

Adaptive Collaboration of Arena-Based Argumentative LLMs for Explainable and Contestable Legal Reasoning

Hoang-Loc Cao^{†*§}, Phuc Ho^{†§}, Truong Thanh Hung Nguyen[‡],
Phuc Truong Loc Nguyen[°], Dinh Thien Loc Nguyen[†], Hung Cao[‡]
[†]University of Science, VNU-HCM, Ho Chi Minh City, Vietnam
[‡]University of New Brunswick, Canada
[°]Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany

Abstract

Legal reasoning requires not only high accuracy but also the ability to justify decisions through verifiable and contestable arguments. However, existing Large Language Model (LLM) approaches, such as Chain-of-Thought (CoT) and Retrieval-Augmented Generation (RAG), often produce unstructured explanations that lack a formal mechanism for verification or user intervention. To address this limitation, we propose *Adaptive Collaboration of Argumentative LLMs (ACAL)*, a neuro-symbolic framework that integrates adaptive multi-agent collaboration with an Arena-based Quantitative Bipolar Argumentation Framework (A-QBAF). ACAL dynamically deploys expert agent teams to construct arguments, employs a clash resolution mechanism to adjudicate conflicting claims, and utilizes uncertainty-aware escalation for borderline cases. Crucially, our framework supports a Human-in-the-Loop (HITL) contestability workflow, enabling users to directly audit and modify the underlying reasoning graph to influence the final judgment. Empirical evaluations on the LegalBench benchmark demonstrate that ACAL outperforms strong baselines across Gemini-2.5-Flash-Lite and Gemini-2.5-Flash architectures, effectively balancing efficient predictive performance with structured transparency and contestability. Our implementation is available at: <https://github.com/loc110504/ACAL>.

Keywords: Legal Reasoning, Large Language Models, Multi-Agent Systems, Contestable AI, Computational Argumentation, Neuro-Symbolic AI

1. Introduction

Legal reasoning is the intricate process by which legal principles and rules are applied to specific facts to reach a justified outcome [1, 2]. This professional responsibility encompasses various essential sub-tasks such as identifying pertinent issues, recalling relevant laws, and interpreting complex statutes or precedents [2, 3]. Practitioners in this field must apply established rules to specific facts, draw logical conclusions, and present persuasive arguments to support their claims. In recent years, Large Language Models (LLMs) have demonstrated remarkable progress on these legal reasoning tasks. Researchers have shown that equipping these models with domain knowledge and specific reasoning strategies can significantly boost performance [3–5]. Recently, researchers have also begun exploring multi-agent LLM frameworks that simulate debates or judicial deliberations to tackle decision-making [6, 7]. By having multiple agents critique and refine each other’s arguments, these approaches can improve factual accuracy and overall robustness [8, 9].

However, existing approaches to legal reasoning with LLMs fall short of these requirements in different ways. Prompting-based methods, including few-shot prompting and chain-of-thought (CoT) reasoning [3, 4, 10], often produce free-form explanations that are difficult to verify or systematically contest. Retrieval-augmented generation (RAG) improves factual grounding by injecting external legal texts [7, 10, 11], but the final decision logic remains implicit and opaque. Multi-agent debate (MAD) frameworks introduce diversity of perspectives, yet their performance gains are inconsistent across tasks, and debate transcripts themselves do not constitute a formally contestable reasoning structure.

Furthermore, frontier models often produce invalid arguments or irrelevant citations when facing complex reasoning demands. This limitation raises concerns about trust, fairness, and

[§]These authors contributed equally.

*Corresponding author: chloc22@clc.fitus.edu.vn

accountability in high-stakes domains. Meanwhile, regulators increasingly require transparency and contestability as reflected in global directives such as the EU AI Act [12], and Canada’s Directive on Automated Decision-Making [13]. Consequently, prior approaches either optimize predictive performance without accountability or provide explanations that lack a principled mechanism for inspection and dispute.

To address these challenges, we propose *Adaptive Collaboration of Argumentative LLMs (ACAL)*, a structured, neuro-symbolic, multi-agent AI framework for explainable, contestable legal reasoning. Our contributions are summarized as follows:

- (1) We propose ACAL, an adaptive multi-agent legal reasoning framework that integrates argumentative LLMs (ArgLLMs) [14] with adaptive role selection in *arena-based (quantitative) bipolar argumentation (A-(Q)BAF)* framework, enhanced by *clash resolution (CR)* and *uncertainty-aware escalation (UAE)* to yield more accurate and decisive legal judgments.
- (2) Our proposed ACAL supports the human-in-the-loop (HITL) contestation workflow, enabling users to directly interrogate and revise the underlying reasoning process, with changes formally propagated to influence the final outcome.
- (3) We empirically demonstrate on LegalBench [guha2023legalbench] that ACAL outperforms strong prompting, RAG, CoT, and MAD baselines in predictive performance, while simultaneously providing structured transparency and contestability not supported by prior methods.

