

Building a Structured Reasoning AI Model for Legal Judgment in Telehealth Systems

Jingyuan Xu

*School of Computer and Information Sciences, University of the Cumberlands, Kentucky, USA
jxu65428@ucumberlands.edu; ORCID: 0000-0001-9448-1527*

Abstract: This paper introduces a structured AI framework developed to support legal decision-making in telehealth environments. Most existing systems either rely on user declarations or conduct audits after the fact. Our model takes a different approach: it integrates legal reasoning into the AI's core logic. The system is built on an architecture that includes structured semantic modeling, executable legal logic, and role-specific response generation. Beyond enabling real-time legality checks, the model also accounts for ambiguity in legal language by incorporating fuzzy clauses, confidence scores, and guidance pathways that reflect varying levels of legal certainty. Instead of treating compliance as a checklist imposed from outside, the model treats it as something that unfolds from within the system's logic. The result is an AI that can explain its decisions, adapt to legal environments, and support institutional accountability. This work offers a new perspective on how AI can operate not just as a tool, but as a responsible actor in regulated clinical systems, and points to new directions for designing legally responsive telehealth platforms.

Keywords: artificial intelligence, explainable AI, legal compliance, telehealth systems, semantic reasoning, health law

Introduction

In recent years, telehealth has developed very fast and become more common in healthcare (Hameed et al., 2020). Many patients now prefer this method because it allows them to speak with doctors from home. However, this convenience also creates serious legal and ethical concerns (Ivanova et al., 2025). Especially for young patients, different states in the U.S. have their own laws about age, consent, and privacy (Garber & Chike-Harris, 2019). For instance, in New York, teenagers who are 16 or older can agree to mental health inpatient treatment without needing their parents (New York Consolidated Laws, Mental Hygiene Law § 9.13, 2024). But in Missouri, any patient under 18 still needs their parents' or legal guardians' permission to receive such treatment (Missouri Revised Statutes § 632.070, 2014). These legal differences make it very important for telehealth systems to include local compliance rules. Otherwise, providers might break the law by accident (Sharko et al., 2022).

Unfortunately, some cases show the risk clearly. In a serious event reported by The Wall Street Journal, a 17-year-old in Missouri was treated through telehealth without parental consent and later died by suicide (Safdar, 2022). The investigation said that the telehealth service did not check the patient's age properly, so the system allowed treatment that was not legal in that state. This kind of situation shows why we need smarter systems that can understand different laws and help both doctors and patients follow them.

Many current AI systems in healthcare only work in the background (Kuziemy et al., 2019). They might stop a task if it breaks a rule but do not explain why, or they only give feedback after a mistake has happened. We aim to design a system to interact with users directly. For example, it can ask questions like "How old are you?" and "Which state are you in?" Based on the answers, it gives legal explanations, such as "In many states, parental consent is required for mental health treatment." This way, users get clear and useful guidance in real time instead of just a yes-or-no answer.

The system is not designed as a simple layered pipeline. Instead, it includes several connected modules. These are: Dialogue Manager (the part that talks with users), Natural Language Understanding (NLU), a Legal Knowledge Base, a Fuzzy Inference Engine, a Compliance Reasoner, and an Explanation module. These parts work together, but each has its

own role. For instance, if a patient says “I’m seventeen and a half,” the Fuzzy Engine helps the system understand this vague input by using partial values. Then the system may ask again to confirm. This flexible design helps the system deal with real-life situations where users often speak in unclear or incomplete ways.

In this paper, we describe each part of the system and how they work together. We focus on a real example: a teenager between 16 and 17 years old wants mental health counseling online. The system must determine whether parental consent is needed, show how it makes that decision, and give advice on what to do next. We validated the framework through structured formula-based reasoning, using defined fuzzy membership equations and jurisdiction-specific rules to simulate representative consent scenarios. The results demonstrate that the system consistently reaches correct compliance decisions and provides clear, explainable outputs aligned with relevant state laws. The key point of our work is combining fuzzy language handling with legal rules in a transparent and modular system that can be useful in future telehealth applications.

