to a qualified feature carries a rate (R) and is typed (e.g., complication, symptom, lab, risk factor). These typed connections with

associated rates are data structures—structured records associating investigation nodes with condition nodes through condition-specific attributes—satisfying each requirement of claim 11.

279. The ‘883 Patent specification further supports this reading. The specification describes the Investigation Group as a data structure between the condition node and the EDI nodes that holds condition-dependent attributes such as ranking (first line, second line, emerging) and ordering information. Ex. 1001, FIG. 4C; 21:56–61; 22:12–16. But as discussed above, the specification makes clear that these data structures “need not necessarily be implemented . . . as nodes.” Ex. 1001 at 18:5–10. Goldberg’s rule records—structured data entries with condition-specific probability, strength, and source attributes on the path between investigation nodes and condition nodes—implement the same functional concept using the representation of Goldberg’s knowledge base.

280. To the extent Goldberg’s disclosure is found not to satisfy claim 11, Dang discloses this additional element through its feature nodes. Dang’s FIG. 1B discloses a three-tier knowledge graph comprising: (1) factual nodes (symptom nodes $\{s1 \dots sm\}$, context nodes $\{c1 \dots cp\}$, and test nodes $\{t1 \dots tj\}$); (2) feature nodes $\{f1, f2 \dots fx\}$; and (3) disease nodes $\{d1, d2 \dots dn\}$. Dang (Ex. 1008) at ¶¶ [0028]–[0029]; FIG. 1B.

100B**FIG. 1B**

Ex. 1008 at FIG. 1B

281. The feature nodes are a structurally distinct intermediate tier positioned between the factual-tier nodes, including test nodes (investigation nodes) and the disease nodes (condition nodes). Each feature node satisfies claim 11. First, each feature node is a data structure different from the investigation node: the test nodes and feature nodes are expressly designated as separate node types occupying distinct

tiers of Dang’s knowledge graph. *Id.* Second, each feature node is traversable to both the investigation node and the condition node: the feature tier receives connections from the factual tier (including test nodes) and connects to disease nodes through conditional probability edges $P(f|d)$. Ex. 1008 at ¶¶ [0030]–[0031]; FIG. 1B. Third, each feature node includes attributes pertaining to the associated medical condition: the conditional probability $P(f|d)$ on each feature-to-disease connection indicates the likelihood of the feature occurring when a specific disease is present — a condition-dependent attribute. *Id.* at ¶ [0031].

282. Dang’s diagnostic process confirms this three-tier traversal at runtime. The graph-based diagnosis process detects factual nodes at step 502, including “symptom nodes, test nodes, context nodes.” Ex. 1008 at ¶ [0047]. The system then infers evidential nodes from the detected factual nodes (step 504), infers disease feature nodes from the evidential nodes (step 506), and infers possible disease nodes from the feature nodes (step 508). *Id.* at ¶¶ [0047]–[0050]; FIG. 5. This traversal path from test nodes through feature nodes to disease nodes demonstrates that Dang’s feature nodes function as intermediate data structures on the path between investigation nodes and condition nodes, carrying condition-dependent probability attributes at each connection.

283. Dang’s claim 7 independently confirms this multi-tier traversal by reciting a method that “detect[s] a set of factual nodes in a knowledge graph data

structure; infer[s] a set of evidential nodes in the set of factual nodes; infer[s] a set of disease feature nodes in the set of evidential nodes; and infer[s] a set of possible disease nodes in the set of diseases feature nodes.” Ex. 1008, claim 7.

284. In the combination, Dang’s three-tier graph topology including distinct feature nodes positioned between test nodes and disease nodes would be added to Goldberg’s hybrid knowledge graph. A POSITA would have been motivated to incorporate Dang’s feature-node architecture into the Goldberg/Gnanasambandam combination for the same reasons discussed in the motivation to combine section: Dang’s three-tier knowledge graph formalizes the intermediate relationships between investigation nodes and condition nodes through dedicated feature nodes with quantified conditional probabilities. *Supra* § X.A.1. Incorporating Dang’s feature-node tier into Goldberg’s hybrid knowledge graph would have been a predictable application of known graph architecture principles to organize the condition-dependent attributes that Goldberg already associates with its investigation-to-condition connections.

285. A POSITA would have been motivated to incorporate Dang’s feature-node architecture into the Goldberg/Gnanasambandam combination because Dang’s feature nodes beneficially provide for organizing condition-dependent attributes on the path between investigation and condition nodes. In Goldberg’s knowledge graph, the condition-specific attributes like probability values, strength-of-source ratings,

and source citations are associated with the connections between investigation nodes and condition nodes through structured rule records. Ex. 1006, FIG. 1; 6:65–7:2. Dang’s feature nodes formalize this intermediate relationship as a dedicated node tier within the graph database itself, where each feature node aggregates the condition-dependent information that links factual-tier nodes (including test nodes) to disease nodes through quantified conditional probabilities. Ex. 1008 at ¶¶ [0028]–[0031]; FIG. 1B.

286. A POSITA would have recognized that implementing Goldberg’s condition-dependent attributes as dedicated intermediate nodes yields concrete benefits for graph-based diagnostic systems: it enables the diagnostic algorithm to traverse through the intermediate data structure as a discrete step in the reasoning process, as Dang’s diagnostic method demonstrates at steps 504–508. Ex. 1008 at ¶¶ [0047]–[0050]; FIG. 5. It also allows the system to reuse the same feature node across multiple disease contexts, improving the modularity and maintainability of the knowledge graph. The result would have been predictable because both Goldberg and Dang already associate condition-dependent probability attributes with the connections between clinical observations and diseases; Dang’s feature-node tier merely formalizes this association as a structurally distinct intermediate node type within the same class of knowledge graph that Goldberg builds.

13.Dependent Claim 12

287. Dependent claim 12 recites: “A method as recited in claim 1 wherein at least one of said investigation nodes is traversable to at least two of said medical condition nodes and wherein said at least one of said investigation nodes does not include information specific to medical conditions associated with said at least two of said medical condition nodes.” I incorporate by reference my discussion of independent claim 1 above. *Supra* § X.A.2.

288. Claim 12 combines two requirements: (1) at least one investigation node must be traversable to at least two different medical condition nodes, and (2) that investigation node must not include information specific to the medical conditions associated with those condition nodes. Both requirements are satisfied by Goldberg’s knowledge graph for the reasons already established in connection with element [1d] and claim 8.

289. As to the first requirement, Goldberg’s unified hybrid knowledge graph includes a set of laboratory test nodes linked through connectors to a plurality of diseases. This is confirmed by the claims. Goldberg’s claim 9 recites “a fifth set of nodes representing the plurality of laboratory tests” as part of a medical knowledge graph that also includes disease nodes, symptom nodes, complication nodes, and treatment nodes. Ex. 1006 at 16:8–24. Goldberg’s claim 11 recites that the system determines “a specific laboratory test in the plurality of laboratory tests to be

administered to the patient” from pathways connecting test nodes to disease nodes. Ex. 1006 at 16:60–17:2. A POSITA reviewing Goldberg’s disclosures would understand that the test node is traversable to multiple disease nodes for the diagnostic algorithm to select the same laboratory test from pathways associated with different suspected diseases.

290. FIG. 4 of Goldberg illustrates this multi-condition architecture concretely. FIG. 4 is “a sample of a Medical Features Data Retrieval and Comparison, illustrating a feature by feature comparison for two selected diseases (CROHN’s and CELIAC).” Ex. 1006, FIG. 4; 9:1–5. The table shows clinical features—including symptoms, complications, associated conditions, and laboratory tests—alongside their likelihood values for both Celiac disease and Crohn’s disease. Multiple features appear with likelihood values for both diseases, demonstrating that the same clinical feature is traversable to at least two medical condition nodes. For example, “Chronic diarrhea (disorder)” has a likelihood of 10–20 for Celiac disease and 5–10 for Crohn’s disease; “Iron deficiency anemia (disorder)” has a likelihood of 5–10 for Celiac disease and 5–15 for Crohn’s disease; and “Failure to thrive (disorder)” has a likelihood of 60–70 for Celiac disease and 30–40 for Crohn’s disease. Ex. 1006, FIG. 4. Although FIG. 4’s laboratory-test entries are populated only for Celiac disease, Goldberg expressly describes retrieving “all the diseases that share a combination of one or more lab test results,” and claims pathways from the

“fifth set of nodes [laboratory tests]” used to select a test. Ex. 1006 at 8:54–58; 16:8–24; 16:55–17:5.

291. Goldberg further confirms that investigation nodes are traversable to multiple conditions. Goldberg discloses that “multiple pathways may be searched for when looking up Pneumonia” and “[t]he system may subsequently show all associations with findings, lab tests, etc.” Ex. 1006 at 11:55–67. This runtime traversal across the hybrid knowledge graph demonstrates that the same laboratory test node may be reached from multiple disease nodes through the graph’s pathway mechanism.