2. Related Work

2.1. LLM-based Approaches for Legal Reasoning

The application of AI in the legal domain has transitioned from early expert systems to modern data-driven approaches leveraging LLMs. Specialized models such as ChatLaw [15] and LegalMind [16] have refined these capabilities by integrating knowledge graphs and reinforcement learning to improve factuality and process optimization. Foundational techniques like RAG are now standard for grounding LLM responses in authoritative external sources. Systems use RAG to query everything from ethical rules [11] and tax codes [10] to complex statutory frameworks and case law precedents [17]. Concurrently, CoT prompting is widely used to elicit multi-step reasoning, with some frameworks developing Law-specific CoT variants [4] or logical-semantic integration models [5] to cultivate more robust deductive pathways. To further improve robustness, researchers have introduced multi-agent architectures in which agents assume distinct roles, such as judge, plaintiff, and defendant, to simulate adversarial legal discourse [9]. These systems, which range from simulating courtroom debates [18] to collaborative law-making [19], leverage iterative critique to refine arguments and reach a consensus.

However, despite these advancements, significant performance gaps remain. Recent studies highlight a “reasoning paradox” in which models that consume more computational resources on hierarchical tasks often degrade in performance, struggling to distinguish between surface-level facts and deeper legal distinctions [20]. Furthermore, while these advanced methods can improve accuracy over single-model baselines [6], the debate transcripts or free-form explanations they produce lack a formal structure. This makes their reasoning difficult to systematically verify or challenge, creating a critical accountability gap that purely performance-driven architectures fail to address.

2.2. Explainability and Contestability in LLM-based Legal Reasoning Systems

Explainability in legal AI is fundamentally concerned with providing rational justification for stakeholders, a standard that is distinct from the technical goal of describing a model’s

internal mechanics. Current frameworks attempt to provide this justification by making their reasoning process transparent. For instance, systems grounded in RAG offer explainability by providing a citation mechanism that links claims to specific source documents [7, 21]. Other approaches utilize CoT techniques to generate structured, step-by-step reasoning chains that mimic judicial logic [17, 20]. More advanced systems even formalize interactions into computational argumentation graphs using frameworks like Toulmin’s model to create a verifiable reasoning trail [9, 22]. However, the presence of these explanatory artifacts does not guarantee their reliability or provide a mechanism for recourse. One recent evaluation of frontier models found that over 60% of generated judicial analyses contained invalid arguments and more than half included irrelevant citations, despite their structural coherence [17]. Consequently, there is a growing demand to move beyond passive explainability towards contestability, an interactive paradigm that empowers users to actively challenge and correct the model’s reasoning process rather than merely observing it [11, 22].

Contestable AI (CAI) extends the principles of XAI by incorporating safeguards and providing explicit pathways for users to challenge and revise a system’s conclusions [23, 24]. The integration of computational argumentation has become central to operationalizing contestability, specifically within legal AI. Frameworks such as ALEX [22] employ formal schemes, such as Toulmin’s model and the ASPIC+ [22] framework, to structure interactions between opposing arguments in criminal cases. This allows users to scrutinize specific nodes in a reasoning graph. Similarly, neuro-symbolic approaches such as SOLAR [8] leverage structured ontological representations to bridge the gap between natural language and symbolic reasoning for statutory analysis. This provides verifiable pathways that can withstand adversarial probing. Other work has explored contestability through structured multi-ply dialectical schemes for case-based reasoning [25] or by deriving policy guidance from regulatory mandates to help users identify the right questions to ask [21]. However, these existing legal CAI approaches primarily focus on generating static reasoning structures for post-hoc inspection. They typically lack a quantitative mechanism to dynamically weigh the strength of conflicting arguments. Furthermore, they do not support a fully interactive workflow where human interventions, such as modifying argument weights or relations, mathematically propagate to update the final legal judgment.

Our work bridges this gap by integrating MAD within a formal quantitative argumentation framework that is directly auditable by an HITL. This transforms static explanations into dynamic and contestable decision objects.

3. Adaptive Collaboration of Argumentative LLMs (ACAL)

As shown in Figure 1, ACAL comprises four modules: adaptive expert team selection, multi-agent argument generation with arena-based clash resolution, quantitative reasoning via A-QBAF, and HITL contestation with uncertainty-aware decision consensus.

3.1. Aspect Identification and Team Selection

Legal reasoning faces many challenges from different domains that require specialized knowledge. A single expert might not be able to capture all perspectives in a complex case that needs different professionals to analyze it effectively. To deal with this limitation, we propose an adaptive multi-agent architecture that dynamically assembles expert teams tailored to specific legal tasks.