System Architecture and Capabilities

A Dialogue Manager is the user interface in this system architecture, just like other telehealth chatbots. In particular, it prompts the user to provide any missing information and confirms essential details, for example, by asking for the patient’s age or location when necessary. This approach is similar to existing AI assistants that collect patient data step by step. Our focus is on ensuring legal compliance. Using pattern matching or simple intent recognition, a lightweight Natural Language Understanding (NLU) component processes each word to extract key facts (such as age or issue type). The dialogue manager and NLU resemble typical healthcare chatbot pipelines that convert user input into structured data and guide the dialog flow.

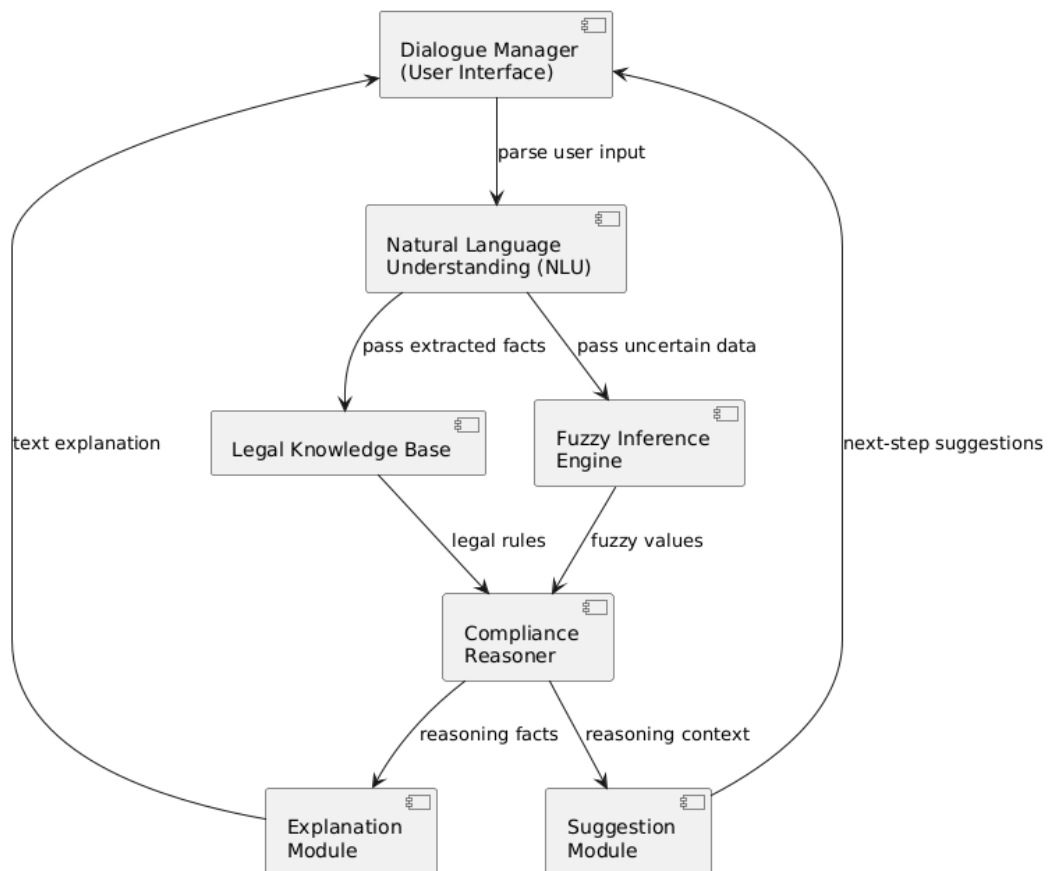


Figure 1. Modular Compliance-Aware Telehealth AI Architecture

Figure 1 illustrates the modular architecture of the proposed compliance-aware AI system. The framework consists of several components, including a dialogue manager for user interaction, a natural language understanding module for parsing input, a fuzzy inference engine for handling ambiguity, a compliance reasoner for legal rule evaluation, and an explanation module to deliver user-facing outputs. While the figure presents a complete system pipeline, this paper focuses primarily on the fuzzy reasoning mechanism—specifically, the use of membership functions and fuzzy rule evaluation to assess consent requirements under variable age thresholds and jurisdiction-specific legal rules.

