Advances in Computational Systems, Algorithms, and Emerging Technologies (2022), 7, 19–34

ORIGINALRESEARCH

Rule-Based Expert Systems for Automated Legal Reasoning and Contract Analysis: A Case Study in Knowledge Representation

Alejandro Pérez Gómez

Instituto de Tecnologia do Ceará, Faculdade de Engenharia de Computação, Avenida Dom Luís, Aldeota, Fortaleza, Ceará

Abstract

Rule-based expert systems play a significant role in providing automated insights and decision-making capabilities across complex domains. Within legal practice, such systems have garnered considerable attention because they offer a structured approach to interpreting statutes, regulations, and contract clauses. By integrating domain-specific rules, these technologies guide both practitioners and automated agents through intricate logical paths to arrive at reasoned conclusions or recommendations. At their core, they reduce the cognitive burden of manually parsing lengthy and interconnected legal documents, particularly in the context of contract analysis, compliance checks, and risk mitigation. In this work, we present a comprehensive discussion of rule-based expert systems tailored for legal reasoning with an emphasis on contract analysis. Our exploration delves into fundamental principles of knowledge representation and logical rule structuring to accommodate nuanced legal requirements. We provide a detailed methodology, including formal definitions of relevant symbolic notations and an examination of logic-based inferential mechanisms. By illustrating these components in a case study, we demonstrate how explicit modeling of legal knowledge fosters consistency, transparency, and efficiency. Moreover, we investigate the practical implications of designing and deploying these systems, highlighting the methods to ensure verifiability and maintainability. This research endeavors to underline the critical importance of rule-based expert systems in automating legal reasoning processes within organizations seeking robust and reliable contract evaluation.

1. Introduction

The evolution of information technology has significantly influenced traditional knowledge-intensive professions, with legal practice witnessing the profound impact of computational methodologies designed to improve efficiency, reduce errors, and enhance the consistency of legal reasoning (Sedaghatbaf and Azgomi 2018). Among these methodologies, rule-based expert systems have garnered substantial attention due to their capacity to encapsulate domain-specific knowledge within a formally structured framework, thereby enabling systematic legal guidance. These systems facilitate decision support across various legal domains, including contract drafting, regulatory compliance, and dispute resolution, by encoding inference rules that embed statutory guidelines, judicial precedents, and contextual requirements into coherent logical models. The formalization of such legal reasoning processes enhances the reliability and replicability of legal interpretations, ultimately contributing to the automation of complex legal analyses.

A pivotal challenge in developing rule-based expert systems for legal applications is the representation of legal knowledge in a manner that balances expressiveness with computational feasibility

(Strecker et al. 2019). This requirement necessitates the use of symbolic logic, discrete mathematical structures, and domain-specific ontologies that can capture intricate legal nuances while remaining amenable to algorithmic processing. Contracts, as legally binding agreements, introduce additional layers of complexity, given that the interpretation of contractual clauses, obligations, and remedies often demands sophisticated logical expressions. The formulation of these expressions must account for ambiguities inherent in natural language while preserving the intended legal semantics. Ensuring the accuracy of these representations requires rigorous validation against real-world legal cases, doctrinal interpretations, and statutory provisions. (Singh et al. 2019)

One effective approach for structuring legal knowledge is the use of formal logic systems such as first-order logic (FOL) and description logic (DL). These logical frameworks enable the explicit encoding of legal principles, constraints, and relationships between contractual entities. First-order logic, in particular, provides a robust foundation for modeling legal arguments by allowing the specification of predicates, quantifiers, and inference rules that capture conditional obligations and entitlements. Description logic, on the other hand, facilitates the construction of ontologies that define the hierarchical relationships between legal concepts, thereby supporting advanced reasoning mechanisms such as subsumption and consistency checking (Valja et al. 2020). The integration of these logical paradigms within rule-based expert systems enables automated reasoning about legal scenarios, improving the consistency and traceability of legal decisions (Forbus et al. 2007).