292. As to the second requirement, the investigation nodes in Goldberg’s knowledge graph do not include information specific to any medical condition, for the reasons discussed in connection with claim 8. Goldberg’s laboratory test nodes are derived from LOINC—a condition-independent ontology that defines tests by what they measure, not by any particular disease. Ex. 1006 at 6:17–28. The LOINC-coded test identity (e.g., “Amylase [Enzymatic activity/volume] in Serum, Plasma or Blood”) contains no information specific to Celiac disease, Crohn’s disease, or any other condition. The condition-specific information (the probability values and result qualifiers) resides in the connections between the test nodes and the condition nodes, not in the test nodes themselves. Ex. 1006, FIG. 1; *supra* § X.A.9.

293. The entire design principle of Goldberg’s knowledge graph is that condition-independent clinical concepts (including laboratory tests) are linked through typed, probabilistic connections to multiple diseases, enabling the diagnostic algorithm to traverse the graph and generate differential diagnoses across conditions. This is precisely what claim 12 requires.

294. To the extent Goldberg is found not to disclose this element, Dang does. Dang’s disease nodes connect not only to feature nodes but to “one or more other disease nodes,” with $P(dn1|dn2)$ indicating “the likelihood of the disease $dn1$ occurring when the disease $dn2$ is observed.” The depicted connections include “ $P(f4|d2)$, $P(d3|d2)$, $P(f6|d2)$... $P(dn|d3)$.” Ex. 1008 at ¶ [0031].

295. Dang’s test nodes carry no disease-specific information for the reasons given for claim 8. *Supra* § X.A.9. A POSITA would have been motivated to use Dang’s layout for the same reason Goldberg shares a single laboratory-test node across disease pathways: one investigation node can inform the diagnosis of several diseases.

14. Dependent Claim 13

296. Dependent claim 13 recites: “A method as recited in claim 1 wherein said knowledge graph further includes on a path between each medical condition node and each of said investigation nodes to which said each medical condition node is traversable, a ranking value indicating how likely it is that an investigation

represented by said each investigation node is indicative of a medical condition represented by said each medical condition node.” I incorporate by reference my discussion of independent claim 1 above. *Supra* § X.A.2.

297. Goldberg expressly discloses this element. Goldberg’s FIG. 1 is a table depicting “a sample of knowledge base rules for Celiac disease.” Ex. 1006, FIG. 1; 6:65–7:2. Each rule entry in FIG. 1 includes a “Probability” column containing a numerical value that indicates the likelihood that the associated feature is indicative of the disease. For example, the LABORATORY TESTS entry for “Amylase [Enzymatic activity/volume] in Serum, Plasma or Blood: HIGH” has a probability value of 5–10, indicating the likelihood that this laboratory test result is indicative of Celiac disease. Ex. 1006, FIG. 1. This probability value is on the path between the condition node (Celiac disease) and the investigation node (the Amylase test) and expressly indicates how likely it is that the investigation is indicative of the medical condition—which is precisely what claim 13 requires.

298. FIG. 4 of Goldberg confirms that these ranking values are present across multiple conditions. FIG. 4 is “a sample of a Medical Features Data Retrieval and Comparison, illustrating a feature by feature comparison for two selected diseases (CROHN’s and CELIAC).” Ex. 1006, FIG. 4; 9:1–5. The table shows clinical features alongside their “Likelihood” values for both Celiac disease and Crohn’s disease. For example, “Chronic diarrhea (disorder)” has a likelihood of 10–

20 for Celiac disease and 5–10 for Crohn’s disease; “Iron deficiency anemia (disorder)” has a likelihood of 5–10 for Celiac disease and 5–15 for Crohn’s disease. Ex. 1006, FIG. 4. Each of these likelihood values is a ranking value on the path between the condition node and the associated clinical feature, indicating how likely the feature is indicative of the respective condition.

299. FIG. 11 of Goldberg further illustrates the hybrid model with rules describing “probabilistic relationships between medical concepts as they appear in medical texts.” Ex. 1006, FIG. 11; 10:28–31. Each rule connecting a medical concept to a qualified feature carries a rate (R) and is typed (e.g., complication, symptom, lab, risk factor). These rates are ranking values on the paths between investigation-related nodes and condition nodes, indicating how likely the investigation or finding is associated with the condition.

300. Goldberg describes the knowledge base as holding “a probabilistic association between various disease and disorder features” in which each feature’s association with a disease is described by “a number of such associations indicating the probability which links each feature with the disease.” Ex. 1006 at 5:67–6:7. The diagnostic algorithm uses these probabilistic values directly: it “uses the inputs and their corresponding associations within the knowledge base to run a mathematical inference algorithm that finds the most likely explanation for the patient’s condition,” taking into account “the strength of the various associations as indicated in the

knowledge base.” Ex. 1006 at 7:32–34; 8:16–19. The probability values in FIG. 1 and the likelihood values in FIG. 4 are the ranking values that drive this inference and they indicate, for each condition-investigation pair, how likely the investigation is indicative of the condition.

301. Goldberg’s claims further confirm the presence of ranking values on the paths between conditions and investigations. Goldberg’s claim 10 recites that “a probabilistic value associated with at least one of the plurality of one-sided arrows is displayed.” Ex. 1006 at 16:57–59. Goldberg’s claim 13 recites connectors comprising “arrows denoting probabilistic linkages between nodes in the plurality of nodes” and determining “an aggregated connection strength for each of the one or more pathways.” Ex. 1006 at 17:10–18:10. These are express ranking values on paths between nodes in the knowledge graph, including the paths between condition nodes and investigation nodes.

302. Dang provides additional support. Each disease-to-feature connection in Dang’s graph “may have a corresponding conditional probability P ” indicating “the likelihood of the feature f_x occurring when the disease d_n [is] observed.” Ex. 1008 at ¶ [0031]; FIG. 1B. That probability sits on the path between a test node and a disease node (*supra* § X.A.13), and it indicates how likely the associated finding is to occur for that disease.

15. Dependent Claim 14

303. Dependent claim 14 recites: “A method as recited in claim 1 wherein said knowledge graph further includes on a path between each medical condition node and each of said investigation nodes to which said each medical condition node is traversable, at least one condition-dependent attribute of a medical condition represented by said each medical condition node.” I incorporate by reference my discussion of independent claim 1 above. *Supra* § X.A.2.

304. Claim 14 thus requires that the knowledge graph includes, on the path between each condition node and each investigation node to which the condition node is traversable, at least one attribute that is dependent on the medical condition represented by that condition node. Claim 14 is broader than claim 13: while claim 13 requires a specific type of condition-dependent attribute (a ranking value indicating likelihood), claim 14 requires only that at least one condition-dependent attribute exists on that path.

305. Goldberg discloses multiple condition-dependent attributes on the path between condition nodes and investigation nodes. As discussed above in connection with claim 13, Goldberg’s FIG. 1 is a table depicting knowledge base rules for Celiac disease, in which each rule entry includes multiple fields that are specific to the association between a particular investigation and a particular condition. Ex. 1006, FIG. 1; 6:65–7:2. For the LABORATORY TESTS entry linking the Amylase test to

Celiac disease, these condition-dependent attributes include: (1) the Probability value of 5–10, indicating the likelihood that the test result is indicative of Celiac disease specifically; (2) the Strength of source value, indicating the reliability of the clinical evidence supporting the association between this test and Celiac disease; and (3) the Link to source, providing the citation to the medical literature underlying this condition-specific association. Ex. 1006, FIG. 1. Each of these attributes is condition-dependent because it pertains to the relationship between the investigation and a specific medical condition, and each may vary across different conditions.

306. FIG. 4 of Goldberg confirms that these attributes are condition-dependent by showing that they vary across conditions. FIG. 4 compares clinical features for Celiac disease and Crohn’s disease, displaying different likelihood values for the same features across the two conditions. Ex. 1006, FIG. 4; 9:1–5. For example, “Chronic diarrhea (disorder)” has a likelihood of 10–20 for Celiac disease and 5–10 for Crohn’s disease; “Iron deficiency anemia (disorder)” has a likelihood of 5–10 for Celiac disease and 5–15 for Crohn’s disease. Ex. 1006, FIG. 4. The fact that the same clinical feature carries different probability values for different conditions demonstrates that these probability values are condition-dependent attributes—they are specific to the medical condition to which the investigation or feature is linked.

307. Goldberg further confirms that its knowledge base stores condition-dependent attributes on the paths between conditions and clinical features. Goldberg describes the knowledge base as holding “a probabilistic association between various disease and disorder features” in which “a chapter in a medical textbook or an article describing a certain disease may be translated into a number of such associations indicating the probability which links each feature with the disease.” Ex. 1006 at 5:67–6:7. The specification also describes that each association within the knowledge base “has a source such as textbook, article etc.” and that “[t]he knowledge base will contain the source along ranking to indicate its relative strength.” Ex. 1006 at 6:35–41. These source-strength rankings and probabilistic associations are condition-dependent attributes on the path between condition nodes and investigation nodes.

308. FIG. 11 of Goldberg further illustrates that the hybrid model includes condition-dependent attributes on the paths between medical concepts. FIG. 11 shows rules describing “probabilistic relationships between medical concepts as they appear in medical texts,” with each rule carrying a rate (R) and being typed by the nature of the clinical relation (e.g., complication, symptom, lab, risk factor). Ex. 1006, FIG. 11; 10:28–31. The rate and the relation type are condition-dependent attributes: the rate varies depending on which condition the feature is associated

with, and the relation type characterizes the nature of the feature's association with the specific condition.