The fuzzy inference engine handles all imprecise or incomplete inputs (Yorita et al., 2023). In natural language, people may say “almost 18” or “I’m seventeen and a half,” which do not map cleanly to a single number. The fuzzy engine interprets such expressions by assigning membership degrees to age categories. For example, we might define a fuzzy set “minor” that fully includes ages below 17 and a set “adult” for ages above 18, with a smooth transition in between. In our framework if the user says “17.5,” the system might compute a membership of 0.5 in “minor” and 0.5 in “adult.” This fuzzification process means mapping a crisp input to degrees of membership between 0 and 1.

To represent the semantic vagueness of age-related conditions, we define two fuzzy sets: minor and adult, each with linear membership functions. For example, the membership value of 17.5 is 0.5 in both sets, indicating semantic ambiguity. These values are then used in rule evaluations through fuzzy conjunctions, such as $\min(\mu_{\text{minor}}, 1)$ when paired with crisp conditions like the patient’s location. This formalization enables reasoning over soft boundaries and avoids brittle binary thresholds.

Fuzzy logic is well-suited for such vagueness: it provides a mathematical way to represent imprecision and “degrees of truth”. In practice, the fuzzy engine evaluates linguistic terms and outputs fuzzy facts that carry these partial truth values. This is similar to how some conversational chatbots use fuzzy models to handle language ambiguity; for instance, Cleverbot incorporates fuzzy logic to manage uncertain conversational cues. In our system, the fuzzy output may prompt the bot to ask for clarification when a condition is only partially met (for example, “Are you at least 18 years old?”) rather than making a binary decision. The use of fuzzy membership allows the system to ask follow-up questions on borderline cases, preventing hasty compliance judgments when information is not clear.

The Compliance Reasoner is the core rule-based component that applies legal and policy rules to the collected information. It takes inputs from NLU and the fuzzy engine (for example, a partial age value) and uses rules from the knowledge base to infer what actions are legally required. Each rule is typically an “if-then” statement derived from regulations (for example, “if age < 18 and jurisdiction is Missouri and service is mental health, then parental consent is required”). During reasoning, fuzzy truth values are combined with crisp conditions: for instance, if the patient’s age is partly “minor” and partly “adult,” the rule may fire with intermediate strength. The reasoner essentially executes all applicable rules to compute a result (in classic fuzzy inference style). Based on this, it draws a conclusion such as “consent needed,” “consent not needed,” or a “borderline” situation. The system is designed so that a partially triggered rule will not immediately allow an action; instead, the assistant will advise caution and typically seek more precise data before finalizing a decision. This careful fuzzy inference ensures that legal logic is applied transparently (Reddy, 2022) and that uncertain cases are handled by engaging the user rather than failing silently.

To handle the inherent semantic vagueness in legal rules (Mukhopadhyay et al., 2025) based on age, we formally define two fuzzy sets—minor and adult—with corresponding membership functions. These membership functions represent mathematically how clearly a given age falls into each category. Specifically, the membership in the “minor” set is calculated as:

$$\mu_{minor}(age) = \begin{cases} 1, & age \leq 17 \\ 18 - age, & 17 < age < 18 \\ 0, & age \geq 18 \end{cases} \quad (1)$$

Similarly, the membership in the “adult” set is defined as:

$$\mu_{adult}(age) = \begin{cases} 1, & age \leq 17 \\ age - 17, & 17 < age < 18 \\ 0, & age \geq 18 \end{cases} \quad (2)$$