The implementation of computational legal reasoning also necessitates mechanisms for handling uncertainty and vagueness, which are pervasive in legal texts. Probabilistic logic and fuzzy logic offer viable solutions for addressing these challenges. Probabilistic logic extends classical logic by incorporating probability distributions over legal assertions, thereby allowing for the quantification of uncertainty in legal arguments (Soomro et al. 2021). Fuzzy logic, in contrast, provides a means to represent gradations of truth, which is particularly useful when dealing with ambiguous or imprecise legal terms such as "reasonable effort" or "material breach." By integrating these techniques, rule-based expert systems can accommodate the inherent uncertainties present in legal interpretations, thereby enhancing their applicability to real-world legal practice.

An important consideration in the design of rule-based expert systems is the usability and interpretability of the generated recommendations. Legal practitioners must be able to comprehend and justify the reasoning process underlying automated decisions. To achieve this, many expert systems employ explanation facilities that provide human-readable justifications for the conclusions reached (Alani 2021). These explanations often leverage argumentation frameworks, which represent legal reasoning as a structured debate in which competing interpretations are evaluated based on predefined criteria. Such frameworks not only enhance transparency but also align with the adversarial nature of legal practice, where conflicting positions must be rigorously examined before reaching a resolution.

From a practical standpoint, the deployment of rule-based expert systems in legal practice necessitates careful consideration of their integration with existing legal information systems and workflows. Many legal organizations rely on document management systems, case law databases, and regulatory repositories to access relevant legal information (Kehren et al. 2021). The interoperability of expert systems with these resources is crucial for ensuring that automated analyses are based on up-to-date and authoritative legal sources. Standardization efforts, such as the development of legal markup languages (e.g., LegalRuleML), play a vital role in facilitating such interoperability by providing structured representations of legal rules and arguments. By leveraging these standards, rule-based expert systems can seamlessly interact with legal data sources, thereby enhancing their utility and effectiveness in practice.

To illustrate the application of rule-based expert systems in legal practice, consider a scenario involving contract analysis (Galdi et al. 2019). A rule-based system can be designed to automatically assess the enforceability of contractual clauses by evaluating their compliance with statutory requirements and established case law. The system operates by applying a set of predefined inference rules that encode legal principles governing contract formation, performance, and termination. For example, a rule might specify that a contract lacking consideration is unenforceable unless an exception applies. By systematically applying such rules, the system can generate an assessment report highlighting potential legal issues and suggesting remedial actions. (Zubkova 2018)

Table 1. Example of Inference Rules for Contract Analysis

Rule Name	Condition	Inference
Lack of Consideration	If a contract lacks consideration and no exceptions apply	Contract is unenforceable
Unconscionability	If a contract clause is excessively one-sided and results in unfair advantage	Clause may be voided
Statute of Frauds	If a contract falls within the statute of frauds but lacks a written form	Contract is unenforceable
Parol Evidence Rule	If extrinsic evidence contradicts a fully integrated contract	Evidence is inadmissible

Beyond contract analysis, rule-based expert systems have also been applied to regulatory compliance, where businesses must adhere to complex legal requirements across multiple jurisdictions. Compliance systems utilize predefined regulatory rules to assess an organization's adherence to legal standards and identify potential violations. For instance, financial institutions are subject to anti-money laundering (AML) regulations that require diligent customer due diligence and reporting of suspicious transactions. A rule-based compliance system can automate the detection of compliance risks by cross-referencing transactional data with regulatory requirements (Laranjeira et al. 2020). Such systems improve regulatory oversight while reducing the manual effort required for compliance monitoring (Sharma and Forbus 2012).

Table 2. Regulatory Compliance Rules in Financial Transactions

Rule Name	Condition	Inference
Know Your Customer (KYC)	If customer identity verification is incomplete	Flag transaction for review
Suspicious Activity Reporting (SAR)	If transaction pattern matches known fraudulent behaviors	Generate report for regulatory authorities
Transaction Threshold Reporting	If transaction amount exceeds legal reporting threshold	Require mandatory disclosure
Sanctions Screening	If customer is listed on a sanctions watchlist	Block transaction and notify compliance team