309. Dang's $P(fx|dn)$ values are also condition-dependent attributes: each is defined "for a pair of disease and feature (dn, fx)," so the same feature carries a separate value for each disease to which it connects. Ex. 1008 at ¶ [0031].

16. Dependent Claim 17

310. Dependent claim 17 recites: "A method as recited in claim 1 wherein said software application causes said output to be displayed on a computer to a user." I incorporate by reference my discussion of independent claim 1 above. *Supra* § X.A.2.

311. Goldberg discloses this element. Goldberg's platform includes "a user interface to enable an end user to view results output by the data processing engine." Ex. 1006, Abstract. The Patient Management App 230 is "designed to assist physicians in managing a patient through the diagnosis process" and "offers physicians a way to make an updated, evidence-based clinical decision related to the evaluation and management of a patient throughout the diagnosis process." Ex. 1006 at 7:19–24. The application generates and displays a list of suspected diseases and complications, along with recommended next-best steps, on the user interface for viewing by the clinician. Ex. 1006 at 7:27–32.

312. FIG. 2 of Goldberg illustrates this architecture: the inference engine 210 is communicatively coupled to the knowledge base 220, and the Patient Management App 230 provides the user interface through which results are displayed to the end user. Ex. 1006, FIG. 2. The Patient Management App 230 displays the differential diagnoses and next-step suggestions returned by inference engine 210. *Id.* at 7:25–42. The app can also show a “path to decision,” i.e., “the clinical reasoning path of associations which were used in order to make a certain diagnosis or recommend a next step including details or links to original medical textbook or academic paper sources,” and “may also display a ‘strength score’ illustrating the sources’ reliability.” *Id.* at 8:41–47. Thus, Goldberg’s software application causes the output of the diagnostic algorithm to be displayed on a computer to a user.

313. Gnanasambandam independently discloses this element. Gnanasambandam’s autonomous multipurpose application presents user interfaces on the display of user device 104, including a home user interface 3300 summarizing the user’s “Chronic conditions” and “Acute issues.” Ex. 1007 at 63:55–64:10; FIG. 33. When the platform identifies a possible medical condition, “[a] user interface 2800 of the computing device of the patient may present the first information ‘Possible medical condition: Type 2 Diabetes Mellitus’ at the first time” *Id.* at 59:60–60:25; FIG. 28. Likewise, the processing device “may provide the diagnosis to a computing device of a patient and/or a physician for presentation on the computing

device.” *Id.* at 62:1–30; FIG. 31 (block 3108). For medical personnel, the “clinic viewer may display a condition with which a patient has been diagnosed” and “may display a care plan for the patient.” *Id.* at 16:55–17:25. Thus, Gnanasambandam’s software application causes the output to be displayed on a computer to a user.

314. Both Goldberg and Gnanasambandam independently satisfy the display limitation of claim 17. In the combination, Goldberg’s diagnostic outputs, including ranked conditions, next-best actions, and associated clinical reasoning, would be delivered through Gnanasambandam’s API architecture and displayed on computing devices to end users through the user interfaces that both references expressly disclose.

17.Dependent Claim 18

315. Dependent claim 18 recites: “A method as recited in claim 1 wherein said output is output to said software application in a data structure suitable for further processing.” I incorporate by reference my discussion of independent claim 1 above. *Supra* § X.A.2.

316. Gnanasambandam discloses this element. Gnanasambandam’s platform is built around an API architecture in which the cognitive intelligence platform publishes APIs that deliver structured analytics results to consuming software applications. Ex. 1007 at 100:24–34; 101:12–37. The autonomous multipurpose application subscribes to these APIs and receives output that it further

processes. For example, the diagnosis is provided “to a computing device of a patient and/or a physician for presentation on the computing device.” *Id.* at 62:1–30. The platform compares a patient graph to the knowledge graph for the condition “to generate a care plan.” *Id.* at 15:1–35. The cognitive agent “may determine whether any parameter is missing respective data in the data structure” and, if so, “may cause a notification to be presented to the user on user interface.” *Id.* at 97:38–98:40; FIG. 87. An API response is a data structure suitable for further processing. APIs exist precisely to enable software applications to receive and programmatically act upon structured data. Gnanasambandam’s API architecture thus satisfies the requirement of claim 18.

317. Goldberg independently supports this element. Goldberg’s inference engine 210 returns differential diagnoses and next-step suggestions to the Patient Management App 230. Ex. 1006, FIG. 2; 7:25–42. The application can “rank likely differential diagnosis results by their probabilistic score” and “recommend test A over test B.” *Id.* at 8:10–32. Output that the application ranks and acts upon is output in a data structure suitable for further processing.

318. In the combination, Gnanasambandam’s API architecture formalizes this structured data delivery. Goldberg’s diagnostic outputs—already structured for consumption by the Patient Management App—would be delivered through Gnanasambandam’s published APIs as structured API responses to consuming

software applications, satisfying the requirement that the output is in a data structure suitable for further processing.

18. Dependent Claim 19

319. Dependent claim 19 recites: “A method as recited in claim 1 further comprising: in response to said output, receiving at said knowledge graph a new patient diagnostic factor known to be present or absent in said patient; and outputting, from said knowledge graph and responsive to said new patient diagnostic factor, a new output.” I incorporate by reference my discussion of independent claim 1 above. *Supra* § X.A.2.

320. Goldberg discloses this iterative diagnostic loop. Goldberg’s application generates “a list of suspected diseases and complications” and then “suggests the next best step to narrow down that list by ruling in some diseases or ruling out other high-risk diseases.” Ex. 1006 at 7:25–35. The next step may be a recommended test, because “[a] positive or negative answer for test A would increase the probability of suspected diagnosis.” *Id.* at 8:20–32. Goldberg describes “a constantly updated differential diagnosis list along with patient-tailored best next step recommendation,” such that “the app essentially creates on the fly a protocol for managing the diagnosis process for each patient.” *Id.* at 7:29–42. A POSITA would have understood that a differential list that is “constantly updated” in response to recommended tests is updated with the results of those tests. That is the loop claim

19 requires: an output (the differential and next step), new patient information obtained in response (a finding now known to be present or absent), and a new output.

321. Dang confirms this iterative design. Dang generates “a best next inquiry question, for use in a medical diagnosis process,” and reports that its method reached a diagnosis with an average of 13.5 questions, compared with 27.7 and 16.6 for benchmark methods. Ex. 1008 at ¶¶ [0003], [0059]. A method measured by the number of questions needed to reach a diagnosis necessarily receives the answer to each question and produces a new output before asking the next. A POSITA would have been motivated to implement Goldberg’s updating loop using Dang’s next-inquiry method for the reasons discussed above. *Supra* § X.A.1; § X.A.8.

322. The “receiving at said knowledge graph” language of claim 19 is satisfied because Goldberg’s application runs its inference over the knowledge base using “the inputs and their corresponding associations within the knowledge base,” so a newly entered finding is processed against the same knowledge graph to produce the updated differential. Ex. 1006 at 8:14–24.

323. Gnanasambandam independently supports this iterative loop. Gnanasambandam’s autonomous multipurpose application “may be implemented as one or more application programming interfaces (API).” Ex. 1007 at 11:20–35. In response to a user’s query, the platform “performs a first round of analysis” and

“determines a first set of follow up questions,” which are “dynamically formulated” “to elicit a response that will inform or fulfill an identified parameter.” *Id.* at 30:15–31:60. “When answers are received from the user in response to a dynamically formulated question, the cognitive intelligence platform 102 inserts the answer into the workspace,” and “[f]or a complex originating question, the cognitive intelligence platform 102 can go through several rounds of analysis.” *Id.* at 31:46–32:60; *see also id.* at 32:60–33:30 (FIG. 4 example posing successive questions and answering after receiving the responses). FIG. 7A likewise shows that if sufficient data is not present, “the method proceeds to block 712 and the method performs another round of analysis.” *Id.* at 35:55–36:20; FIG. 7A. Gnanasambandam thus discloses receiving new patient information in response to an output and producing a new output.

19. Dependent Claim 20

324. Dependent claim 20 recites: “A method as recited in claim 1 further comprising: matching a portion of said medical query to at least one of said medical condition nodes or to at least one of said finding nodes to produce said output.” I incorporate by reference my discussion of independent claim 1 above. *Supra* § X.A.2.

325. Goldberg expressly discloses this element. Goldberg’s diagnostic algorithm receives patient condition data including “history, symptoms, signs and other clinical features” and “uses the inputs and their corresponding associations

within the knowledge base to run a mathematical inference algorithm that finds the most likely explanation for the patient's condition.” Ex. 1006 at 7:25–40; 8:15–25. The system matches the patient's clinical data to finding nodes (symptom nodes, sign nodes) and condition nodes (disease nodes) in the knowledge graph to generate differential diagnoses. This matching operation produces the output: “a list of suspected diseases and complications” along with “the next best step to narrow down that list.” Ex. 1006 at 7:27–32. A POSITA would understand that this matching process occurs to generate differential diagnoses.

326. Gnanasambandam independently supports this element. Gnanasambandam's platform identifies possible medical conditions “by identifying a similarity between the indicia and a knowledge graph representing knowledge pertaining to the possible medical condition.” Ex. 1007 at 57:23–27. This similarity identification is a matching operation between a portion of the medical query (the extracted indicia) and the nodes in the knowledge graph.