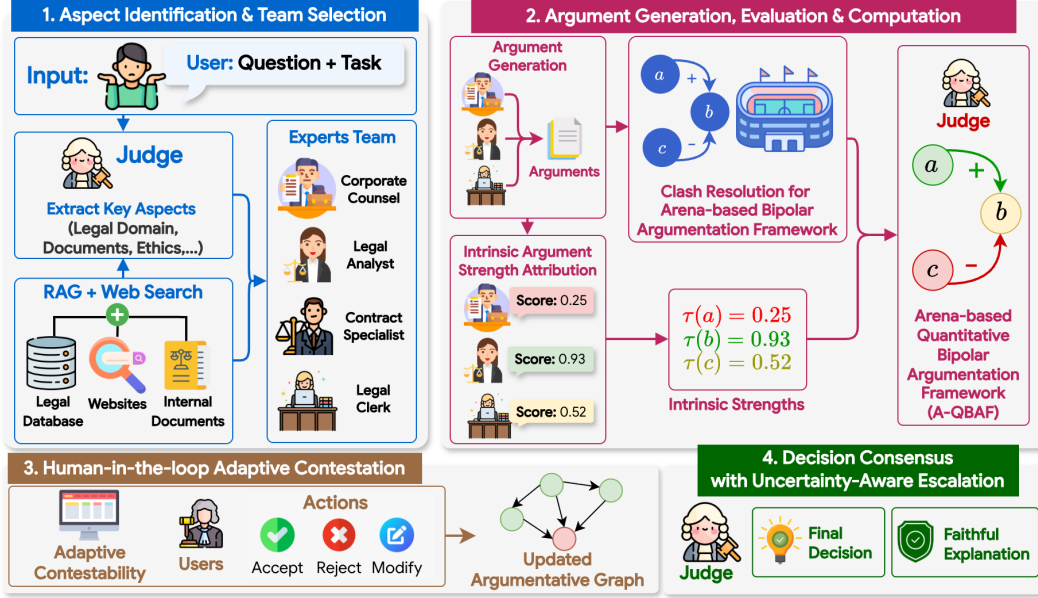


Figure 1. Our proposed Adaptive Collaboration of Argumentative LLMs (ACAL) Architecture for Legal Reasoning.

3.1.1. Legal Agent Pool Definition

We define a comprehensive pool of legal agents $\mathcal{A} = \{a_1, a_2, \dots, a_n\}$, where each agent a_i is characterized by a tuple r_i, E_i, P_i, S_i as follows: (i) *Role* r_i : The professional designation; (ii) *Expertise Areas* E_i : A set of domain specializations; (iii) *Focus Priorities* P_i : Task-oriented objectives that guide reasoning; (iv) *Argument Style* S_i : The characteristic reasoning approach. As shown in Table 1, our framework implements 10 distinct legal roles organized into functional categories [26].

Table 1. Legal roles grouped by functional category.

Category	Roles
Adjudication	Judge, Law Clerk / Judicial Clerk
Litigation & Advocacy	Private Practice Lawyer, Prosecutor, Public Defender
Advisory & Transactional	Corporate Counsel, Compliance Officer, IP Attorney
Research & Support	Legal Analyst, Paralegal

3.1.2. Adaptive Agent Selection

Rather than employing all agents for every case, which would be computationally expensive and potentially introduce noise when using a non-expert for a specific task, we implement an adaptive selection mechanism. Given a legal task \mathcal{T} with context c and claim ϕ , the system selects two subsets:

$$\mathcal{A}^+ = \text{Select}(\mathcal{A}, \mathcal{T}, c, \text{support}); \quad \mathcal{A}^- = \text{Select}(\mathcal{A}, \mathcal{T}, c, \text{attack}), \quad (3.1)$$

where \mathcal{A}^+ are agents whose expertise aligns with constructing supporting arguments for that specific case context, \mathcal{A}^- are agents suited for generating counter-arguments. The selection function leverages an LLM to match agent expertise profiles against case characteristics.

3.2. Argument Generation, Evaluation and Computation

The core of our framework lies in the structured generation and computation of argument strengths using A-QBAF, which is built upon traditional QBAF [27].

3.2.1. Contextual Grounding via Hybrid RAG

To ensure arguments are grounded in authoritative legal standards rather than generic knowledge, we integrate a hybrid retrieval module. Prior to argument generation, the system queries both a vectorized legal database for legislative texts and an external web search for recent case law. The top- k most semantically relevant passages are aggregated to form the evidentiary context c , which conditions the subsequent reasoning of all agents to minimize hallucination and ensure accuracy.

3.2.2. Multi-Agent Argument Generation

Each selected agent generates arguments directly addressing the central claim ϕ in the context c . For an agent $a \in \mathcal{A}^+ \cup \mathcal{A}^-$:

$$\text{Args}(a) = \text{LLM}(a, \phi, c, \mathcal{C}, \text{type}(a)), \quad (3.2)$$

where $\text{type}(a) \in \{\text{support}, \text{attack}\}$ determines whether the agent constructs evidence that the claim is $\text{True}(\text{support})$ or $\text{False}(\text{attack})$. Each agent typically generates 2-5 arguments, with the quantity self-determined based on available evidence, the complexity of the case.

3.2.3. Intrinsic Strength Attribution

Each argument α_i receives an intrinsic strength score $\tau(\alpha_i)$ from LLM. We designed a scoring criterion (see Table 2) that enforces strict differentiation to prevent score saturation, requiring evaluators to penalize generic statements and reward case-specific legal precision.

Table 2. LLM-based argument strength scoring rubric.

Score Range	Interpretation
0.1 – 0.2	Incorrect legal analysis or misidentification of key elements
0.3 – 0.4	Partially correct but missing critical components or overly generic
0.5 – 0.6	Sound analysis with minor gaps or insufficient case-specific application
0.7 – 0.8	Strong analysis with specific facts and correct legal reasoning
0.9 – 1.0	Exceptional precision with authoritative citations and flawless logic

3.2.4. LLM-based Inter-Argument Relation Identification

Before constructing the A-QBAF graph, we must determine the pairwise relations \mathcal{R}^- and \mathcal{R}^+ among arguments. In heuristic mode, it follows a simple rule whereby arguments sharing the same stance (both support or both attack) are treated as mutual supporters, while arguments with opposing stances are treated as mutual attackers. However, this oversimplifies legal reasoning in which two arguments on opposite sides may address entirely different legal aspects, making them logically independent rather than conflicting.