Despite their advantages, rule-based expert systems are not without limitations. One key challenge is the maintenance and updating of rule sets, particularly in areas of law that undergo frequent changes. Ensuring that the system remains current requires continuous monitoring of legal developments and systematic updates to the encoded rules (Almustafa 2020). Another challenge is the potential rigidity of rule-based approaches, as they may struggle to accommodate novel legal arguments or unforeseen case-specific nuances. Hybrid approaches that integrate rule-based reasoning with machine learning techniques offer a promising avenue for addressing these limitations. By leveraging natural language processing (NLP) and case-based reasoning, such systems can dynamically adapt to evolving legal landscapes while maintaining the interpretability of rule-based frameworks.

the application of rule-based expert systems in legal practice represents a significant advancement in the automation of legal reasoning (Raj et al. 2020). By formalizing legal knowledge within structured inference frameworks, these systems enhance efficiency, consistency, and accessibility in legal decision-making. However, their effectiveness depends on careful knowledge representation, robust reasoning mechanisms, and seamless integration with legal information systems. As legal technology continues to evolve, the synergy between rule-based logic and data-driven techniques holds great promise for the future of computational legal reasoning.

To achieve this, researchers have explored various theoretical and practical approaches, including forward and backward chaining methods for rule execution, the integration of domain ontologies to frame legal concepts, and rigorous systems for verifying the correctness of rule implementations (Nagasawa et al. 2021). One of the paramount challenges remains the translation of textual legal documents, often filled with ambiguous language, into a structured and unambiguous representation. This step necessitates a bridge between human-driven legal interpretation and machine-readable logic. In essence, a successful rule-based expert system for contract analysis must seamlessly merge comprehensive legal expertise with computational efficiency, ensuring that every relevant detail is accurately captured.

In this paper, we present a comprehensive framework for designing, implementing, and validating rule-based expert systems geared toward legal reasoning, with a focus on contract analysis (Neha et al. 2021) (Sharma 2011). We discuss the theoretical underpinnings of knowledge representation in the context of the legal domain and then detail a case study that exemplifies how carefully curated rules can automate and streamline contract review. Our goal is to demonstrate not only the feasibility of these systems but also their significance in fostering consistency, transparency, and real-time responsiveness to regulatory changes.

2. Knowledge Representation in Rule-Based Expert Systems

The success of a rule-based expert system hinges on its ability to accurately embody domain knowledge and systematically manipulate that knowledge through logical inferences. In the context of legal reasoning, knowledge representation becomes especially critical, given the complexity and formality inherent in legal documents (Böhm, Menges, and Pernul 2018). Traditional symbolic approaches typically rely on well-defined predicates, inference rules, and domain-specific vocabularies to construct a robust framework for reasoning. A predicate such as

$$\mathrm{HasClause}(C, \mathrm{Arbitration}) \rightarrow \mathrm{RequiresProcedure}(C, \mathrm{ArbitrationProtocol})$$

can encode how a contract C that contains an arbitration clause triggers the obligation to follow a specific arbitration protocol. In this manner, each rule explicitly captures an aspect of legal practice, facilitating modularity and clarity.

Symbolic Representation

Symbolic representation serves as a foundational element in many rule-based expert systems, particularly in the legal domain, where precise definitions and logical structures are necessary for effective reasoning (Malavasi et al. 2018). By employing symbolic notation, legal concepts such as "obligation," "liability," and "remedy" can be explicitly defined within a formal system, ensuring systematic and consistent interpretation. A fundamental example of such representation is the formalization of contractual obligations. In a logical framework, an obligation can be denoted as a binary relation $O(a, b)$, indicating that a party a is obliged to perform an action b. This notation allows an expert system to process and evaluate legal obligations efficiently (Zheltov and Kos'yanchuk 2018). Logical inference mechanisms can be employed to analyze contractual performance, as demonstrated in the following formal rule:

$$\forall a \, \forall b \, (O(a, b) \land \neg \mathrm{Performed}(a, b) \rightarrow \mathrm{Breach}(a, b))\,.$$

This rule encapsulates the notion that if a party a is contractually obligated to execute an action b but fails to do so, a breach of contract is inferred. Such formalizations enable automated reasoning, allowing legal expert systems to detect violations, evaluate liability, and suggest remedial actions based on predefined legal principles.