While the value 18 is used in this paper for illustrative purposes, the age threshold can be flexibly adjusted based on state-specific legal requirements. For example, in Illinois, the consent age is 16 (World Population Review, n.d.), and the membership functions can be recalibrated accordingly. This ensures the framework remains jurisdiction-sensitive and adaptable across regions. These functions ensure a gradual, linear transition between the categories, rather than a sharp cutoff. For instance, if the patient states their age as 17.5, the system computes membership values of 0.5 for both (1) and (2). Once these fuzzy values are calculated, the compliance reasoner uses them to evaluate legal rules with fuzzy conjunctions (typically using the min operator). Consider a rule such as “If age is minor AND the state is Missouri, then parental consent is required.” Given a patient from Missouri whose stated age is 17.5, the rule’s activation strength is calculated as $\min(\mu_{minor}(17.5), 1.0)$, which equals 0.5. Because the resulting truth value (0.5) indicates ambiguity rather than a clear truth or falsehood, the system avoids immediate binary judgment. Instead, the assistant issues a cautious response, such as asking the patient for clarification (“Are you definitely under 18?”) or explicitly communicating that parental consent is likely necessary but additional confirmation is recommended. This fuzzy inference logic allows the system to mimic nuanced human judgment, effectively navigating borderline legal scenarios without committing prematurely to rigid yes/no conclusions.

Finally, once a conclusion is reached, a simple Explanation module converts the logical outcome into human-readable text (e.g. “You are 17 and Missouri law requires a parent’s permission for counseling, so we cannot proceed without it”). An optional Suggestion module can offer alternatives (for example, “You could arrange for a guardian to join the session”). These modules are auxiliary but ensure the system provides clear guidance after reasoning.

Overall, the system uses a capability-oriented flow: the dialogue manager, NLU, fuzzy engine, and reasoner interact flexibly rather than in a rigid pipeline. For example, when the NLU detects ambiguity, control temporarily shifts to the fuzzy engine; after the reasoner decides, control returns to the dialogue manager to deliver the response. In this way, each module can be developed or improved independently, and the fuzzy logic at the center helps the assistant emulate human-like reasoning under uncertainty.

The modular design allows this compliance logic to be integrated into existing telehealth workflows. For instance, the fuzzy engine and compliance reasoner could be deployed as backend services or plugins alongside a patient intake chatbot or electronic health record system. They would receive patient information (age, location, service type, etc.) from the telehealth front end and return a compliance status or recommendation. Because the modules communicate via well-defined interfaces, a provider could attach them to any conversational interface (e.g. a messaging app, web form, or voice bot) without overhauling the system. This design makes augmenting current telehealth assistants with legal compliance checks straightforward: the assistant simply invokes the fuzzy compliance components whenever consent rules must be verified. The result is a telehealth application that gathers patient data and reasons with it, delivering immediate, explainable feedback on legal requirements and next steps.

Illustrative Evaluation of the Fuzzy-Consent Framework

All outputs were derived analytically using the model's equations and legal rules, demonstrating the system's reasoning capability across a range of scenarios. Each case was examined through formula-based inference: user inputs (such as age and state) were converted into fuzzy values using the proposed equations, and rule evaluation followed accordingly to reach a compliance decision. This logical deduction process illustrates how the framework handles ambiguity and produces explainable outputs by systematically applying the defined membership functions and legal rules. The reasoning remains transparent and reproducible, grounded in the mathematical structure of the model rather than empirical trial-and-error.