To build a better logical structure, we implemented LLM-based semantic relation identification mode. For each pair (α_i, α_j) with $i < j$, we prompt an LLM to classify the relationship into one of three categories attack, support or neutral, along with a confidence score. To ensure robustness, any support or attack relation with confidence score < 0.6 is demoted to neutral. All established relations are made bidirectional, ensuring symmetric

treatment in the subsequent QBAF propagation. Neutral pairs introduce no edge, keeping the graph sparse and interpretable. For efficiency, argument pairs are analyzed in batches of b pairs per LLM call (default $b = 10$), reducing the number of API calls from $\binom{n}{2}$ to $\lceil \frac{\binom{n}{2}}{b} \rceil$

3.2.5. Clash Resolution via Arena Debating Round

When opposing arguments have similar base scores (difference $< \delta$, default $\delta = 0.2$), we propose a *clash resolution* (CR) mechanism that adjudicates between them, as illustrated in Figure 2. For each conflicting pair (α_s, α_a) where α_s supports and α_a attacks the claim:

- (1) Present both arguments to an LLM acting as a legal reasoning expert.
- (2) Evaluate which argument is stronger, given case-specific facts and legal standards.
- (3) Adjust scores based on win/loss outcomes across all clashes.

The adjustment follows a symmetric formula based on the win rate w :

$$\Delta\tau(\alpha) = \beta \cdot (2w - 1), \quad (3.3)$$

where β is the base adjustment magnitude and $w \in [0, 1]$ is the proportion of clashes won. This ensures winners receive bonuses while losers receive proportional penalties, maintaining score differentiation.

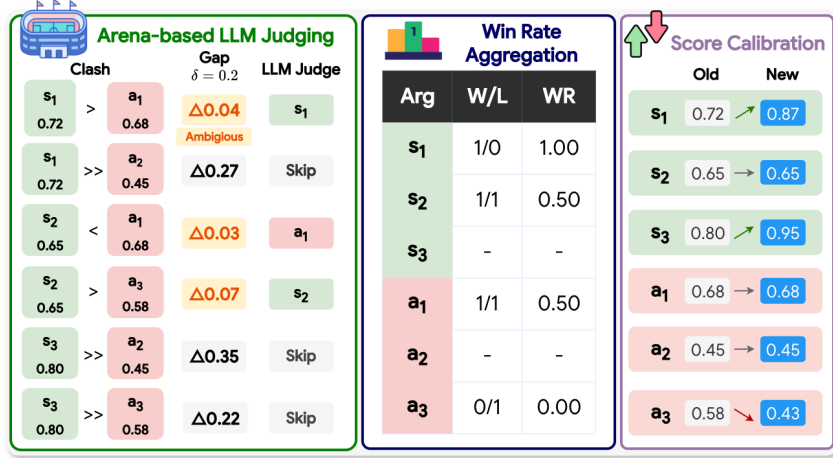


Figure 2. Illustration of an Arena Debating Round through Clash Resolution (CR).

3.2.6. Arena-based Quantitative Bipolar Argumentation Framework (A-QBAF)

We formalize the argumentation structure as a QBAF $\langle \mathcal{X}, \mathcal{R}^-, \mathcal{R}^+, \tau \rangle$ where:

- $\mathcal{X} = \{\phi\} \cup \{\alpha_1, \dots, \alpha_n\}$: The set of arguments including the central claim ϕ ,
- $\mathcal{R}^- \subseteq \mathcal{X} \times \mathcal{X}$: Attack relations,
- $\mathcal{R}^+ \subseteq \mathcal{X} \times \mathcal{X}$: Support relations,
- $\tau : \mathcal{X} \rightarrow [0, 1]$: Base strength function.

Graph Construction: The central claim node ϕ is initialized with a base score $\tau(\phi) = 0.5$ (neutral). Each argument α_i connects to ϕ via:

- $(\alpha_i, \phi) \in \mathcal{R}^+$ if α_i is a support argument,
- $(\alpha_i, \phi) \in \mathcal{R}^-$ if α_i is an attack argument.

We compute final argument strengths using the *Quadratic Energy (QE)* semantics [28], which models score propagation as a continuous dynamical system to ensure convergence and axiomatic stability, as follows:

Energy Calculation: For an argument $j \in \mathcal{X}$, the total energy E_j is the linear sum of the current strengths of its supporters minus its attackers

$$E_j = \sum_{i \in \text{Sup}_j} \sigma_i - \sum_{i \in \text{Att}_j} \sigma_i, \quad (3.4)$$

where σ_i represents the strength of argument i at time t .

Impact Function: The impact of energy is governed by a quadratic function $h(x)$, which is continuously differentiable and strictly monotonic for positive values:

$$h(x) = \frac{\max\{x, 0\}^2}{1 + \max\{x, 0\}^2}, \quad (3.5)$$

Equilibrium Strength: The final propagated strength σ_j^* (also denoted as $\sigma(j)$ in subsequent sections) is the equilibrium state $\sigma^* = \lim_{t \rightarrow \infty} \sigma(t)$. At equilibrium, the strength satisfies:

$$\sigma_j^* = \tau(j) + (1 - \tau(j)) \cdot h(E_j) - \tau(j) \cdot h(-E_j), \quad (3.6)$$

where $\tau(j)$ is the argument’s initial base weight. This ensures that attacks and supports have a symmetrical (dual) impact on the argument’s strength.