Beyond obligations, other legal constructs can be systematically represented using symbolic logic (Patalas-Maliszewska, Dudek, and Kłos 2019). For example, liability may be formulated using a ternary relation $L(a, b, c)$, where a is liable to b for failing to fulfill obligation c. Remedies, in turn, can be formally expressed as mappings from breach conditions to compensatory actions. These representations facilitate structured legal reasoning by establishing clear relationships between entities and their respective legal responsibilities.

The domain-specific nature of legal reasoning necessitates the development of sophisticated ontologies that categorize contractual roles, legal actions, and normative statements (Alizadehsani et al. 2021). A well-structured ontology distinguishes between different contractual participants, such as "buyer," "seller," "landlord," and "tenant," as well as legal actions such as "payment," "delivery," and "notice." By explicitly defining these concepts and their interrelations, an expert system can minimize ambiguity and ensure that legal terms are consistently interpreted across various documents. Consider the following example ontology for contractual transactions:

Table 3. Ontology of Legal Roles and Actions in Contracts

Legal Concept	**Definition**	**Example Usage**
Buyer	An entity that agrees to purchase goods or services in exchange for consideration	"The buyer must render payment within 30 days of delivery."
Seller	An entity that agrees to transfer ownership of goods or provide services	"The seller guarantees that the goods are free from defects."
Obligation	A duty imposed by contract or law that a party must fulfill	"The landlord is obligated to maintain the premises in a habitable condition."
Remedy	A legal consequence imposed in response to a breach of obligation	"In the event of non-payment, the seller may seek damages."
Delivery	The act of transferring possession of goods or rendering services as per the contract	"Delivery shall occur within five business days of order confirmation."

Ontologies such as the one illustrated above enable automated contract analysis by providing expert systems with a structured understanding of legal relationships. These representations allow expert systems to extract semantic meaning from legal documents, map obligations to responsible parties, and determine whether contractual provisions align with statutory requirements (Islam et al. 2021). By leveraging ontological structures, rule-based systems can also standardize terminology across diverse legal sources, reducing discrepancies in interpretation.

One of the primary advantages of employing formal symbolic representations in legal expert systems is the ability to automate compliance verification. Consider a regulatory framework governing financial transactions, where institutions are required to conduct due diligence on their clients. A symbolic rule for compliance verification might be represented as follows: (Yadav 2021)

$$\forall x \, (\text{Transaction}(x) \land \text{Amount}(x) > \theta \rightarrow \text{DueDiligenceRequired}(x)) \, .$$

This rule states that for any transaction x, if the transaction amount exceeds a predefined threshold θ, then due diligence procedures must be executed. By embedding such rules into an expert system, financial institutions can automate the identification of high-risk transactions and ensure compliance with legal requirements.

Similarly, rule-based expert systems can facilitate the analysis of liability by applying legal doctrines to factual scenarios. Consider the principle of vicarious liability, where an employer may be held responsible for the actions of an employee (Goh et al. 2021). This principle can be formalized in a logical rule as follows:

$$\forall e \forall a \, (\text{Employee}(e) \land \text{Action}(a, e) \land \text{WithinScope}(a) \rightarrow \text{EmployerLiable}(e, a)) \, .$$

This formulation states that if an individual e is an employee and performs an action a within the scope of their employment, then the employer is liable for that action. Such formal representations enable expert systems to evaluate cases involving employment disputes, workplace injuries, and corporate liability.

Furthermore, symbolic representations extend to judicial precedent analysis, where expert systems can assess past court decisions to determine their applicability to current cases (L. Yang et al. 2020). A case-based reasoning system might encode legal precedents as logical predicates, such as:

$$\text{Precedent}(p) \land \text{SimilarFacts}(c, p) \rightarrow \text{Binding}(c, p).$$

This rule asserts that if a precedent p exists and a new case c shares materially similar facts with p, then the precedent is binding on c. By incorporating such inference rules, legal expert systems can assist in legal research, helping practitioners identify relevant precedents and predict case outcomes based on prior rulings.

