# Employing Hybrid Reasoning to Support Clinical Decision-Making

Sabbir M. Rashid[1][0000-0002-4162-8334]

Rensselaer Polytechnic Institute, Troy, NY, 12180, USA

**Abstract.** Clinical reasoning, involving abstraction, abduction, deduction, and induction, is the primary tool that physicians use when making clinical decisions. To support them, we focus on the creation of an AI system that is able to emulate clinical reasoning. We leverage Semantic Web technologies to perform a set of AI tasks involving the various forms of inference associated with clinical reasoning strategies. In particular, for the scope of this work, we focus on clinical problems that require differential diagnosis techniques. For a given clinical scenario, overlapping reasoning types and strategies may be employed by a physician in conjunction, signifying the need for our AI system to perform hybrid reasoning. Therefore, we consider the construction of a hybrid reasoner that is compatible with description logics. For medical scenarios where description logics may not have some needed expressivity, we consider possible extensions that will allow for the representation of such a scenario. The reasoning system, clinical rule representation, and the resulting recommendations will be evaluated based on domain expert consultation in order to determine whether the recommendation aligns with what the expert would recommend.

**Keywords:** Hybrid Reasoning · Deduction · Abduction · Temporal Reasoning · Knowledge Representation · Diabetes · Ontology · Explainable AI · Inference · Disease · Informatics · Differential Diagnosis · Clinical Reasoning Strategies · Rule Representation

## 1 Problem statement

This thesis focuses on the design and implementation of a clinical decision support system that is able to perform hybrid reasoning by emulating how physicians reason. Clinical reasoning is employed in many medical tasks, such as those involving information comprehension, decision-making, and medical error identification. Approaches for clinical reasoning use four distinct types of inference: abstraction, abduction, deduction, and induction [1, 15]. We base our approach on the Select and Test Model (ST-Model), an epistemological framework for medical reasoning in which an expert chooses a plausible hypothesis which is subsequently confirmed or falsified through testing [19]. The cyclic nature of the ST-Model demonstrates how the different types of reasoning can be applied in conjunction during a clinical reasoning framework. Given a set of initial patient data/information, abstraction is applied to identify problem features, which can be used, through abduction, to determine a set of plausible hypotheses. Alternatively, inductive techniques can be applied to arrive at hypotheses directly from the data. In applying this model for clinical reasoning, we deviate from the traditional ST-Model and instead of including the induction step, we use the observed data to validate diagnostic hypotheses. The resulting hypotheses are then ranked in terms of likelihood as well as clinical relevance, and are used for the deductive inference of expected outcomes or consequences. In order to support clinical decision-making, we note that three reasoning strategies commonly employed by physicians include differential diagnosis, treatment planning, and plan critiquing. In this paper, we focus primarily on differential diagnosis.

To illustrate the applicability of this revised ST-Model framework using a simple example, consider a scenario where a patient gained a significant amount of weight over a short period of time. The observed data that the patient gained 30 pounds over 2 months is abstracted to be a clinical problem. Abduction is used to determine four likely hypotheses for the patient unintentionally gaining weight, such as 1) a decrease in the amount of activity, 2) a change in the patient's diet, 3) a medicinal effect, or 4) a biological effect.[1] These hypotheses can be ranked in terms of importance, where medicinal or biological effects are more clinically relevant than a decrease in activity or a change of diet. Existing data in the system can be used to rule out or validate the hypotheses. For example, if Internet of Things (IoT) data pertaining to step counts shows that the patient's amount of activity has not decreased, hypothesis 1 can be ruled out. Hypothesis 3 regarding medicinal effects could be broken into sub-hypotheses, such as due side-effects of a single medicine or contraindications between multiple medications. This would cause the system to check for symptoms of the drugs that the patient is prescribed. The biological effect of hypothesis 4 could be a decrease in the patient's metabolism, which in turn may be explained by a disease like hypothyroidism, and would require more data for validation. Under the open-world assumption, this hypothesis would not be ruled out, and instead the system would recommend a test to gather more data, such as a TSH (thyroid-stimulating hormone) test to check for thyroxine (a thyroid hormone) levels.