20. Dependent Claim 21

327. Dependent claim 21 recites: “A method as recited in claim 1 wherein said knowledge graph includes more than twenty thousand nodes.” I incorporate by reference my discussion of independent claim 1 above. *Supra* § X.A.2.

328. Goldberg discloses this element. Goldberg's hybrid knowledge graph integrates SNOMED, which Goldberg describes as “an encompassing Ontology

covering hundreds of thousands of medical terms arranged as an ontology.” Ex. 1006 at 1:54–2:1. Hundreds of thousands of medical terms is well in excess of twenty thousand nodes. Additionally, Goldberg integrates LOINC and clinical probability data linking concepts across these ontologies. Ex. 1006 at 6:17–28. A POSITA would understand, and would have found obvious, that the resulting hybrid knowledge graph, which combines SNOMED, LOINC, and clinical rules, contains more than twenty thousand nodes.

329. Additionally, the number of nodes recited in claim 21 would have been obvious to a POSITA as a matter of routine scaling. The only difference between a knowledge graph with fewer than twenty thousand nodes and one with more than twenty thousand nodes is the number of medical conditions, findings, and investigations represented, i.e., additional instances of the same types of nodes that Goldberg already discloses. Goldberg’s architecture is expressly designed to accommodate an expanding body of medical knowledge: Goldberg’s Medical knowledgebase 310 “is curated and maintained by a community of medical specialists,” because “[t]he amount of available medical knowledge and its growth rate require a coordinated effort by a large group of contributors,” and “[t]he Knowledgebase 310 may be updated on an ongoing basis.” Ex. 1006 at 7:48–8:5; FIG. 3; *see also id.* at 3:28–32 (data curation modules including “a crowdsourcing module, a community-based editing module”); 11:1–20 (“as crowdsourcing users

add data into the knowledge graph”). Adding more conditions, each with their associated finding nodes and investigation nodes, is merely a duplication of the same components that Goldberg already discloses.

330. Here, scaling the knowledge graph from any number of nodes to more than twenty thousand nodes would produce no new or unexpected result—the system would operate in exactly the same manner, applying the same diagnostic algorithm, the same probabilistic associations, and the same graph-traversal techniques to a larger set of the same types of nodes. A POSITA would have understood that a clinically useful knowledge graph covering a meaningful range of medical conditions would naturally exceed twenty thousand nodes, given that SNOMED alone contains hundreds of thousands of terms and Goldberg expressly integrates SNOMED into its hybrid knowledge graph. Ex. 1006 at 1:54–2:1; 6:5–28.

21.Dependent Claim 22

331. Dependent claim 22 recites: “A method as recited in claim 1 wherein said knowledge graph includes more than one hundred thousand nodes.” I incorporate by reference my discussion of independent claim 1 above. *Supra* § X.A.2.

332. Goldberg discloses this element. Goldberg’s hybrid knowledge graph integrates SNOMED, which Goldberg describes as “an encompassing Ontology covering hundreds of thousands of medical terms arranged as an ontology.” Ex. 1006

at 1:54–2:1. Hundreds of thousands of medical terms would result in well more than one hundred thousand nodes. Additionally, Goldberg integrates LOINC, which is “an ontology dedicated to laboratory tests and procedures” (Ex. 1006 at 6:17–20), and clinical probability data linking concepts across these ontologies. A POSITA would understand, and would have found obvious, that the resulting hybrid knowledge graph, which combines SNOMED, LOINC, and clinical rules, contains more than one hundred thousand nodes.

333. Additionally, the number of nodes recited in claim 22 would have been obvious to a POSITA as a matter of routine scaling. The only difference between a knowledge graph with fewer than one hundred thousand nodes and one with more than one hundred thousand nodes is the number of medical conditions, findings, and investigations represented, i.e., additional instances of the same types of nodes that Goldberg already discloses. Goldberg’s architecture is expressly designed to accommodate an expanding body of medical knowledge: Goldberg’s Medical knowledgebase 310 can be, for example, “curated and maintained by a community of medical specialists,” because “[t]he amount of available medical knowledge and its growth rate require a coordinated effort by a large group of contributors,” and “[t]he Knowledgebase 310 may be updated on an ongoing basis.” Ex. 1006 at 7:48–8:5; FIG. 3; *see also id.* at 3:28–32 (data curation modules including “a crowdsourcing module, a community-based editing module”); 11:1–20 (“as

crowdsourcing users add data into the knowledge graph”). Adding more conditions, each with their associated finding nodes and investigation nodes, is merely a duplication of the same components that Goldberg already discloses.

334. Scaling the knowledge graph from any number of nodes to more than one hundred thousand nodes would produce no new or unexpected result—the system would operate in exactly the same manner, applying the same diagnostic algorithm, the same probabilistic associations, and the same graph-traversal techniques to a larger set of the same types of nodes. A POSITA would have understood that a clinically useful knowledge graph covering a meaningful range of medical conditions would naturally exceed one hundred thousand nodes, given that SNOMED alone contains hundreds of thousands of terms and Goldberg expressly integrates SNOMED into its hybrid knowledge graph. Ex. 1006 at 1:54–2:1; 6:5–28.

22.Independent Claim 23

335. Each element of independent claim 23 of the ’883 Patent is disclosed or rendered obvious by the combination of Goldberg and Gnanasambandam. I address each element below.

a. Element [23pre]

336. Element [23pre] recites: “A system for providing reasoning capability to integrated medical evidence from medical evidence sources, the system comprising.”

337. To the extent the preamble is limiting, the combination of Goldberg and Gnanasambandam discloses a system for providing reasoning capability to integrated medical evidence from medical evidence sources, for the same reasons discussed above with respect to element [1pre]. *Supra* § X.A.2.a. Goldberg’s platform includes a healthcare knowledge graph database, a diagnostic algorithm, and user interfaces that collectively provide reasoning capability. Ex. 1006, Abstract; 3:5–18; 7:25–50. Gnanasambandam’s API platform publishes APIs that perform analytics on medical knowledge graphs. Ex. 1007 at 1:32–42; 100:24–34.

b. Element [23a]

338. Element [23a] recites: “a database that stores a knowledge graph including a plurality of nodes, which include a plurality of medical condition nodes, a plurality of finding nodes, and a plurality of investigation nodes.”

339. Element [23a] requires a database that stores a knowledge graph with the claimed node types. This element corresponds to element [1b] with the additional requirement of a database. Goldberg discloses “a database for storage of medical hybrid knowledge graph.” Ex. 1006 at 12:45–50. The knowledge graph stored in this database includes the six node sets discussed above with respect to element [1b]: diseases, associated conditions, symptoms, complications, laboratory tests, and treatments. Ex. 1006 at 16:8–24. Thus, Goldberg discloses element [23a] for the same reasons discussed above for element [1b], with the additional database

limitation met by Goldberg's express disclosure of a database for storing the hybrid knowledge graph. *Supra* § X.A.2.c.

c. Element [23b]

340. Element [23b] recites: “wherein at least one of said finding nodes in said knowledge graph is traversable to a first medical condition node to indicate that a finding represented by said at least one of said finding nodes is associated with a first medical condition represented by said first medical condition node.”

341. Element [23b] is the same as element [1c]. Thus, Goldberg discloses element [23b] for the same reasons as element [1c]. *Supra* § X.A.2.d.

d. Element [23c]

342. Element [23c] recites: “wherein at least one of said investigation nodes in said knowledge graph is traversable to a second medical condition node to indicate that an investigation represented by said at least one of said investigation nodes is associated with a second medical condition represented by said second medical condition node.”

343. Element [23c] is the same as element [1d]. Thus, Goldberg discloses element [23c] for the same reasons as element [1d]. *Supra* § X.A.2.e.

e. Element [23d]

344. Element [23d] recites: “wherein one or more of said nodes in said knowledge graph are extracted from a plurality of medical evidence sources and are linked to a medical terminology ontology.”

345. Element [23d] is the same as element [1e]. Thus, Goldberg discloses element [23d] for the same reasons as element [1e]. *Supra* § X.A.2.f.

f. Element [23e]

346. Element [23e] recites: “wherein said knowledge graph is configured to support reasoning capability by allowing traversal of ancestor-descendant relationships among one or more of said medical condition nodes, said finding nodes, or said investigation nodes, and.”

347. Element [23e] is the same as element [1f]. Thus, Goldberg discloses element [23e] for the same reasons as element [1f]. *Supra* § X.A.2.g.

g. Element [23f]

348. Element [23f] recites: “wherein said medical terminology ontology includes a Systemized Nomenclature of Medicine—Clinical Terms (SNOMED-CT) database, or a Logical Observation Identifiers, Names, and Codes (LOINC) database.”

349. Element [23f] is the same as element [1g]. Thus, Goldberg discloses element [23f] for the same reasons as element [1g]. *Supra* § X.A.2.h.

h. Element [23g]

350. Element [23g] recites: “a computer system comprising one or more processors and memory that stores instructions, wherein the instructions, when executed by the one or more processors, cause the one or more processors to.”