In practical applications, the effectiveness of symbolic legal reasoning hinges on the ability to integrate logical rules with natural language processing (NLP) techniques (Kottner et al. 2019) (Abhishek and Basu 2005). Since legal documents are typically written in complex, domain-specific language, expert systems must be equipped with NLP capabilities to extract structured legal concepts from unstructured text. Machine learning-based entity recognition, syntactic parsing, and semantic role labeling can enhance the accuracy of legal text interpretation, enabling expert systems to apply symbolic logic more effectively.

Despite their numerous advantages, symbolic rule-based systems face challenges related to legal ambiguity and evolving jurisprudence. Many legal concepts, such as "reasonable effort" or "material breach," lack precise definitions and require contextual interpretation (Sonkusare et al. 2021). Addressing this issue requires hybrid approaches that combine rule-based logic with probabilistic reasoning, enabling expert systems to handle uncertainty while maintaining interpretability. For example, fuzzy logic can be introduced to handle gradations of legal terms:

$$\text{ReasonableEffort}(a) = \mu_{\text{effort}}(a),$$

where $\mu_{\text{effort}}(a)$ is a membership function that assigns a degree of truth to whether an action a constitutes "reasonable effort." Such formulations enable expert systems to navigate legal indeterminacy while preserving formal structure.

Another challenge is ensuring that symbolic representations remain up-to-date with legal developments. Legal rules and interpretations evolve over time due to legislative amendments, judicial rulings, and regulatory changes (Trifa, Hedhili, and Chaari 2018). Continuous monitoring and updating of rule bases are essential to maintain the accuracy and relevance of expert systems. One approach to addressing this challenge is to integrate automated legal text analysis tools that detect and extract new legal provisions, updating logical rule sets accordingly.

symbolic representation is a cornerstone of rule-based legal expert systems, enabling precise encoding of legal concepts and facilitating automated reasoning over contractual obligations, liabilities, and regulatory compliance. By employing well-structured ontologies, formal logic, and integration with natural language processing techniques, these systems enhance legal analysis, improve decision-making consistency, and reduce the complexity of legal research (Portelli 2020). However, challenges related to legal ambiguity, dynamic jurisprudence, and integration with real-world legal data necessitate ongoing refinement of symbolic methodologies to ensure their continued effectiveness in legal practice.

Logical Foundations

While numerous logic formalisms exist, classical first-order logic (FOL) remains a common choice for high-level rule representation in legal expert systems. First-order logic allows the system to express universal statements ($\forall$), existential statements ($\exists$), and logical connectives ($\wedge, \vee, \rightarrow, \neg$), which collectively enable a rich language for encoding complex legal scenarios. However, legal contexts often require modal extensions or deontic logic to capture concepts like "permissible," "obligatory," and "prohibited," since these notions are pervasive in contracts and regulations. (Bock et al. 2019)

Nonetheless, implementing full-fledged modal or deontic logic can be computationally intensive, especially for large-scale contract analyses that span multiple jurisdictions or legal frameworks. To maintain tractability, many practitioners opt for rule-based systems that encode only essential deontic aspects or approximate them through constraints and conditionals. This compromise balances expressiveness and computational efficiency.

Frames and Semantic Networks

Beyond purely logical representations, frames and semantic networks offer an alternative structure for representing legal knowledge (Sayyouh et al. 2019). A frame-based representation consists of slots (attributes) and fillers (values) that define key aspects of an object or concept. For example, a "Contract" frame might include slots for "Parties," "EffectiveDate," "GoverningLaw," and so on. By instantiating these frames for each contract, the system can systematically query relevant slots and apply rules accordingly. A sample rule could check whether the "GoverningLaw" of the contract is set to "StateX," triggering the application of the specialized legal rules for that jurisdiction. (Miloslavskaya and Tolstoy 2018)

Semantic networks, on the other hand, rely on a graph-based representation, where nodes represent legal concepts and edges represent relationships such as "hasObligation" or "isLinkedTo." In a legal reasoning context, this approach is particularly helpful in tracing dependencies among clauses. For instance, an "IndemnificationClause" node might link to a "LiabilityLimitClause" node if the system identifies a textual or logical dependency between those two parts of a contract.