3.3. Human-in-the-loop (HITL) Contestation

While structured argumentation improves transparency, legal decision-making additionally requires *contestability*: affected parties and legal professionals must be able to challenge the system’s reasoning, correct errors, and meaningfully influence the outcome. Accordingly, we introduce a role-based HITL contestation module that turns our A-QBAF graph from a static explanation into an editable, revisable decision object.

3.3.1. Contestation Artifacts

Given the argumentation structure $\langle \mathcal{X}, \mathcal{R}^-, \mathcal{R}^+, \tau \rangle$ and the computed degrees $\sigma(\cdot)$, the system generates a *participation summary* and an *argument dashboard*. Each argument card displays: (i) the argument text and stance (support/attack); (ii) originating agent role r_i ; (iii) links to the specific evidence passages used (statutes, precedents, case facts); (iv) the intrinsic score $\tau(\alpha)$ and propagated score $\sigma(\alpha)$; and (v) its local neighborhood in the A-QBAF (incoming supporters/attackers), plus a short “why it matters” trace describing its influence on $\sigma(\phi)$.

3.3.2. Unified Contestation Workflow

To support contestability in a single, user-facing workflow, the interface provides a guided contestation process centered on *what is being challenged* in the system’s reasoning: (a) *factual issues* (incorrect or missing facts/evidence), (b) *legal rule issues* (misapplied statute/standard or wrong test), (c) *precedent issues* (irrelevant/overruled/misread authority), (d) *missing exceptions/defenses* (e.g., consent, necessity, limitation period), and (e) *procedural fairness concerns* (lack of notice, opportunity to respond, or mitigation). Selecting a contestation type triggers targeted prompts that either (1) request additional supporting materials from the user or (2) instruct the relevant agents to regenerate, refine, or rebut arguments under the user’s contestation claim. All resulting updates are reflected in the editable argument record (arguments, relations, and scores) with an auditable change log, ensuring contestations can materially affect the propagated claim score rather than remaining as post-hoc feedback.

3.3.3. Allowed Human Interventions

During contestation, users can provide feedback on the argument record by (i) accepting or rejecting arguments; (ii) editing arguments to correct facts or refine the legal rationale; and (iii) adding missing arguments (with citations when available). Users may also suggest adjustments to an argument’s strength or its support/attack relation when the structure is clearly mis-specified.

Update and recomputation. All accepted edits are applied to yield an updated A-QBAF $(\mathcal{X}^H, \mathcal{R}^{+H}, \mathcal{R}^{-H}, \tau^H)$, after which we re-run the same propagation semantics to obtain updated scores $\sigma^H(\cdot)$ and the revised claim score $\sigma^H(\phi)$. Thus, contestation is treated as a model input that can change the final decision rather than a post-hoc annotation.

Oversight and logging. For high-uncertainty or high-impact cases, the system can trigger an additional review step, and all contestation actions are recorded in an audit log (who changed what, and how it affected $\sigma(\phi)$) to support traceability.

3.4. Decision Consensus with Uncertainty-Aware Escalation

The final answer is derived from the claim score:

$$\text{Answer} = \begin{cases} \text{Yes,} & \text{if } \sigma(\phi) \geq \theta, \\ \text{No,} & \text{otherwise,} \end{cases} \quad (3.7)$$

where θ is the decision threshold (default $\theta = 0.5$). However, in practice, the initial LLM-based base scoring may yield near-ties between support and attack arguments (e.g., both appear comparably plausible), causing the propagated claim score to collapse toward the neutral region. Thus, threshold-based decisions near the boundary risk becoming arbitrary, with outcomes driven by marginal scoring noise rather than meaningful legal distinctions. Accordingly, prior work [29] recommends deferring or abstaining from predictions in high-uncertainty situations, often by escalating to a stronger model or a human reviewer.

Therefore, we propose the *uncertainty-aware escalation (UAE)* method for borderline cases where $\sigma(\phi) \in [0.49, 0.51]$ ($\sigma(\phi) \approx 0.5$) by bypassing the base-score-limited decision rule and invoking a *Final Judge* agent. The Judge performs an *independent* legal analysis of the case (re-evaluating evidence, legal standards, and resolving key conflicts) and outputs a *binding* decision. This mechanism specifically addresses the same-score saturation issue between attack vs. support arguments and ensures a decisive outcome under high uncertainty.

4. Experiment Setup

4.1. Dataset

We evaluate our approach on **LegalBench** [guha2023legalbench], a collaboratively built benchmark comprising a diverse suite of legal reasoning tasks designed for in-context evaluation of large language models. In this work, we focus on two classification tasks:

- **Hearsay** (`hearsay`): given a short statement, predict whether it is hearsay under FRE 801(c) (`hearsay` vs. `not_hearsay`).
- **Courts** (`learned_hands_courts`): given a real-world narrative (e.g., a forum-style post), predict whether it belongs to the `courts` category in the Learned Hands taxonomy (`yes` / `no`).