351. Goldberg discloses one or more processors. Each of Goldberg’s independent claims recites steps performed “by at least one processor.” Claim 1, for example, recites “diagnosing, by at least one processor, a patient with a disease,” including “defining, by the at least one processor, one or more pathways on the medical knowledge graph,” “receiving, by the at least one process[or], the health data,” and “generating, by the at least one processor, a statistical correlation between each suspected disease . . . and the patient health data.” Ex. 1006 at 14:28–67 (claim 1); *see also id.*, claims 9, 13. Goldberg’s specification describes “a healthcare diagnostics and treatment platform . . . comprising a computerized model that runs on a database for storage of medical hybrid knowledge graph, a diagnostic algorithm or software which gets inputs from the hybrid knowledge graph and from given sets of patient findings or test results data.” *Id.* at 12:43–55. The platform includes “a cloud-based inference engine 210, communicatively coupled to a knowledge base 220.” *Id.* at 7:10–20; FIG. 2. Its modules “may be implemented using one or more front-end and server-side techniques, software environments and development languages.” *Id.* at 8:6–13.

352. To the extent Goldberg does not expressly say that its system includes a memory, a POSITA would have readily understood that Goldberg includes memory that stores the instructions its processor executes, and at a minimum would have found this obvious. First, Goldberg describes its application as software that “runs a file with instructions to execute commands to manage a patient through a diagnosis process.” Ex. 1006 at 2:61–65. A processor executes instructions by retrieving them from memory. Second, Goldberg’s “database for storage of medical hybrid knowledge graph” and its “diagnostic algorithm or software” both require storage from which the processor reads data and code. *Id.* at 12:43–55. Computers used for such applications always included a processor and memory storing the programs the processor executes. Implementing Goldberg’s processor-executed software with instructions stored in memory would therefore have been the ordinary, predictable way to build Goldberg’s system, not a design choice requiring any inventive step.

353. Gnanasambandam expressly discloses this conventional architecture. Its computing device 1400 “can be used to implement the various components described herein” including “the several computing devices implementing the cognitive intelligence platform 102.” Ex. 1007 at 40:13–21. Device 1400 includes “a processor 1402 that represents a microprocessor or controller for controlling the overall operation of the computing device 1400,” and “a Random-Access Memory

(RAM) 1420” that “stores instructions related to the operation of processes and applications executing on the computing device.” *Id.* at 40:20–27, 40:49–58; FIG. 14; *see also id.*, claim 15.

354. Dang confirms the same. Its computer system/server 12 includes “one or more processors or processing units 16, a system memory 28,” and program modules stored in memory 28 “carry out the functions and/or methodologies of embodiments of the invention.” Ex. 1008 at ¶¶ [0063], [0066]–[0067].

355. In the combined system, the functions recited in elements [23h]–[23j] would be carried out by instructions stored in memory and executed by the system’s processors.

i. Element [23h]

356. Element [23h] recites: “receive, via an application programming interface (API), a medical query from a software application for querying said knowledge graph.”

357. Element [23h] corresponds to element [1a] and requires that the computer system receive, via an API, a medical query from a software application. Gnanasambandam and Goldberg disclose this element for the same reasons as element [1a]. *Supra* § X.A.2.b.

j. Element [23i]

358. Element [23i] recites: “identify, using said reasoning capability, a set of ancestor-descendant relationships of said knowledge graph that are relevant to said medical query, wherein said reasoning capability including traversing said knowledge graph to identify one or more of said medical condition nodes, said finding nodes, or said investigation nodes that are relevant to said medical query; and.”

359. Element [23i] corresponds to element [1h] and requires that the computer system identify, using the reasoning capability, ancestor-descendant relationships relevant to the medical query. Goldberg and Gnanasambandam disclose this element for the same reasons as element [1h]. *Supra* § X.A.2.i.

k. Element [23j]

360. Element [23j] recites: “output, responsive to said medical query, an output to said software application, wherein said output includes data related to one or more of said medical condition nodes, said finding nodes, or said investigation nodes that are relevant to said medical query in order for said software application to generate a medical determination.”

361. Element [23j] corresponds to element [1i] and requires that the computer system output data to the software application for generating a medical determination. Gnanasambandam and Goldberg disclose this element for the same reasons as element [1i]. *Supra* § X.A.2.j.

23. Dependent Claim 24

362. Dependent claim 24 recites: “A system as recited in claim 23, wherein each of said investigation nodes is linked to a corresponding result node holding a result of said each investigation node, each of said corresponding result nodes being linked to said at least one of said medical condition nodes to which said each investigation node is linked.” I incorporate by reference my discussion of independent claim 23 above. *Supra* § X.A.22.

363. Dependent claim 24 contains the same additional limitation as dependent claim 2. Thus, my analysis for dependent claim 2 applies equally to the additional limitation in claim 24, and I incorporate by reference that discussion herein. *Supra* § X.A.3.

24. Dependent Claim 25

364. Dependent claim 25 recites: “A system as recited in claim 23, wherein said finding nodes are diagnostic factor nodes.” I incorporate by reference my discussion of independent claim 23 above. *Supra* § X.A.22.

365. Dependent claim 25 contains the same additional limitation as dependent claim 4. Thus, my analysis for dependent claim 4 applies equally to the additional limitation in claim 25, and I incorporate by reference that discussion herein. *Supra* § X.A.5.

25. Dependent Claim 26

366. Dependent claim 26 recites: “A system as recited in claim 23, wherein said each investigation node does not include information specific to a medical condition associated with one of said medical condition nodes to which said each investigation node is traversable.” I incorporate by reference my discussion of independent claim 23 above. *Supra* § X.A.22.

367. Dependent claim 26 contains the same additional limitation as dependent claim 8. Thus, my analysis for dependent claim 8 applies equally to the additional limitation in claim 26, and I incorporate by reference that discussion herein. *Supra* § X.A.9.

26. Dependent Claim 27

368. Dependent claim 27 recites: “A system as recited in claim 23, wherein each of said investigation nodes is disambiguated across said plurality of medical condition nodes.” I incorporate by reference my discussion of independent claim 23 above. *Supra* § X.A.22.

369. Dependent claim 27 contains the same additional limitation as dependent claim 9. Thus, my analysis for dependent claim 9 applies equally to the additional limitation in claim 27, and I incorporate by reference that discussion herein. *Supra* § X.A.10.

27. Dependent Claim 28

370. Dependent claim 28 recites: “A system as recited in claim 23, wherein each of said investigation nodes is independent of said at least one of said medical condition nodes to which said each investigation node is traversable.” I incorporate by reference my discussion of independent claim 23 above. *Supra* § X.A.22.

371. Dependent claim 28 contains the same additional limitation as dependent claim 10. Thus, my analysis for dependent claim 10 applies equally to the additional limitation in claim 28, and I incorporate by reference that discussion herein. *Supra* § X.A.11.

28. Dependent Claim 29

372. Dependent claim 29 recites: “A system as recited in claim 23, wherein said knowledge graph further includes for each of said investigation nodes, a data structure different from said each investigation node that is traversable to both said each investigation node and said at least one of said medical condition nodes to which said each investigation node is traversable, said data structure including attributes pertaining to said at least one of said medical condition nodes to which said each investigation node is traversable.” I incorporate by reference my discussion of independent claim 23 above. *Supra* § X.A.22.

373. Dependent claim 29 contains the same additional limitation as dependent claim 11. Thus, my analysis for dependent claim 11 applies equally to the

additional limitation in claim 29, and I incorporate by reference that discussion herein. *Supra* § X.A.12.

29.Dependent Claim 30

374. Dependent claim 30 recites: “A system as recited in claim 23, wherein at least one of said investigation nodes is traversable to at least two of said medical condition nodes and wherein said at least one of said investigation nodes does not include information specific to medical conditions associated with said at least two of said medical condition nodes.” I incorporate by reference my discussion of independent claim 23 above. *Supra* § X.A.22.

375. Dependent claim 30 contains the same additional limitation as dependent claim 12. Thus, my analysis for dependent claim 12 applies equally to the additional limitation in claim 30, and I incorporate by reference that discussion herein. *Supra* § X.A.13.

30.Dependent Claim 31

376. Dependent claim 31 recites: “A system as recited in claim 23, wherein said knowledge graph further includes on a path between each medical condition node and each of said investigation nodes to which said each medical condition node is traversable, a ranking value indicating how likely it is that an investigation represented by said each investigation node is indicative of a medical condition

represented by said each medical condition node.” I incorporate by reference my discussion of independent claim 23 above. *Supra* § X.A.22.

377. Dependent claim 31 contains the same additional limitation as dependent claim 13. Thus, my analysis for dependent claim 13 applies equally to the additional limitation in claim 31, and I incorporate by reference that discussion herein. *Supra* § X.A.14.

31.Dependent Claim 32

378. Dependent claim 32 recites: “A system as recited in claim 23, wherein said knowledge graph further includes on a path between each medical condition node and each of said investigation nodes to which said each medical condition node is traversable, at least one condition-dependent attribute of a medical condition represented by said each medical condition node.” I incorporate by reference my discussion of independent claim 23 above. *Supra* § X.A.22.

379. Dependent claim 32 contains the same additional limitation as dependent claim 14. Thus, my analysis for dependent claim 14 applies equally to the additional limitation in claim 32, and I incorporate by reference that discussion herein. *Supra* § X.A.15.