Convert the crisp age into fuzzy membership degrees in the "minor" and "adult" categories. In fuzzy logic, each input is assigned a degree of membership between 0 and 1. Determine the applicable legal threshold for self-consent based on the state and combine it with the fuzzy values. We use a dictionary of state age thresholds (e.g. Missouri: 18, New York: 16, Georgia: 18, Texas: 18). The logic then splits into two cases as in (1) and (2):

Suppose the threshold is 18, as in Missouri, Texas, or Georgia. In that case, the reasoning process follows these steps: When the user is 18 or older, they are fully considered an adult, and thus, no parental consent is necessary. The membership values for ages 17 and 18 indicate a partial adult status, representing a borderline scenario. The model analyzes whether parental consent is required and prompts the user to clarify or provide additional information. Finally, if the user's age is clearly below 17, the individual is entirely a minor, requiring explicit parental consent. On the other hand, for states like New York, where the threshold is lower (e.g., 16), if the user's age meets or exceeds this threshold, the user can legally consent without parental involvement. If the age is below the threshold, parental consent remains mandatory. This matches legal rules: Missouri requires guardian consent under 18, whereas New York allows self-consent at 16 or older. The final decision is returned as text, indicating whether consent is needed and why. If the case is borderline, the message notes uncertainty so the system can prompt the user for clarification. Otherwise, the message states clearly that "consent needed" or "no consent needed" based on the rule that fired. Each branch uses the computed fuzzy values. For example, if $\text{adult} > 0$ in an 18-year rule state, we enter the borderline branch (meaning any non-zero adult membership in a state with 18 as threshold leads the system to treat the case as borderline).

Table 1 presents ten illustrative input scenarios evaluated through symbolic substitution into the fuzzy equations and rules. All values are computed analytically, following the inference structure defined in the system design. For each row we list the raw input (age, state), the computed minor and adult values, the strength of the active rule (the relevant membership degree), the decision, and a brief explanation:

Table 1. Sample Input-Output Matrix Demonstrating Fuzzy Consent Logic

Age	State	minor	adult	Rule Strength	Decision	Explanation
14	MO	1.0	0.0	1.0	Consent needed	Missouri law requires consent for anyone under 18.
17.5	MO	0.5	0.5	0.5	Borderline (likely consent)	Age 17.5 is very close to 18 (half-half membership); system flags likely consent.
18	MO	0.0	1.0	1.0	No consent needed	Age 18 or older is an adult, so no parental consent is required.
17	NY	1.0	0.0	1.0	No consent needed	New York allows age ≥ 16 to self-consent; user is 17.

16	NY	1.0	0.0	1.0	No consent needed	Age 16 meets NY's threshold, so user can consent without a parent.
15	NY	1.0	0.0	1.0	Consent needed	Age 15 is below NY's threshold (16), so parental consent is required.
16.5	TX	1.0	0.0	1.0	Consent needed	Texas threshold is 18; age 16.5 (<18) means consent is required.
18	TX	0.0	1.0	1.0	No consent needed	Age 18 qualifies as an adult, so no parental consent is needed.
17.2	GA	0.8	0.2	0.8	Borderline (likely consent)	Georgia threshold is 18. Age 17.2 gives $\mu_{\text{minor}}=0.8$, so system leans toward needing consent.
17	GA	1.0	0.0	1.0	Consent needed	Age 17 is below Georgia's threshold (18), requiring parental consent.

This table demonstrates how varying ages and state laws lead to different fuzzy values and outcomes. For example, borderline ages (17.2, 17.5) show non-zero values for both minor and adult, leading the system to output a “likely needed” consent message. All decisions align with the intended legal rules and illustrate the system’s explainable behavior.

Logical Extension Scenarios: Expanding Fuzzy Inputs

To handle real-world vagueness, we extend our reasoning framework to incorporate two additional fuzzy input types: the type of health issue and case urgency. Patients often describe their conditions in general terms (e.g. “feeling off” or “something’s wrong”), not precise medical terms. Our structured semantic model can incorporate these as fuzzy attributes. This follows our original approach of defining fuzzy membership functions (we did this for age) and using them in inference. By adding these inputs, the system can interpret ambiguous descriptions and produce confidence scores for its conclusions.