4.2. Benchmark

We conduct all experiments using two Google models: Gemini-2.5-Flash-Lite and Gemini-2.5-Flash. For the retrieval-augmented baseline, we implement a standard RAG pipeline in

which documents are embedded with OpenAI `text-embedding-3-large`, extracting the top- k ($k = 5$) most relevant chunks to provide additional context to the generator before making a prediction. Regarding our proposed framework, we set the base adjustment magnitude $\beta = 0.15$ for the clash resolution mechanism, a value determined empirically to yield optimal performance.

We compare against commonly-used prompting and multi-agent baselines for legal and general reasoning: **SP** [30] (few-shot standard prompting) applies a fixed instruction template with a small set of in-context demonstrations; **CoT** [31] encourages intermediate reasoning steps before producing the final label; **RAG** [32] augments the input with retrieved legal evidence (e.g., relevant statutory or doctrinal snippets) and conditions predictions on this external context; and **MAD** [33] uses a multi-agent debate setting with three agents over two rounds of critique and revision to reach a final decision.

For both tasks, we report Accuracy (Acc), Precision (Prec), Recall (Rec), and Macro-F1 (F1). Macro-F1 is computed by averaging class-wise F1 scores, ensuring equal weight across classes and making the evaluation more robust to class imbalance.

5. Results

Table 3. Comparative results on Learned Hands Courts and Hearsay from LegalBench. All metrics are reported as percentages. The highest scores are in **bold**.

Model	Method	Learned Hands Courts				Hearsay			
		Acc	Prec	Rec	F1	Acc	Prec	Rec	F1
Gemini-2.5-Flash-Lite	SP	57.8	63.1	57.8	53.1	69.2	73.6	71.5	68.9
	CoT	69.3	69.3	69.2	69.2	69.2	74.7	71.8	68.7
	RAG	70.3	70.5	70.3	70.3	71.3	70.8	70.9	70.9
	MAD	69.8	69.8	69.8	69.8	74.5	75.8	72.4	72.7
	ACAL (Ours)	70.8	71.2	70.8	70.7	74.5	77.1	76.3	74.4
Gemini-2.5-Flash	SP	65.1	72.5	65.1	62.0	75.5	75.7	74.2	74.5
	CoT	72.3	74.1	72.6	71.3	77.2	77.4	75.6	75.9
	RAG	75.0	78.1	75.0	74.3	75.5	79.5	72.8	73.0
	MAD	75.5	76.8	75.5	75.2	77.6	78.5	76.1	76.5
	ACAL (Ours)	75.5	76.4	75.5	75.3	76.7	77.1	77.6	76.7

5.1. Quantitative Analysis

Table 3 presents the comparative results of our proposed framework, ACAL, against four baselines (SP, CoT, RAG, and MAD) on the *Learned Hands Courts* and *Hearsay* tasks from LegalBench. We evaluate performance across two backbone models: Gemini-2.5-Flash-Lite and Gemini-2.5-Flash. As shown in the table, ACAL demonstrates superior or highly competitive performance across both model architectures.

Gemini-2.5-Flash-Lite: In the resource-constrained setting, ACAL achieves the best overall performance. On the *Learned Hands Courts* dataset, our method surpasses all baselines, achieving the highest Accuracy (70.8%) and F1-score (70.7%). Similarly, on the *Hearsay* dataset, ACAL outperforms the strongest baseline (MAD) in Precision (77.1%), Recall (76.3%), and F1-score (74.4%), while matching the highest Accuracy (74.5%).

Gemini-2.5-Flash: Scaling to the larger model, ACAL maintains its robustness. On *Learned Hands Courts*, ACAL matches the top accuracy of the Multi-Agent Debate (MAD)

Table 4. Ablation study of our two proposed modules, Clash Resolution (CR) and Uncertainty-Aware Escalation (UAE), on Gemini-2.5-Flash-Lite (Hearsay).

CR	UAE	Acc	Prec	Rec	F1
✗	✗	64.9	69.8	67.5	64.4
✗	✓	62.8	72.5	66.4	61.2
✓	✗	72.8	75.1	74.5	72.8
✓	✓	74.5	77.1	76.3	74.4

Table 5. Ablation study of parameter β (base adjustment magnitude) on Gemini-2.5-Flash-Lite (Learned Hands Courts).

β	Acc	Prec	Rec	F1
0.05	69.2	69.2	69.1	69.1
0.10	70.0	70.4	69.9	69.8
0.15	70.8	71.2	70.8	70.7
0.20	70.2	70.5	70.2	70.0
0.25	69.1	69.6	68.1	68.0

baseline (75.5%) and achieves a superior F1-score (75.3%) compared to RAG (74.3%). Notably, on the *Hearsay* dataset, ACAL achieves the highest Recall (77.6%) and F1-score (76.7%) among all compared methods. These results indicate that ACAL consistently outperforms standard prompting methods and remains competitive with complex retrieval and debate-based baselines.

5.2. Explainability and Contestability Analysis

Beyond predictive performance, ACAL addresses the opacity of traditional legal AI. Unlike baselines such as CoT or MAD, which produce unstructured text or debate transcripts, ACAL generates an A-QBAF, a structured graph where decisions are mathematically derived from explicit arguments and intrinsic scores. Furthermore, ACAL advances from passive explainability to active contestability. While RAG systems offer static citations, our HITL contestation workflow empowers users to directly audit and modify the reasoning graph. These interventions are mathematically propagated to update the final judgment, ensuring a transparent and verifiable decision-making process.