32.Dependent Claim 33

380. Dependent claim 33 recites: “A system as recited in claim 23 wherein said knowledge graph includes more than twenty thousand nodes.” I incorporate by reference my discussion of independent claim 23 above. *Supra* § X.A.22.

381. Dependent claim 33 contains the same additional limitation as dependent claim 21. Thus, my analysis for dependent claim 21 applies equally to the additional limitation in claim 33, and I incorporate by reference that discussion herein. *Supra* § X.A.20.

33. Dependent Claim 34

382. Dependent claim 34 recites: “A system as recited in claim 23 wherein said knowledge graph includes more than one hundred thousand nodes.” I incorporate by reference my discussion of independent claim 23 above. *Supra* § X.A.22.

383. Dependent claim 34 contains the same additional limitation as dependent claim 22. Thus, my analysis for dependent claim 22 applies equally to the additional limitation in claim 34, and I incorporate by reference that discussion herein. *Supra* § X.A.21.

B. Ground 2: Claims 15-16 are Obvious Over Goldberg, Gnanasambandam, Dang, and Datla

384. In my opinion, claims 15-16 of the '883 Patent are obvious over Goldberg, Gnanasambandam, Dang, and Datla.

1. Motivation to Combine

385. As discussed for Ground 1, the combination of Goldberg and Gnanasambandam provides a knowledge-graph-based diagnostic system accessible via published APIs. Goldberg supplies the hybrid knowledge graph integrating clinical data with SNOMED and LOINC ontologies, the diagnostic algorithm that traverses pathways to generate differential diagnoses, and the patient management application through which clinicians interact with the system. Ex. 1006, Abstract; 6:5–28; 7:25–50. Gnanasambandam supplies the API deployment architecture through which external software applications submit medical queries and receive structured diagnostic outputs. Ex. 1007 at 100:24–40; 101:12–37.

386. Claims 15 and 16 require the software application that queries the knowledge graph to be a natural language processing (NLP) engine that receives verbatim text and produces the medical query. Goldberg’s patient management application receives patient data through structured user interfaces, and Gnanasambandam’s autonomous multipurpose application subscribes to and invokes APIs, but neither reference expressly characterizes the querying application as an NLP engine that receives unstructured verbatim text and converts it into structured medical queries. Datla discloses this element.

387. Datla discloses a knowledge-graph-based clinical diagnosis system built around an NLP engine that receives natural language input describing a patient’s scenario and produces structured medical queries used to identify patient

conditions from a knowledge graph. Ex. 1005, Abstract; ¶¶ [0006]–[0008], [0028]; claim 1. Datla’s NLP engine occupies a single, non-overlapping functional layer—the natural-language front-end—that does not duplicate or interfere with the contributions of Goldberg or Gnanasambandam.

388. A POSITA would have been motivated to incorporate Datla’s NLP engine into the Goldberg/Gnanasambandam combination because Datla, like Goldberg and Gnanasambandam, is directed to healthcare knowledge-graph-based clinical decision support. Ex. 1005 at ¶¶ [0001], [0006]. Datla addresses the specific problem of accepting natural language input from a clinician describing a patient’s scenario and converting it into structured medical queries used to identify conditions from a knowledge graph. *Id.* at ¶¶ [0005]–[0006]. The Goldberg/Gnanasambandam combination already provides an API architecture for receiving and routing medical queries (Gnanasambandam), a hybrid knowledge graph with condition, finding, and investigation nodes linked to SNOMED and LOINC (Goldberg), and a diagnostic algorithm that generates differential diagnoses from patient inputs (Goldberg). Datla provides an NLP engine that receives verbatim patient text and converts it into structured medical queries. Datla’s NLP engine receives natural language, extracts patient symptoms, and uses them to query a knowledge graph. Ex. 1005 at ¶¶ [0052], [0055]; claim 1.

389. The result of this combination—a knowledge-graph-based diagnostic engine with an NLP front-end accessible via published APIs—would have been entirely predictable. Datla’s NLP engine produces structured output (“one or more patient symptoms 316 and one or more demographic parameters 318”) that is directly translatable to API query parameters. Connecting Datla’s NLP output to Gnanasambandam’s published APIs would require no inventive modification because the NLP engine’s structured output is the natural input for the API-accessible diagnostic engine. Each reference is used for its known purpose: Goldberg provides the knowledge graph and diagnostic algorithm, Gnanasambandam provides the API deployment layer, and Datla provides the NLP engine that converts unstructured clinical text into structured medical queries.

390. Each reference provides a well-described, implemented system. Goldberg describes a deployed clinical platform with a hybrid knowledge graph, diagnostic algorithm, and multiple user interfaces. Ex. 1006, Abstract; 3:5–18; 7:25–50. Gnanasambandam describes a working API architecture that publishes APIs performing AI-based analytics on medical knowledge graphs. Ex. 1007 at 100:24–34. Datla’s NLP engine produces structured output (“one or more patient symptoms 316 and one or more demographic parameters 318”) that is directly translatable to API query parameters. Ex. 1005, Abstract; ¶¶ [0006]–[0008]; claim 1. A POSITA

would have had a high expectation that combining these complementary, individually validated systems would produce a combined working system.

391. A POSITA would have been motivated to combine all four references (Goldberg, Gnanasambandam, Datla, and Dang) because they operate in the same field of computer-assisted medical diagnosis using knowledge graphs, and each contributes a non-overlapping functional layer: Goldberg (knowledge graph and diagnostic reasoning), Gnanasambandam (API platform), Dang (graph topology and next-inquiry generation), and Datla (NLP-based text processing and query generation). A POSITA would have had a high expectation that combining these four complementary, individually validated systems would produce a working combined system, because the NLP engine (Datla) simply provides a front-end text-processing layer that feeds structured queries into the same API architecture (Gnanasambandam) and knowledge graph system (Goldberg, supplemented by Dang) discussed in Ground 1.

2. Dependent Claim 15


392. Claim 15 recites: “A method as recited in claim 1 wherein said software application is a natural language processing (NLP) engine that receives verbatim text and produces said medical query that includes a patient diagnostic factor or a patient medical condition, and wherein said matching a portion includes matching said patient diagnostic factor to said at least one of said finding nodes or matching said

patient medical condition to said at least one of said medical condition nodes to produce said output.” I incorporate by reference my discussion of independent claim 1 above. *Supra* § X.A.2.

393. Claim 15 refers to “said matching a portion.” Claim 1 does not recite any matching step. I understand from counsel that a missing antecedent does not necessarily prevent a POSITA from understanding a claim where its meaning is reasonably clear in context. Here, the clause in which “said matching a portion” appears states that the matching includes: “matching said patient diagnostic factor to said at least one of said finding nodes or matching said patient medical condition to said at least one of said medical condition nodes to produce said output.” For purposes of my analysis, I read claim 15 as requiring that the method include matching the patient diagnostic factor in the medical query to at least one of the finding nodes, or matching the patient medical condition in the medical query to at least one of the medical condition nodes, to produce the output.

394. Claim 15 thus adds three requirements beyond independent claim 1: (1) the software application that queries the knowledge graph is an NLP engine; (2) the NLP engine receives verbatim text and produces the medical query including a patient diagnostic factor or patient medical condition; and (3) matching the patient diagnostic factor to finding nodes or matching the patient medical condition to condition nodes to produce the output.

395. Datla expressly discloses each of these requirements. As to the first requirement, Datla's system includes an NLP engine. Ex. 1005 at ¶¶ [0006], [0028], [0052], [0065], [0075], [0081]; claim 1; FIGS. 3 (element 314), 4, 5 (element 550), claim 1 ("a processor comprising a natural language processing engine configured to extract the at least one patient symptom and at least one demographic parameter from the received natural language input."). For example, NLP Engine 314 is shown in FIG. 3 (annotated below).



Ex. 1005 at FIG. 3 (annotated)

396. As to the second requirement, Datla's NLP engine is configured to "receive natural language input from a user, the input comprising information about at least one patient symptom and at least one demographic parameter about the patient." Ex. 1005 at ¶ [0007]; claim 1. The input may be spoken natural language

converted to text via speech-to-text: “a clinician, medical professional, or patient provides information to the automated system via a user interface. The information is provided in natural language.” *Id.* at ¶¶ [0050], [0074]. The input may also be typed text from “a mobile device, laptop, desktop, wearable device, home computing device, or any other computing device.” *Id.* Datla’s system thus receives verbatim text—unstructured natural language from a clinician or patient describing the patient’s scenario.

397. Datla’s NLP engine produces a medical query that includes a patient diagnostic factor or a patient medical condition. The NLP engine “extract[s] the at least one patient symptom and at least one demographic parameter from the received natural language input.” Ex. 1005 at ¶ [0007]; claim 1. The NLP engine further “extract[s] keywords related to symptoms, such as signs, lab results, procedures and demographic information.” *Id.* at ¶¶ [0006], [0028], [0052]. The extracted symptoms and conditions are then used to query the knowledge graph: “the system queries the knowledge [graph] using the extracted one or more patient symptoms.” *Id.* at ¶ [0055]. The extracted patient symptoms, together with the other extracted keywords such as signs and lab results, form the medical query, and each extracted symptom is a “patient diagnostic factor” within the meaning of claim 15.