Both frames and semantic networks can coexist with logical approaches, as each paradigm offers a unique perspective on knowledge representation. Frames and semantic networks excel at capturing relationships and hierarchical structures, whereas logical representations are more adept at specifying formal rules and inference procedures. In practice, legal expert systems often integrate both to leverage their complementary advantages (Prakash, Manconi, and Loew 2021) (Basu et al. 2006).

Inference Mechanisms

Once the knowledge is represented, inference mechanisms drive the actual reasoning process. Forward chaining begins from known facts (e.g., "Party A is a seller") and applies rules to infer new facts (e.g., "Party A must deliver goods by Date X"). Conversely, backward chaining starts with a goal (e.g., "Is Party A in breach?") and attempts to identify which facts and rules must hold for the goal to be true. In legal expert systems, the choice between forward and backward chaining often depends on the typical usage scenario: compliance checks (where forward chaining might suffice) or legal question answering (which might favor backward chaining) (30th annual computational neuroscience meeting: cns*2021–meeting abstracts. 2021).

Logic-based inference processes can be complemented by heuristic rules that guide the order or selection of inferences, particularly in systems dealing with large rule sets. For instance, if certain contract clauses are known to be highly consequential for risk assessment, the system might prioritize checking them first. Such heuristics help reduce computational load and improve the user experience by rapidly surfacing critical issues.

3. Legal Reasoning and Contract Analysis

Legal reasoning is characterized by the interplay of structured logic, interpretative flexibility, and domain-specific expertise (Kwon, Lee, and Mun 2020). Contracts, as foundational documents in various transactions, present a prime arena for applying rule-based expert systems. They are replete with stipulations, obligations, conditions, and possible exceptions that can be captured in the form of rules, each referencing a specific contractual context. By mapping textual clauses to a suite of logical rules, an expert system can analyze contract compliance, highlight inconsistencies, and even propose revisions.

Logical Interpretation of Contract Provisions

Contractual clauses frequently employ conditional language such as "if," "when," and "provided that," mirroring the structure of logical implications (Kumar, Fujita, and Singh 2019). A simple clause might say, "If the buyer fails to make payment within 30 days, then interest shall accrue at a rate of 5% per annum." This can be directly modeled as:

$$(\mathrm{FailPayment}(\mathrm{buyer}, 30\mathrm{days})) \to \mathrm{AccrueInterest}(\mathrm{buyer}, 5\%).$$

In more intricate scenarios, multiple antecedents must be satisfied for a particular legal consequence to arise. Clauses regarding indemnification may incorporate a variety of triggers, from breach of representation to third-party claims. These can be expressed via: (Mastoi et al. 2021)

$$(\mathrm{BreachRep}(b) \lor \mathrm{ThirdPartyClaim}(b)) \land \mathrm{LossIncurred}(\mathrm{a}) \to \mathrm{DutyToIndemnify}(b, a).$$

The precise structure of such clauses often determines the complexity of the resulting inference rules, underscoring the importance of rigorous logical design.

Hierarchy and Priority of Clauses

Legal documents, especially multi-jurisdictional contracts, may contain conflicting clauses or layered provisions. A carefully designed rule-based system must account for the hierarchy and priority rules inherent to the contract's legal environment. For instance, a boilerplate clause may be superseded by a custom provision within the same document (Critical care canada forum 2019 abstracts. 2019) (Sharma, Witbrock, and Goolsbey 2016). To encode this priority scheme, we may utilize a partial order:

$$\mathrm{CustomClause}(C) \prec \mathrm{BoilerplateClause}(C),$$

signifying that a custom clause in contract C takes precedence over conflicting boilerplate clauses. During inference, the system checks these ordering constraints to resolve conflicts, ensuring that the most authoritative or relevant clause prevails.

At a more global level, certain jurisdictions or regulatory frameworks may override contractual stipulations (Lee et al. 2018). The system thus requires meta-rules that capture external legal mandates. If a mandatory regulatory provision conflicts with a contractual clause, the system should identify the clause as invalid or subservient to the higher-level rule. This layered approach is essential in industries such as finance or healthcare, where legal oversight is stringent.