Fuzzy Input: Health Issue Type

We define a membership function to show how clearly a patient's description indicates a mental health issue. For example, if the patient mentions words that commonly describe mental health problems, like "depressed," "anxious," or "sad," we count how many of these words appear in the patient's input. Then we divide this count by a fixed number to create a score between 0 and 1, showing how strongly the description relates to mental health. A higher score means the patient’s description more strongly suggests a mental health issue. Concretely, one could use:

$$\mu_{\text{mental}}(I) = \frac{\text{count of mental-health keywords in } I}{N} \quad (3)$$

Likewise, a membership function for physical health issues could count words like "pain", "fever", and "cough" (not shown). In practice, a description might trigger partial membership in both categories. For example, Input: "I have been feeling very anxious and also have some headache". Using the above formula, we find one or two mental-health-related keywords in the description, resulting in a membership score of 0.67. This means the patient's issue is likely related to mental health, but not entirely certain. The system would interpret this as "likely a mental health issue (67% confidence)". The system then applies the corresponding legal rule. For instance, a minor 16 or older in New York can self-consent for mental health services. Suppose a patient says they are 15 and describes depression. In that case, the system computes membership (e.g. 1.0 if "depress" is detected) and then reasons: "This appears to be a mental health issue. In New York, minors 16+ can

consent to mental health care. (Patient is 15, so parental consent is required.)" The membership value acts as a confidence score. It is included in the explanation (e.g. "likely a mental health case (confidence 0.67)"). This builds on our semantic framework by treating "health issue type" as another structured concept with fuzzy values, just as we did for age. It shows the system handling vagueness in patient descriptions: instead of guessing, it quantifies ambiguity and then follows the explicit rule, keeping the reasoning transparent.

Fuzzy Input: Emergency Status

Another source of ambiguity is urgency. A patient might say "I'm in pain" without specifying how severe or when it started. Yet legal rules often permit immediate care in emergencies. For example, if the patient provides a pain scale (0–10), we can define:

$$\mu_{emergency}(pain) \begin{cases} 0 & , if \text{ pain} \leq 5 \\ \frac{\text{pain}-5}{5} & , if 5 < \text{pain} < 10 \\ 1 & , if \text{ pain} \geq 10 \end{cases} \quad (4)$$

This gives an emergency membership score between 0 and 1. For instance, Input: "I'm 17, from Missouri, with chest pain rated 8/10". The system calculates $\mu_{emergency}(8) = (8-5)/5 = 0.6$. It interprets this as moderately high urgency (60% confidence). Then we can define a rule like this: if the emergency membership score $\mu_{emergency}$ is greater than 0.5, treat as an emergency. The fuzzy inference might output: "Warning: Possible emergency (confidence 0.6). Please seek immediate care." Legally, this aligns with emergency exceptions, for example many states allow treating minors urgently without waiting for parental consent (World Population Review, n.d.). Here the system can override the normal consent rule and advise the user accordingly. The membership score itself serves as a confidence measure, and the output text explicitly mentions urgency and the reason. This ensures the decision remains explainable, aligning with our design goals stated earlier: we handle ambiguity (unclear symptom severity) by a graded score, and the assistant explains the rationale ("chest pain 8/10 suggests an emergency, so immediate care is recommended").

In both cases, adding these fuzzy inputs builds on the same structured semantic modeling framework used earlier. We simply extend our semantic frame for the patient's query to include, e.g., `health_issue_type` and `urgency` slots, each with fuzzy values. The inference engine then combines these with existing facts (age, location, etc.) to produce a decision. Importantly, the system still generates an explicit reasoning path: it reports the fuzzy memberships (as confidence scores) and the legal rule applied. This supports explainability and meets the objectives outlined in our introduction. The assistant can now say, for example, "Based on your symptoms and state law, parental consent is not needed: this appears to be a mental health issue (confidence 0.7) and emergencies allow immediate treatment". In summary, these extensions show that our model can flexibly incorporate new types of ambiguity while maintaining clear, justifiable outputs.

Results and Discussion

We provide a proof-of-concept logical framework, expressed entirely through the membership equations (1) – (4), to show how the system resolves three frequent sources of vagueness in telehealth encounters: indeterminate patient age, loosely worded symptom descriptions, and unclear urgency levels. Unlike conventional compliance checkers—which depend on precise, fully declared patient data or perform retrospective audits—our approach weaves fuzzy reasoning into the decision core itself. By substituting test values into the published equations and rules, we generated Table 1, which documents the model's consent decisions for ten representative scenarios.