398. As to the third requirement, Datla matches extracted patient diagnostic factors to nodes in the knowledge graph and produces the output. The knowledge

graph has “symptoms as root nodes connected by edges to the diseases and medical conditions associated with those symptoms.” Ex. 1005 at ¶ [0036]; *see also id.* at ¶ [0031]. The system generates a “diagnosis graph as a subset of the knowledge graph” via an activation-decay process: extracted symptoms are assigned activation weights, the graph expands outward from those symptom root nodes through connected nodes, and the system identifies “a ranked list of one or more medical conditions, diagnoses, treatments, and/or tests for the patient from the diagnosis graph.” *Id.* at ¶¶ [0007]–[0008], [0055]–[0057]; claim 1; FIG. 3. This constitutes matching the patient diagnostic factor to finding nodes (symptom root nodes) and matching the patient medical condition to condition nodes (disease nodes) to produce the output (ranked conditions, diagnoses, tests, and treatments).

399. In the combination, Datla’s NLP engine serves as the “software application” that receives verbatim text, extracts patient diagnostic factors and medical conditions, and sends structured medical queries via Gnanasambandam’s API to the Goldberg/Gnanasambandam knowledge graph system. Gnanasambandam supplies the API architecture through which Datla’s NLP output is delivered to the knowledge graph. Ex. 1007 at 101:12–37, 102:22–24, 100:24–34. The knowledge graph system matches the extracted factors to finding nodes and conditions to condition nodes using Goldberg’s diagnostic algorithm, and returns its output back through the API to the NLP engine, which can then present the results.

3. Dependent Claim 16

400. Claim 16 recites: “A method as recited in claim 15 wherein said NLP engine is initialized with at least one of said plurality of finding nodes from said knowledge graph in order to facilitate said matching.” I incorporate by reference my discussion of dependent claim 15 above. *Supra* § X.B.2.

401. Datla’s extracted symptoms are matched to nodes of its knowledge graph. Datla assigns the symptom weight as an activation weight “to a node of the knowledge graph comprising the extracted one or more patient symptoms,” and its initial forest starts from “root nodes associated with the extracted and weighted one or more symptoms.” Ex. 1005 at ¶¶ [0055], [0067]. For that step to work, the NLP engine’s output must correspond to the graph’s symptom nodes.

402. A POSITA would have found it obvious to initialize Datla’s NLP engine with the finding nodes of the knowledge graph it queries. Goldberg’s knowledge graph integrates SNOMED and LOINC. Ex. 1006 at 1:54–56; 6:5–28. Datla separately states that, “[a]ccording to an embodiment, the system may extract . . . weighted topical keywords . . . and map them to categories represented in one or more ontologies, including but not limited to the following controlled clinical ontologies: SNOMED CT for diagnoses, LOINC for tests, and/or RxNorm for treatments.” Ex. 1005 at ¶ [0054].

403. In the combination, a POSITA would have initialized Datla's NLP engine with the finding nodes of Goldberg's hybrid knowledge graph. Datla's query begins at the graph nodes that match the extracted symptoms (Ex. 1005 at ¶¶ [0055], [0067]), so an extracted term that does not match a node cannot start the query. Loading the graph's finding-node terms into the NLP lexicon was the simplest, routine way to make sure extracted concepts matched existing nodes.

404. Datla's knowledge graph further confirms that finding nodes inform NLP operation. Datla's activation-decay process begins at the symptom root nodes that match the extracted concepts and spreads to connected disease nodes. Ex. 1005 at ¶¶ [0036], [0055]–[0057], [0067]. This confirms that Datla's NLP output matches the graph's finding nodes, which is the purpose of the initialization recited in claim 16.

C. Ground 3: Claims 1-34 are Directed to Ineligible Subject Matter

405. In my opinion, claims 1-34 of the '883 Patent are directed to ineligible subject matter under 35 U.S.C. § 101 because they claim the abstract idea of collecting, organizing, storing, and providing access to medical data using conventional computer components and well-known techniques, without any inventive concept.

1. The Challenged Claims are Directed to an Abstract Idea

406. Independent claim 1 recites a method comprising the following functional steps: (1) receiving a medical query from a software application via an API; (2) using a knowledge graph that includes medical condition nodes, finding nodes, and investigation nodes extracted from medical evidence sources and linked to a medical terminology ontology (SNOMED-CT or LOINC); (3) identifying, using a “reasoning capability,” ancestor-descendant relationships in the knowledge graph relevant to the medical query; and (4) outputting data to the software application for generating a medical determination.

407. Independent claim 23 recites the same functional steps as claim 1 in system form. The only additional elements are “a database that stores a knowledge graph” and “a computer system comprising one or more processors and memory that stores instructions” that, when executed, perform those steps. Ex. 1001, claim 23. These are generic computing components that the specification describes in conventional terms, and the examiner treated them as such during prosecution. Ex. 1001 at FIGS. 11A–11B; Ex. 1002 at 88–89. They do not change what claim 23 is directed to and claim 23 is therefore directed to the same abstract idea as claim 1 for the reasons discussed below.

408. At their core, these claims are directed to collecting medical data from evidence sources, organizing that data into a structured representation (a knowledge graph), reasoning over the structured data by traversing hierarchical relationships

among medical concepts, and outputting the results. The focus of the claims is not on any specific improvement to computer technology but rather on the functional result of making medical evidence accessible and queryable. This process could be, and historically has been, achieved through manual processes.

409. The core activities described in the claims, including extracting diagnostic concepts from medical literature, organizing those concepts into hierarchical structures linked to medical ontologies such as SNOMED-CT and LOINC, traversing hierarchical relationships to identify related concepts, and using the results to support clinical decision-making, are processes that medical professionals and informatics specialists have performed, in substance, for decades. Before the earliest claimed priority date of the '883 Patent, clinicians and medical informaticists routinely reviewed medical literature, identified diagnostic factors and laboratory tests associated with medical conditions, organized those relationships in structured taxonomies and databases, and used hierarchical medical terminologies to reason about relationships among medical concepts. The challenged claims automate these well-established processes using conventional computing components but do not improve any of those components.

410. The knowledge graph recited in the claims is a conventional data structure. The specification describes the database as “[a]ny of a wide variety of other types of databases . . . such as a relational database or object store or object-

oriented database,” including, for example, a conventional, off-the-shelf JanusGraph graph database. Ex. 1001 at 8:10–22. Knowledge graphs organizing medical conditions, findings, and investigations were well-established in the art before the ‘883 Patent's priority date. As discussed in my Technology Background (*Supra* § VI.A), SNOMED-CT and LOINC are pre-existing, third-party medical ontologies that have been used in clinical informatics systems for decades. Linking clinical concepts to these ontologies was standard practice.

411. The “reasoning capability” limitation involving traversing ancestor-descendant relationships among nodes in the knowledge graph is a form of hierarchical, associative reasoning that medical professionals perform when they consult a medical reference ontology and follow parent-child relationships to generalize or specialize a clinical query. For example, a clinician who observes “double quotidian fever” and looks up the broader concept “fever” in a medical taxonomy to identify associated conditions is performing the same type of hierarchical traversal that the claims describe. SNOMED-CT’s classification structure, which organizes clinical concepts into a polyhierarchy where each concept may have multiple parent and child concepts, was designed expressly to support this type of traversal-based reasoning, and reasoning systems have long exploited these relationships. *Supra* § VI.F. The ‘883 Patent does not claim a new traversal

algorithm or a new type of ontological reasoning; it claims the application of well-known graph-traversal techniques to a pre-existing medical ontology.

412. Regarding the API limitation involving receiving a medical query via an API and outputting results to a software application, an API is a generic software interface that enables communication between software components. The specification describes only generic, off-the-shelf APIs, including FHIR, USCDI, CDS Hooks, and vendor-specific APIs, and does not disclose any new or non-conventional API technology. Ex. 1001 at 6:4–6. Exposing a knowledge-graph-based system through an API was a well-established design pattern in healthcare IT before the ‘883 Patent's priority date. As discussed in my Technology Background (*Supra* § VI.E), clinical decision support systems routinely used APIs to enable programmatic access to medical knowledge bases.

413. The prosecution history confirms that these claims are directed to an abstract idea. The patent examiner initially rejected all claims under Section 101, finding that apart from the generic recitation of an API and a computer system nothing in the claims precluded the steps from being performed in the human mind, and that the API-based receiving and outputting limitations were mere “insignificant extra-solution activity.” Ex. 1002 at 88–90. The examiner withdrew that rejection only after the applicant amended the claims to add the outputting limitation requiring the software application to “generate a medical determination” and argued that a

human cannot mentally “implement a software application . . . in order to access a knowledge graph via an API.” Ex. 1002 at 58–59.

414. I disagree with the Examiner’s finding. I understand that the relevant question is whether the claim as a whole is directed to an abstract idea, not whether every individual step can be performed mentally. The fact that a human cannot literally implement an API in his or her mind does not make the claim non-abstract because the underlying process involving receiving a query about a patient’s clinical state, consulting a hierarchically organized body of medical knowledge, traversing related concepts to identify relevant information, and providing a result is the same process that clinicians and medical informaticists have always performed.