5.3. Ablation Study

Clash Resolution (CR) and Uncertainty-Aware Escalation (UAE) Mechanism.

To validate the effectiveness of our proposed modules, we conducted an ablation study on the *Hearsay* dataset using Gemini-2.5-Flash-Lite, effectively isolating the impact of the CR mechanism and the UAE strategy. This ablation study also serves as a comparative evaluation against vanilla ArgLLMs (without CR and UAE) [14], while adopting QE semantics for final argument strengths computation.

As shown in Table 4, the results identify CR as the primary driver of performance, yielding a substantial 7.9% increase in accuracy (64.9% \rightarrow 72.8%) when applied independently, which confirms the necessity of resolving score saturation in LLM-generated arguments. Interestingly, deploying UAE in isolation negatively impacts performance (62.8%), suggesting that uncertainty estimation is unreliable without the calibration provided by CR. However, the full ACAL framework achieves the highest accuracy (74.5%) and F1-score (74.4%), demonstrating a complementary effect in which CR establishes the argument structure required for UAE to operate effectively in borderline cases.

Base Adjustment Magnitude β . We investigated the sensitivity of the hyperparameter β (base adjustment magnitude), which governs the intensity of score updates during Clash Resolution. As detailed in Table 5, performance on the *Learned Hands Courts* dataset exhibits a distinct bell-shaped trend, gradually improving as β increases from 0.05 and peaking at $\beta = 0.15$ with the highest Accuracy (70.8%) and F1-score (70.7%). However, increasing β beyond this threshold results in performance degradation. This finding suggests that while a moderate adjustment is essential to effectively differentiate between conflicting arguments, an excessively aggressive magnitude ($\beta \geq 0.20$) introduces volatility, leading the

final propagated scores to become overly sensitive to individual clash outcomes rather than reflecting the holistic argument structure.

6. Conclusion

In this paper, we propose ACAL, a neuro-symbolic framework designed to address the critical need for performance and contestability in automated legal reasoning. By integrating adaptive multi-agent collaboration with an A-QBAF, ACAL successfully transforms unstructured LLM outputs into formal, verifiable reasoning graphs. Our empirical evaluation on the LegalBench benchmark demonstrates that ACAL achieves competitive performance, surpassing robust baselines including CoT and RAG across both Gemini-2.5-Flash-Lite and Gemini-2.5-Flash models. Crucially, beyond standard performance metrics, ACAL bridges the accountability gap by enabling a HITL contestability workflow. This allows stakeholders to directly audit and intervene in the reasoning process, ensuring that legal judgments are not only accurate but also justifiable and aligned with emerging regulatory standards for high-stakes AI. Future work will focus on optimizing the computational efficiency of the multi-agent architecture to reduce inference costs without compromising reasoning depth. Additionally, we aim to validate the ACAL framework’s generalizability by extending the adaptive agent pool to a broader range of complex legal tasks and other high-stakes domains that require contestable decision-making.

Acknowledgment

This work was supported by NSERC Discovery Grant No RGPIN-2025-04478 and NSERC Discovery Supplement Award No DGECR-2025-00129.

References

- [1] E. Linna and T. Linna. “Challenges for generative AI in legal reasoning”. In: *Discover Artificial Intelligence* (2026). ISSN: 2731-0809.
- [2] N. Guha et al. “LegalBench: A Collaboratively Built Benchmark for Measuring Legal Reasoning in Large Language Models”. In: *Advances in Neural Information Processing Systems*. Vol. 36. 2023, pp. 44123–44279.
- [3] Y.-C. Yu et al. “Structured Evaluation of Legal Reasoning in LLMs: Chain-of-Thought Prompting and Human Scoring for Retrieval Robustness”. In: *NII Institutional Repository* (2025).
- [4] S. Yue et al. “LawLLM: Intelligent Legal System with Legal Reasoning and Verifiable Retrieval”. In: *International Conference on Database Systems for Advanced Applications*. Springer. 2024, pp. 304–321.
- [5] R. Yao et al. “Elevating Legal LLM Responses: Harnessing Trainable Logical Structures and Semantic Knowledge with Legal Reasoning”. In: *arXiv preprint arXiv:2502.07912* (2025).
- [6] G. Chen et al. “AgentCourt: Simulating Court with Adversarial Evolvable Lawyer Agents”. In: *Findings of the Association for Computational Linguistics: ACL 2025*. 2025.
- [7] P. Devadiga, O. J. Shetty, and P. Agarwal. “SAMVAD: A Multi-Agent System for Simulating Judicial Deliberation Dynamics in India”. In: *arXiv preprint arXiv:2509.03793* (2025).
- [8] A. Sadowski and J. A. Chudziak. “On Verifiable Legal Reasoning: A Multi-Agent Framework with Formalized Knowledge Representations”. In: *Proceedings of the 34th ACM International Conference on Information and Knowledge Management*. 2025, pp. 2535–2545.
- [9] S. Jung and J. Jung. “Courtroom-LLM: A Legal-Inspired Multi-LLM Framework for Resolving Ambiguous Text Classifications”. In: *Proceedings of the 31st International Conference on Computational Linguistics*. 2025, pp. 7367–7385.
- [10] J. J. Nay et al. “Large language models as tax attorneys: a case study in legal capabilities emergence”. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 382.2270 (Feb. 2024), p. 20230159. ISSN: 1364-503X.