In order to implement the ST-Model framework into an Artificial Intelligence (AI) system, our approach involves semantically annotating and transforming patient data into Resource Description Framework (RDF), individually encoding each reasoning technique as an agent that detects changes and reasons over content in the knowledge graph, and validating ranked hypotheses through test cases and expert consultation.

## 2 Importance

This work is geared towards aiding physicians during the clinical decision-making process by creating an AI system that can reason in a way similar to how a physician reasons. Earlier work on medical expert systems use approaches involving

[1] https://www.webmd.com/diet/ss/slideshow-weight-gain-shockers

decision tree rules or the construction of knowledge bases by leveraging machine learning techniques. Our approach differs from that earlier work in that we are building an AI system using semantic web technologies in a manner that conforms with Description Logics (in particular, with OWL-DL). This allows us to both leverage recent ontologies created in knowledge representation for health care efforts and existing inference engines that reason over OWL-DL ontologies. Furthermore, we note that by limiting expressivity, we can perform tractable reasoning over OWL-DL ontologies. Therefore, a challenge worth exploring involves the representation of medical knowledge such that polynomial time reasoning can be achieved while preserving important domain and patient details.

Furthermore, since part of this work involves the implementation of an abductive reasoning component, we are advancing semantic web techniques to allow for non-monotonic reasoning (where conclusions can be validated or rejected as more data is observed). As Chitsaz defines in her thesis [2], abductive reasoning is "the process of inferring the best explanation from given observations [in accordance with] background knowledge." In the example related to patient weight gain presented in Section 1, abductive reasoning is used to generate several hypotheses. Non-monotonic reasoning allows for each of these hypotheses to exist in a state in which they are not yet confirmed to be true or false, allowing for acceptance or rejection once more data is leveraged for validation. This process, which is necessary for emulating the way a physician reasons, can be recorded through the use of provenance capturing. By creating a framework that is able to perform abduction in this way, we promote traceability in reasoning and move towards the goal of explainable AI systems.

## 3 Related Work

As early as the 1950s, Artificial Intelligence (AI) techniques were beginning to be applied to model clinical reasoning and medical decision-making [8]. By the end of the decade, work on the reasoning foundations for medical diagnosis, in which Bayes' formula was adapted to find conditional probabilities of a disease given a set of symptoms [9], lead to the application of probabilistic approaches for modeling clinical reasoning over the next decade. In addition to probabilistic approaches for diagnostics, research around this time also focused on medical literature retrieval and indexing, such as MEDLARS (Medical Literature Analysis and Retrieval System [21]).

By the 1970s, theoretical research on medical reasoning was beginning to be applied towards the creation of medical expert systems. Approaches related to the Newell-Simon methodology on how humans solve problems in different situations [12] were applied in the creation of expert systems [7]. Also around this time period (1970s-80s) SUMEX-AIM (Stanford University Medical Experimental – AI in Medicine [6]) arose, a computational infrastructure shared by multiple institutions. This lead to discussions and workshops related to biomedical problem solving and clinical decision-making [8].

While work on expert systems continued into the 80s, the importance of having these systems be explainable and justified became more apparent. Explainable expert systems, developed in collaboration between knowledge engineers and domain experts, represented a shift from procedural encoding to declarative knowledge representation [20]. Relevant literature from this time period also includes work on medical reasoning, including the ideas that physicians develop hypotheses early in the diagnostic process [4] and experts differ from novices in approaches involving hypothetico-deductive reasoning [16]. These differences derive both from the features of a knowledge base [5] and highly automated perceptual processes [10] that experts have in contrast with novices.

An approach that builds on this work uses propositional analysis to build rules from knowledge that abstracts medical problem solving into steps involving data gathering, diagnosis, therapeutic planning, and patient management [16]. Another important idea from this work is that medical reasoning often involves both forward and backward reasoning. This direction of research continued throughout the 90s to show that propositional and semantic analyses can improve the validity, usability, and comprehension of Clinical Practice Guidelines (CPGs), mental models could be used as a cognitive framework for how people reason, and reasoning approaches were often data and hypothesis driven [14].