415. Dependent claims 2–22 and 24–34 do not alter the abstract nature of the independent claims. These claims recite additional details about the content and structure of the knowledge graph (claims 2–4, 8–14, and 24–32), the size of the knowledge graph (claims 21–22 and 33–34), the type of software application that sends the query (claims 5 and 15–16), matching patient information to nodes in the knowledge graph (claims 6 and 20), the content, format, or destination of the output (claims 7, 17, and 18), and repeating the query when new patient information becomes available (claim 19).

416. Specifically, claims 2–3 and 24 add result nodes that hold the results of laboratory tests or examinations; claims 4 and 25 specify that the finding nodes are

diagnostic factor nodes; claims 8–10, 12, 26–28, and 30 specify that investigation nodes are independent of, disambiguated across, or free of information specific to the medical conditions to which they are linked; and claims 11, 13–14, 29, and 31–32 add a separate data structure, a ranking value, or a condition-dependent attribute associated with the path between an investigation node and a medical condition node.

417. Claim 5 specifies that the software application is an artificial intelligence application. Claim 15 specifies that the software application is a natural language processing engine that receives verbatim text and produces a query containing a patient diagnostic factor or medical condition, which is matched to a finding node or medical condition node to produce the output. Claim 16 adds that the NLP engine is initialized with finding nodes from the knowledge graph to facilitate that matching.

418. Claim 6 recites matching a diagnostic factor extracted from a patient's electronic health record to a finding node and outputting the associated medical condition. Claim 20 recites matching a portion of the query to a medical condition node or finding node. Claim 7 recites outputting a diagnostic factor that is not yet known to be present or absent in the patient. Claims 17 and 18 recite displaying the

output on a computer or outputting it in a data structure suitable for further processing.

419. None of these limitations changes the fundamental character of the claims. They merely specify the content and scale of the data being collected, organized, and retrieved, how patient information is compared with that data, or what is done with the result, or they recite conventional features of clinical decision-support systems, such as using AI applications or NLP engines, identifying which findings have not yet been elicited from a patient, or displaying results to a user.

2. The Challenged Claims Do Not Recite an Inventive Concept

420. When viewed individually or as an ordered combination, the Challenged Claims do not recite an inventive concept that transforms the abstract idea into patent-eligible subject matter.

421. The claimed knowledge graph and database are generic and conventional. As noted above, the specification discloses that any of a wide variety of databases may be used, including an off-the-shelf JanusGraph graph database. Ex. 1001 at 8:10–22. Building a medical knowledge graph with nodes representing conditions, findings, and investigations was well-established in the art. As I explained in my Technology Background (*Supra* § VI) and in my analysis of the prior art (*Supra* §§ IX, X.A–X.B), the prior art references I discuss above, including

Goldberg, Gnanasambandam, Dang, and Datla disclose medical knowledge graphs with the same types of node structures and relationships that the ‘883 Patent claims.

422. The medical terminology ontology limitation requiring SNOMED-CT or LOINC is not inventive. SNOMED-CT and LOINC are pre-existing, third-party controlled vocabularies that were widely used in clinical informatics systems before the ‘883 Patent’s priority date. Goldberg expressly integrates both SNOMED and LOINC into its hybrid knowledge graph. Ex. 1006 at 1:54–56; 6:5–28. Incorporating pre-existing ontologies into a knowledge graph is not an inventive concept.

423. The “reasoning capability” of traversing ancestor-descendant relationships is not inventive. As discussed above, SNOMED-CT’s polyhierarchy was designed to support exactly this type of hierarchical traversal. The claims merely apply conventional graph-traversal techniques to a pre-existing medical ontology. For example, inferring that a patient with “acute fever” also has “fever” by traversing upward in the SNOMED-CT hierarchy is a well-known technique, not an inventive advance. The ‘883 Patent does not claim a new traversal algorithm or a new type of reasoning; it claims the application of existing techniques to existing ontologies.

424. The claimed API limitations are generic and conventional. The specification describes only standard, off-the-shelf APIs (FHIR, USCDI, CDS Hooks, vendor-specific APIs). Ex. 1001 at 6:4–6. Using an API to provide programmatic access to a medical knowledge base was a well-established design

pattern, as Gnanasambandam demonstrates by publishing APIs for AI-based analytics on medical knowledge graphs. Ex. 1007 at 100:24–34; 101:12–37.

425. The outputting limitation requiring that the output include data for the software application to “generate a medical determination” does not supply an inventive concept. This limitation describes the functional purpose of the output (to enable a downstream application to generate a medical determination) but does not specify any novel mechanism for achieving that purpose. Outputting diagnostic results such as differential diagnoses, next-best actions, and associated clinical reasoning to a consuming software application through an API merely recites the intended use of the data rather than an inventive technical improvement.

426. The prosecution history confirms the absence of an inventive concept. The elements that the applicant added during prosecution to overcome the Section 101 rejection, including the API-based querying and output to a software application for generating medical determinations, were well known in the prior art, as demonstrated by the references analyzed herein. The examiner found, and the applicant did not dispute, that the knowledge-graph structure, node traversal, and reasoning-capability limitations are directed to a mental process. Ex. 1002 at 88–90. The only basis for allowance was the examiner’s conclusion that a human cannot

mentally implement an API. This does not establish an inventive concept because it merely adds a generic computer interface to an otherwise abstract process.

427. The dependent claims fare no better. Many of their additional limitations simply narrow the abstract idea itself by specifying what medical information is stored, compared, or reported, and the remainder recite conventional computing features.

428. Claims 2–4 and 24–25 specify that investigation nodes are linked to result nodes holding the results of laboratory tests or examinations and that the finding nodes are diagnostic factor nodes. These limitations only designate the categories of medical information stored in the graph and where a test result is recorded, which the specification describes as a matter of keeping reference ranges and results separate from the test itself. Claims 8–10, 12, 26–28, and 30 require that investigation nodes be independent of, disambiguated across, or free of information specific to the conditions to which they are linked. Storing a single, reusable entry for each test and placing condition-specific information elsewhere is a standard data-organization choice.

429. Claims 11, 13–14, 29, and 31–32 add a separate data structure, a ranking value, or a condition-dependent attribute associated with the path between

an investigation node and a condition node. Recording how strongly a test indicates a condition adds no technical mechanism.

430. Claims 21–22 and 33–34 recite knowledge graph size thresholds (more than twenty thousand or one hundred thousand nodes), which merely specify the scale of data stored, not a technological improvement. The ‘883 Patent does not claim that these node count thresholds have any technical significance.

431. Claim 5 specifies that the software application is an artificial intelligence application, which is a well-known application type. *Supra* § X.A.6. Claim 15 specifies that the querying application is an NLP engine that converts verbatim text into a query containing a diagnostic factor or medical condition that is matched to the graph, and claim 16 adds that the NLP engine is initialized with finding nodes from the graph. The specification describes its NLP processes as conventional techniques and describes initialization with graph nodes only as a preference, without any particular technique for performing it.

432. Claims 6 and 20 recite matching patient information, including factors extracted from an electronic health record, to finding or condition nodes. Comparing a patient’s findings with known associations is the core of the abstract idea itself, not an additional element.

433. Claim 7 recites outputting a diagnostic factor not yet known to be present or absent in the patient, which is the ordinary next step in differential

diagnosis: identifying which findings have not yet been elicited. The claim recites no particular technique for selecting that factor. Claims 17–18 recite displaying output on a computer and outputting in a data structure suitable for further processing, i.e., generic computing functionality.

434. Claim 19 recites receiving a new diagnostic factor and producing a new output, which is simply repeating the same process with updated information, as the specification describes.

435. Finally, claim 23 recites the same steps as claim 1 in system form, adding only a database that stores the knowledge graph and “a computer system comprising one or more processors and memory that stores instructions,” both of which are generic components. Dependent claims 24–32 add the same limitations as method claims 2, 4, and 8–14, respectively, and claims 33–34 add the same size thresholds as claims 21–22; each lacks an inventive concept for the reasons given for its method counterpart.

436. Considered individually or as an ordered combination, none of these limitations transforms the claims into patent-eligible subject matter.

437. In summary, the Challenged Claims recite the abstract idea of collecting, organizing, and reasoning over medical data using a knowledge graph linked to medical ontologies, and none of the claim elements, individually or as an ordered combination, provides an inventive concept that transforms the claims into

patent-eligible subject matter. Every component of the claimed system was well known and widely used in clinical informatics before the '883 Patent's priority date.

XI. SECONDARY CONSIDERATIONS

438. I have found no evidence either in the '883 patent's prosecution history or elsewhere that supports any secondary considerations arguments, or evidence of nexus to any Challenged Claim. As I have discussed herein, any purported solutions to problems or unexpected results in the '883 patent were already well-known.

439. To the extent Patent Owner asserts the existence of secondary considerations, or presents any additional evidence or arguments, I reserve the right to address any such arguments and evidence.

XII. ADDITIONAL REMARKS

440. The opinions I have expressed in this Declaration are based on the information and materials currently available to me. I reserve the right to supplement, amend, or modify my opinions if additional information becomes available, including any arguments, evidence, or declarations submitted by Patent Owner or any other party in this proceeding.

441. I am prepared to provide testimony on the matters addressed in this Declaration if called upon to do so by the Board or any party in this proceeding.

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true, and further, that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code, and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Dated: October 3, 2026

By: 
Mark A. Musen, M.D., Ph.D.