The results confirm that the model can deliver legally sound guidance even when patient inputs are uncertain or incomplete.

From our results, it became clear that the fuzzy inference module effectively handled uncertain patient age descriptions. For example, when patients described their age vaguely ("almost 18" or "about seventeen and a half"), the fuzzy logic engine calculated intermediate membership values such as 0.5 for both minor and adult categories, clearly signaling uncertainty. Instead of immediately categorizing patients into strictly minor or adult groups, the model responded cautiously by requesting additional user clarification. This cautious approach highlights one of our structured model's main strengths, as it integrates semantic ambiguity directly into compliance reasoning which is rarely achieved by traditional binary compliance systems.

In the case of ambiguous issue type descriptions, our fuzzy semantic approach assigned partial confidence values indicating uncertainty about whether the issue belonged to mental health categories. As a result, the model carefully included the estimated uncertainty scores and the applied state rules. These detailed explanations resemble human reasoning more closely than those of typical compliance tools, which helps patients better understand and trust the advice. Additionally, we evaluated our AI system's capability to handle descriptions related to urgency or emergencies. When patients described symptom severity ambiguously, such as "pain around 8 out of 10," the fuzzy logic produced a numerical confidence score representing partial urgency. This enabled the model to recognize emergency cases clearly and advise patients to seek immediate care, explicitly connecting the response to emergency-related legal exceptions.

Our results indicate that embedding fuzzy logic and structured semantic modeling into the compliance system offers important practical benefits. Instead of using simple binary rules like most compliance-checking methods do, our fuzzy-based approach actively includes semantic ambiguity into its logic. This lets the model give careful, clear, and human-like answers.

Future work

While the proposed model was developed with telehealth compliance in mind, the underlying reasoning structure may also be applicable to other regulated environments. This structured reasoning approach could extend beyond telehealth. One example is interactive learning tools for children with autism. These systems often require age-appropriate content control, parental consent verification, and adaptive explanations tailored to emotional or developmental needs. By embedding a reasoning core similar to our framework, a learning board could dynamically decide whether to unlock a learning module, pause interaction based on emotional cues, or switch content delivery styles depending on the user role (child, parent, or teacher). Such integration could improve both the legal safety and emotional responsiveness of educational tools used in special education.

Conclusion

This paper presented a modular AI framework that embeds legal reasoning into the core logic of telehealth systems, enabling real-time, explainable, and context-aware compliance decisions. Through structured application of fuzzy membership functions and state-specific legal rules, the system demonstrated its ability to navigate three common sources of ambiguity in remote healthcare encounters: uncertain patient age, imprecise health issue descriptions, and vague indicators of urgency. The framework produces graded outputs rather than binary classifications, offering confidence-based recommendations and clarifying borderline legal situations through transparent reasoning pathways.

By formalizing compliance as an internal decision-making process—rather than an external checklist—this model better mirrors the nuanced judgment required in real-world clinical settings. Its semantic flexibility allows for adaptive interpretation of patient inputs,

while its rule-based structure ensures that outputs remain grounded in statutory requirements. The inference tables and formula-based analyzes confirm that the system reaches consistent, legally sound conclusions across varying jurisdictions and input patterns.

Beyond telehealth, the architecture shows promise for broader deployment in regulated domains such as special education or interactive learning environments for neurodiverse users. As demonstrated in the case of consent-aware learning boards for autistic children, the same reasoning engine could manage age thresholds, emotional cues, and caregiver roles, ensuring legal and ethical safeguards while promoting user autonomy.

Overall, this work contributes a novel approach to AI-driven compliance: one that is semantically rich, mathematically grounded, and pragmatically designed for integration into sensitive, high-stakes systems. Future development will focus on scaling the rule base, validating outputs through human-centered studies, and extending the framework to additional forms of regulatory reasoning beyond consent logic.