## 4 Research questions

In conducting this work, we focus on several research questions. We begin by determining the following: *Which forms of reasoning should be incorporated into a clinical decision support system in order to emulate how a physician reasons?* When consulting related literature in order to determine the forms of reasoning that should be incorporated into a clinical decision support system, we find that mentions of abstraction, deduction, induction, abduction, temporal reasoning, and causal reasoning [1, 19, 15]. Since our approach for replicating physician reasoning is based on the ST-Model [19], we focus on the abstraction process to formulate observed data into a representative problem space. For observations occurring over a period of time, we leverage existing approaches for temporal reasoning [13, 18]. Through abduction, we determine potential hypotheses that can be used to explain the clinical problem at hand. We incorporate probabilistic reasoning techniques when ranking abductive hypotheses in order to first address more clinically relevant hypotheses. From each hypothesis, we deductively determine a set of predictions corresponding to expected patient data. While using existing semantic technology for deduction, we do consider the writing of custom inference rules when necessary. While we do not focus on induction, we attempt to validate abductive hypotheses by comparing the expected and observed data.

The next research question we consider relates to design requirements: *What are the necessary implementation considerations for building a hybrid reasoner that conforms with OWL Description Logics (OWL-DL)?* Specifically, we consider satisfiability, consistency, and explainability. In terms of satisfiability, we search for contradictions with respect to the ontologies that are used and between facts that are asserted to be true. To check for consistency, we determine if the facts in the knowledge base are satisfiable. Inconsistencies within the knowledge base result in instances of the owl:Nothing class. If the reasoner does not return such instances, we conclude that the knowledge base is consistent. Finally, we require an explanation for facts derived from inference processes. In order to guarantee that our hybrid reasoner conforms with DL, initially the implementation is limited to Semantic Web compliant technologies. We consider the extent to which the expressivity of the ontologies that we use can be limited in order to allow for tractible reasoning while still allowing the necessary medical knowledge to be accurately represented. As additional functionalities are incorporated, such as non-monotonic abductive reasoning, tests will be written to ensure that the systems is constrained to the bounds of DL. Boundary limit testing will be conducted to determine which scenarios break this constraint.

For our considered usage scenarios, we ask the following: *How can verbal statements made by physicians (captured in written transcripts) be represented as reasoning rules?* As mentioned earlier, our approach is based on the ST-model. In terms of representing verbal statements, we consider representational simplification by abstracting out mentions of conditions, treatments, diseases/diagnoses, risks/symptoms, and goals. In doing so, we must ensure that the resulting rule is both semantically consistent and representative of the original recommendation or statement. For example, in the use case of Section 1, the *symptom* of weight gain may point to a *diagnosis* for hypothyroidism under the *condition* that thyroxine levels are low. A diagnosis of hypothyroidism may result in *treatment* in the form of a levothyroxine prescription (commonly used to treat hypothyroidism) in order obtain the *goal* of decreasing the patient's weight.

## 5 Preliminary results

In order to apply our approach on patient data, we represent the data as RDF in a semantic structure consistent with the execution of the reasoning rules derived from physician transcripts. The Semantic Data Dictionary (SDD) approach [17] has been applied for annotating and converting patient data to RDF from several publicly available biomedical datasets. Additionally, we have received permission to annotate data from two private datasets, StepUp and Limited Claims EHR Dataset (LCED). StepUp contains data about its user's physical activities, including step count, so it can be used to test for hypothesis 1 of the use case of Section 1. LCED contains administrative claims information as well as Electronic Health Record (EHR) data, including information about patient demographics, habits, and prescriptions [3]. Therefore, it can be used for patient characteristic representation for testing using real patient data.

In terms of implementation, a customizable deductive reasoner has been incorporated into the Whyis [11] knowledge management framework. Preliminary work on traceability within this framework has commenced and is being leveraged to aid in the abductive reasoning component. Initial work related to clinical decision support and rule representation includes the categorization of reasoning types and strategies from physicians transcripts, as well as discussions on how to adapt the ST-Model for use within the framework of this thesis.