Risk Assessment and Compliance Verification

Another salient application of rule-based expert systems in contract analysis is risk assessment (Gudkov 2020). Contracts often contain terms that expose parties to potential liabilities or operational difficulties. By systematically enumerating risk-related clauses—such as warranty limitations, indemnities, or confidentiality obligations—an expert system can compute a "risk index" or generate flags indicating potential vulnerabilities. One might define a function:

$$\text{RiskFactor}(c) = \sum_{i=1}^{n} \alpha_i \cdot \text{ClauseRisk}(c_i),$$

where $\text{ClauseRisk}(c_i)$ quantifies the risk associated with clause c_i, and α_i is a weighting parameter reflecting the clause's importance. A higher total risk factor alerts users to carefully review or renegotiate specific terms. (Ahmad et al. 2018)

Compliance verification often goes hand in hand with risk assessment. Contracts can stipulate compliance with internal policies or external regulations, such as data protection laws. By introducing rules that relate contract clauses to regulatory requirements, an expert system can flag non-compliant language or omissions. For example, a data protection compliance rule might be written as: (Funke et al. 2019)

$$\neg\text{DataProtectionClause}(C) \rightarrow \text{ComplianceFailure}(C, \text{PrivacyLaw}).$$

Thus, if contract C lacks an explicit data protection clause, the system identifies a compliance gap.

Dispute Resolution and Enforcement

When contractual obligations are disputed, rule-based expert systems can assist by providing a systematic framework for analyzing claims, defenses, and remedies. If a contracting party alleges breach, the system can automatically verify whether the factual conditions for breach are satisfied. If the system infers a breach, it can then suggest relevant remedies, such as damages, specific performance, or termination, depending on the clauses in question. (Shemberko and Uvarova 2019)

Additionally, the system can store references to applicable arbitration or litigation procedures, ensuring that resolution steps comply with the agreed-upon mechanisms. A logical rule might outline the conditions under which arbitration can be invoked, referencing both the clauses in the contract and statutory guidelines. This structured approach eliminates ambiguity during disputes and ensures uniform enforcement of contractual terms.

4. Case Study: Implementation and Methodology

To illustrate the practical application of a rule-based expert system for automated legal reasoning and contract analysis, we present a case study using a prototype system implemented with a specialized knowledge base and an inference engine. The aim is to showcase how different representational and logical techniques can be cohesively combined into a working system that is both transparent and extensible.

System Architecture

The prototype system is organized into three main layers: the Knowledge Base (KB), the Inference Engine, and the User Interface. The KB stores all relevant legal rules, clauses, and concepts, which are structured into ontologies and frames. The Inference Engine processes user queries or triggered events and performs reasoning by matching known facts with the rules in the KB (Sullivan et al. 2018). Finally, the User Interface allows legal professionals to input contract details, view analysis results, and refine rules as necessary.

- **Knowledge Base**: This component encapsulates a broad collection of rules derived from various contract templates and standard legal guidelines. Each rule is codified in a format that follows a uniform logical schema. The KB also includes an ontology of legal concepts, ensuring consistent labeling for entities like "Buyer," "Seller," "Arbitration Clause," and "Liability Cap."

- **Inference Engine**: Implemented using a backward-chaining logic, the engine evaluates a given query (e.g., "Is this contract legally compliant?") by iteratively determining which rules lead to the conclusion. If the user requests forward-looking analysis (e.g., "Which obligations must Party A fulfill by next month?"), the system can switch to a forward-chaining mode. (Malarvizhi, Selvarani, and Raj 2019)
- **User Interface**: Through a graphical interface, attorneys or contract managers can upload contract text. An internal parser attempts to identify key clauses and map them to the system's legal concepts. Users can also manually label clauses to improve accuracy. Once analysis is complete, the interface presents a summary of potential issues, recommended revisions, and relevant references to statutory or regulatory provisions.