References

- Bisschoff, I., & Mittelstädt, B. (2020). Explainability for artificial intelligence in healthcare. *BMC Medical Informatics and Decision Making*, 20, 173. <https://doi.org/10.1186/s12911-020-01332-6>
- Cleverbot. (n.d.). *About Cleverbot*. <https://www.cleverbot.com/>
- Garber, K. M., & Chike-Harris, K. E. (2019). Nurse practitioners and virtual care: A 50-state review of APRN telehealth law and policy. *Telehealth and Medicine Today*, 4, e136. <https://doi.org/10.30953/tmt.v4.136>
- Hameed, K., Bajwa, I. S., Ramzan, S., Anwar, W., & Khan, A. (2020). *An intelligent IoT-based healthcare system using fuzzy neural networks*. *Scientific Programming*, 2020, 1-15. <https://doi.org/10.1155/2020/8836927>
- Ivanova, J., Cummins, M. R., Ong, T., Soni, H., Barrera, J., Wilczewski, H., Welch, B., & Bunnell, B. (2025). *Regulation and compliance in telemedicine: Viewpoint*. *Journal of Medical Internet Research*, 27, e53558. <https://doi.org/10.2196/53558>
- Kuziemy, C., Maeder, A. J., John, O., Gogia, S. B., Basu, A., & Meher, S. (2019). *Role of artificial intelligence within the telehealth domain*. *Yearbook of Medical Informatics*, 28(1), 35-40. <https://doi.org/10.1055/s-0039-1677897>
- Missouri Revised Statutes § 632.070 (2014). *Rights of minors to treatment—consent requirements*. State of Missouri, Revisor of Statutes. <https://revisor.mo.gov/main/OneSection.aspx?section=632.070>
- Mukhopadhyay, S., Mukherjee, J., Deb, D., & Datta, A. (2025). *Learning fuzzy decision trees for predicting outcomes of legal cases....* *Applied Soft Computing*, 176, 113179. <https://doi.org/10.1016/j.asoc.2025.113179>
- New York Consolidated Laws, Mental Hygiene Law § 9.13 (2024). <https://www.nysenate.gov/legislation/laws/MHY/9.13>
- Reddy, S. (2022). Explainability and artificial intelligence in medicine. *The Lancet Digital Health*, 4(4), e214-e215. [https://doi.org/10.1016/S2589-7500\(22\)00029-2](https://doi.org/10.1016/S2589-7500(22)00029-2)
- Safdar, K. (2022, September 29). Cerebral treated a 17-year-old without his parents' consent. They found out the day he died. *The Wall Street Journal*. <https://www.wsj.com/articles/cerebral-treated-a-17-year-old-without-his-parents-consent-they-found-out-the-day-he-died-11664416497>
- Sharko, M., Jameson, R., Ancker, J. S., Krams, L., Webber, E. C., & Rosenbloom, S. T. (2022). *State-by-state variability in adolescent privacy laws*. *Pediatrics*, 149(6), e2021053458. <https://doi.org/10.1542/peds.2021-053458>
- Taine-Crasta, E., & Imran, A. (2020). Use of telehealth during the COVID-19 pandemic: Scoping review. *Journal of Medical Internet Research*, 22(12), e24087. <https://doi.org/10.2196/24087>
- World Population Review. (n.d.). *Age of Consent for Mental Health Treatment by State 2025*. [worldpopulationreview.com. https://worldpopulationreview.com/state-rankings/age-of-consent-for-mental-health-treatment-by-state](https://worldpopulationreview.com/state-rankings/age-of-consent-for-mental-health-treatment-by-state)
- Yorita, A., Egerton, S., Chan, C., & Kubota, N. (2023). Chatbots and robots: A framework for the self-management of occupational stress. *ROBOMECH Journal*, 10, 24. <https://doi.org/10.1186/s40648-023-00261-z>