## 6 Evaluation

In this section, we begin by discussing some of the evaluation that has been conducted for the aforementioned methodology. We continue by discussing additional evaluation of our clinical decision-making system that we plan to conduct.

The SDD approach, which is leveraged to semantically convert tabular data into RDF, has been evaluated against traditional data dictionaries, data integration approaches, and mapping languages [17]. Reasoning component implementations have been validated using unit tests to check for the expected functionality. As further components are incorporated, additional tests will be written accordingly, including tests of efficiency when reasoning over large data.

In addition to testing the functionality of the implementation, we will use background biomedical knowledge to evaluate clinical reasoning processes. By providing a clinical expert with sets of test patient data and asking them to walk through their diagnostic process, we will evaluate the set of hypotheses that the system generates through comparison with the physician's hypotheses, which we will measure through the use of confidence scores. While the determination of the full evaluation pipeline is still work in progress, in addition to measuring hypothesis confidence, we plan to conduct a comparative evaluation against other hybrid reasoning systems. Additionally, for each individual reasoning component, we will evaluate capabilities against other special purpose reasoners that are able to do the associated reasoning task. The complete set of metrics for such an evaluation is yet to be determined. Abductive explanations will be ranked in terms of likelihood using appropriate heuristic metrics. One such metric corresponds to whether the hypotheses are accurately ranked in terms of clinical importance (such as medicinal implications being more relevant than a change in the patient's diet). The accuracy of these rankings will be evaluated based on importance assigned by a domain expert who is provided the same set of hypotheses. Initially, the clinical expert we will consult is a member of the author's thesis committee who holds both Ph.D. and M.D. degrees in relevant areas.

## 7 Reflections

We have discussed the creation of a decision support system that is similar to some of the expert systems that we encountered during the literature review. While these systems used various AI techniques, they were not implemented using semantic web technologies. Our approach differs from earlier methods as we are creating a system that conforms with OWL-DL. We are able to integrate various forms of inference by leveraging existing efficient reasoners (for example, for performing deduction or temporal reasoning) that also conform with OWL-DL semantics. Additionally, semantic knowledge resources, such as biomedical ontologies and patient data, can be incorporated into our knowledge base that can be used during the decision-making processes. The use of semantic technologies allows us to capture provenance, detect inconsistencies, integrate multiple knowledge resources, and leveraging existing tools for performing deductive and temporal reasoning. Our application can help address specific challenges by aiding physicians with clinical problems involving reasoning strategies.

Future work includes the incorporation of other reasoning strategies, such as therapy planning and plan critiquing. Extensions in this direction include leveraging CPGs to represent existing therapy plan recommendations and test physician conformance to CPGs. Additionally, medical information from biomedical resources that provide recommendations based on scientific evidence, as well as clinical opinions on whether or not to follow guideline recommendation, can be used to model plan critiquing. We mention in Section 4 that the incorporation of inductive reasoning is beyond the scope of this thesis, as the implementation of an abductive component will be sufficiently challenging in itself. Future extensions may include leveraging existing machine learning techniques to incorporate induction capabilities into our framework.

## Acknowledgements

I am advised by Professor Deborah L. McGuinness, Constellation Chair at Tetherless World Constellation (TWC), RPI. I am thankful for the members of TWC who have helped me edit this paper. This work is done as part of the Health Empowerment by Analytics, Learning, and Semantics (HEALS) project and is supported by IBM Research AI through the AI Horizons Network.