- [11] C. G. O’Grady and C. O’Grady. “Agentic Workflows in the Practice of Law—AI Agents as Ethics Counsel”. In: *Arizona Legal Studies Discussion Paper* (2024), pp. 25–03.
- [12] European Union. “The EU Artificial Intelligence Act”. In: *The EU Artificial Intelligence Act* (2024).
- [13] Government of Canada. *Directive on Automated Decision-Making- Canada.ca*. 2019.
- [14] G. Freedman, A. Dejl, D. Gorur, X. Yin, A. Rago, and F. Toni. “Argumentative Large Language Models for Explainable and Contestable Claim Verification”. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 39. 14. 2025, pp. 14930–14939.
- [15] J. Cui et al. “Chatlaw: A Multi-Agent Collaborative Legal Assistant with Knowledge Graph Enhanced Mixture-of-Experts Large Language Model”. In: *arXiv:2306.16092* (2023).
- [16] N. Raju et al. “LegalMind: Agentic AI-Driven Process Optimization and Cost Reduction in Legal Services Using DeepSeek”. In: *IEEE Access* (2025).
- [17] S. S. Han et al. “COURTREASONER: Can LLM Agents Reason Like Judges?” In: *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*. 2025.
- [18] Z. He et al. “AgentsCourt: Building Judicial Decision-Making Agents with Court Debate Simulation and Legal Knowledge Augmentation”. In: *Findings of the Association for Computational Linguistics: EMNLP 2024*. 2024, pp. 9399–9416.
- [19] A. Hota and J. P. Jokinen. “NomicLaw: Emergent Trust and Strategic Argumentation in LLMs During Collaborative Law-Making”. In: *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society* 8.2 (2025), pp. 1278–1289.
- [20] L. Zhang, M. Grabmair, M. Gray, and K. Ashley. “Thinking Longer, Not Always Smarter: Evaluating LLM Capabilities in Hierarchical Legal Reasoning”. In: *arXiv preprint* (2025).
- [21] I. Cheong et al. “(A)I Am Not a Lawyer, But...: Engaging Legal Experts towards Responsible LLM Policies for Legal Advice”. In: *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*. FAccT ’24. New York, NY, USA, 2024.
- [22] S. Park et al. “Objection, your honor!: an LLM-driven approach for generating Korean criminal case counterarguments”. In: *Artificial Intelligence and Law* (2025).
- [23] H. Nguyen et al. “Heart2Mind: Human-Centered Contestable Psychiatric Disorder Prediction System Using Wearable ECG Monitors”. In: *ACM Trans. Comput. Healthcare* (2026).
- [24] L. P. T. Nguyen et al. “Motion2Meaning: A Clinician-Centered Framework for Contestable LLM in Parkinson’s Disease Gait Interpretation”. In: *Proceedings of 9th International Symposium on Chatbots and Human-centred AI (CONVERSATIONS) 2025*. 2025.
- [25] M. Gray, L. Zhang, and K. D. Ashley. “Generating Case-Based Legal Arguments with LLMs”. In: *Proceedings of the 2025 Symposium on Computer Science and Law*. CSLAW ’25. Munich, Germany: Association for Computing Machinery, 2025, 160–168. ISBN: 9798400714214.
- [26] S. Liu et al. “LLM Agents in Law: Taxonomy, Applications, and Challenges”. In: *arXiv preprint arXiv:2601.06216* (2026).
- [27] X. Yin, N. Potyka, and F. Toni. “Argument attribution explanations in quantitative bipolar argumentation frameworks”. In: (2023).
- [28] N. Potyka. “Continuous Dynamical Systems for Weighted Bipolar Argumentation.” In: *KR 2018* (2018), pp. 148–57.
- [29] M. M. Hasan et al. “Survey on leveraging uncertainty estimation towards trustworthy deep neural networks: The case of reject option and post-training processing”. In: *ACM Computing Surveys* 57.9 (2025), pp. 1–35.
- [30] T. Brown et al. “Language models are few-shot learners”. In: *Advances in neural information processing systems* 33 (2020), pp. 1877–1901.
- [31] J. Wei et al. “Chain-of-thought prompting elicits reasoning in large language models”. In: *Advances in neural information processing systems* 35 (2022), pp. 24824–24837.
- [32] P. Lewis et al. “Retrieval-augmented generation for knowledge-intensive nlp tasks”. In: *Advances in neural information processing systems* 33 (2020), pp. 9459–9474.
- [33] Y. Du, S. Li, A. Torralba, J. B. Tenenbaum, and I. Mordatch. “Improving factuality and reasoning in language models through multiagent debate”. In: *Proceedings of the 41st International Conference on Machine Learning*. ICML’24. JMLR.org, 2024.

Appendix A. Case Study of ACAL in Legal Reasoning

This appendix presents an illustrative case study of the ACAL framework applied to a complex legal reasoning scenario for a sample in the Hearsay task from LegalBench. Figure 3 demonstrates the end-to-end neuro-symbolic workflow, including: (1-2) Case Aspect Identification and Adaptive Expert Team recruitment; (3) Multi-Agent Argument Generation; (4) HITL contestation where a user explicitly modifies the reasoning graph; (5-6) A-QBAF Graph Construction and Clash Resolution for score calibration; and (7) Final Answer Generation with faithful explanation.