# Bibliography

[1] Arocha, J.F., Wang, D., Patel, V.L.: Identifying reasoning strategies in medical decision making: a methodological guide. Journal of biomedical informatics **38**(2), 154–171 (2005)
[2] Chitsaz, M.: Approaches to Abductive and Inductive Reasoning in Lightweight Ontologies. Ph.D. thesis, Griffith University (2015)
[3] Choudhury, O., Gkoulalas-Divanis, A., Salonidis, T., Sylla, I., Park, Y., Hsu, G., Das, A.: Differential privacy-enabled federated learning for sensitive health data. arXiv preprint arXiv:1910.02578 (2019)
[4] Elstein, A.S., Shulman, L.S., Sprafka, S.A.: Medical problem solving an analysis of clinical reasoning (1978)
[5] Feltovich, P.J.: Knowledge based components of expertise in medical diagnosis. Tech. rep., PITTSBURGH UNIV PA LEARNING RESEARCH AND DEVELOPMENT CENTER (1981)
[6] Freiherr, G.: The seeds of artificial intelligence: SUMEX-AIM. No. 80, US Department of Health, Education, and Welfare, Public Health Service ... (1980)

[7] Hayes-Roth, F., Waterman, D.A., Lenat, D.B.: Building expert system (1983)
[8] Kulikowski, C.A.: Beginnings of artificial intelligence in medicine (aim): Computational artifice assisting scientific inquiry and clinical art–with reflections on present aim challenges. Yearbook of medical informatics **28**(01), 249–256 (2019)
[9] Ledley, R.S., Lusted, L.B.: Reasoning foundations of medical diagnosis. Science **130**(3366), 9–21 (1959)
[10] Lesgold, A.M., Feltovich, P.J., Glaser, R., Wang, Y.: The acquisition of perceptual diagnostic skill in radiology. Tech. rep., PITTSBURGH UNIV PA LEARNING RESEARCH AND DEVELOPMENT CENTER (1981)
[11] McCusker, J., Rashid, S.M., Agu, N., Bennett, K.P., McGuinness, D.L.: The whyis knowledge graph framework in action. In: International Semantic Web Conference (P&D/Industry/BlueSky) (2018)
[12] Newell, A., Simon, H.A., et al.: Human problem solving, vol. 104. Prentice-Hall Englewood Cliffs, NJ (1972)
[13] O'Connor, M.J., Das, A.K.: A method for representing and querying temporal information in owl. In: International joint conference on biomedical engineering systems and technologies. pp. 97–110. Springer (2010)
[14] Patel, V.L., Arocha, J.F., Diermeier, M., Greenes, R.A., Shortliffe, E.H.: Methods of cognitive analysis to support the design and evaluation of biomedical systems: the case of clinical practice guidelines. Journal of biomedical informatics **34**(1), 52–66 (2001)
[15] Patel, V.L., Arocha, J.F., Zhang, J.: Medical reasoning and thinking. The Oxford handbook of thinking and reasoning pp. 736–754 (2012)
[16] Patel, V.L., Groen, G.J.: Knowledge based solution strategies in medical reasoning. Cognitive science **10**(1), 91–116 (1986)
[17] Rashid, S.M., McCusker, J.P., Pinheiro, P., Bax, M.P., Santos, H.O., Stingone, J.A., Das, A.K., McGuinness, D.L.: The semantic data dictionary – an approach for describing and annotating data. Data Intelligence pp. 443–486 (2020)
[18] Rospocher, M., van Erp, M., Vossen, P., Fokkens, A., Aldabe, I., Rigau, G., Soroa, A., Ploeger, T., Bogaard, T.: Building event-centric knowledge graphs from news. Web Semantics: Science, Services and Agents on the World Wide Web (2016). https://doi.org/10.1016/j.websem.2015.12.004, `http://www.sciencedirect.com/science/article/pii/S1570826815001456`
[19] Stefanelli, M., Lanzola, G., Ramoni, M.: Knowledge acquisition based on an epistemological model of medical reasoning. In: 1992 14th Annual International Conference of the IEEE Engineering in Medicine and Biology Society. vol. 3, pp. 880–882. IEEE (1992)
[20] Swartout, W.R.: Xplain: A system for creating and explaining expert consulting programs. Artificial intelligence **21**(3), 285–325 (1983)
[21] Taine, S.: The medical literature analysis and retrieval system (medlars) of the us national library of medicine. Methods of information in medicine **2**(02), 65–69 (1963)