Representation Strategy

The system employs a hybrid representation strategy that combines logical rules with frame-based structures (Wang et al. 2021). Each contract is instantiated as a frame, with slots corresponding to metadata such as the parties involved, jurisdiction, effective date, and essential clauses. Embedded within each frame are pointers to logic rules that govern the interpretation of that clause. For instance, a "LimitationOfLiability" clause might contain a pointer to a rule of the form:

$$\mathrm{HasClause}(C, \mathrm{LimitationOfLiability}) \land \mathrm{HighRiskService}(C) \rightarrow \mathrm{CheckExclusions}(C).$$

When the Inference Engine identifies that a contract C has both a limitation of liability clause and involves a high-risk service (e.g., chemical manufacturing), it prompts the user to ensure that specific exclusions or indemnities are explicitly stated.

Logic Encoding and Parsing

To handle the textual extraction of clauses, we employ a natural language processing (NLP) subsystem that identifies legally significant patterns such as "the party shall indemnify," "in the event of," or "notice must be given." These patterns are assigned symbolic labels and mapped to the system's ontology (Singh-Bains et al. 2021). For example, a phrase recognized as "BreachOfContractCondition" might trigger a transformation rule that inserts $\mathrm{BreachClause}(C)$ into the KB for the relevant contract C.

During the parsing stage, potential ambiguities are flagged for user confirmation. Legal texts often include language that can fit multiple categories or none conclusively. For instance, "the party must provide timely notice in the event of any dispute arising from the manufacturing process" can simultaneously implicate notice requirements, dispute resolution, and manufacturing obligations. The system thus queries the user, "Does this clause represent a notice requirement, a dispute resolution provision, or both?" Such guided user interaction reduces errors in the representation. (Bozsahin 2018)

Inference Procedures and Results

After populating the system with contract-specific facts, the user initiates one of several inference procedures:

1. **Compliance Check**: The system systematically evaluates each clause against the relevant regulatory requirements, referencing rules that incorporate statutory obligations. Missing or incomplete clauses are flagged with suggestions for remedial language.
2. **Risk Assessment**: By summing assigned risk values, the system produces an aggregate risk score. This score can be broken down by contractual categories (e.g., liability, indemnity, intellectual property), helping the user focus on high-risk elements.

3. **Obligation Timetable**: The system compiles a timeline of obligations and milestones. Forward chaining ensures that any unfulfilled conditions automatically trigger subsequent requirements (Khatri et al. 2020). For example, failing to deliver goods within the specified timeline might activate clauses related to liquidated damages.
4. **Breach Analysis**: By querying "Is there a breach of contract?" the system performs backward chaining to identify the chain of conditions leading to a breach. If the conditions are satisfied, it enumerates the remedies and potential liabilities per relevant clauses.

A typical output might read: "A potential breach has been detected for Party B related to late delivery obligations. Clause 10.2 states that Party B must deliver goods by Date X, but current facts indicate delivery on Date Y (Saheb and Saheb 2020). Under Clause 11.1, this breach activates penalty fees and Clause 12.3 indicates potential termination rights for Party A."

Refinement and Maintenance

As new legal precedents emerge or contract standards evolve, the knowledge base must be updated. The system includes an administrator dashboard for rule modification, where domain experts can insert or revise rules using a structured editor. Additionally, logs of inference processes are stored to aid debugging and system refinement (Büyükkaramikli et al. 2019). If a rule repeatedly produces false positives, domain experts can adjust its logical structure, add exceptions, or refine clause detection patterns.

5. Discussion

Building and deploying a rule-based expert system for automated legal reasoning involves complex interactions of technology, law, and human expertise. Practical deployment yields insights that illuminate the system's strengths and potential pitfalls. A key advantage lies in the system's transparency; because rules are explicitly encoded, every inference can be traced to a discrete logical statement (Vijh, Gaur, and Kumar 2019). This traceability is invaluable in a legal context where justification and defensibility of reasoning are crucial.

Another critical consideration is system scope. While broad legal domains can theoretically be encoded, over-generalization often leads to performance bottlenecks or incomplete coverage of specialized niches. Many organizations find it more productive to build narrower, domain-specific expert systems that address recurring contractual themes—such as non-disclosure agreements or master service agreements—rather than trying to encode the entirety of contract law in a single, monolithic knowledge base. (He and Huisken 2020)

Despite these advantages, rule-based systems face limitations. Ambiguity in legal language can hinder clause detection and classification. Even with NLP assistance, resolving certain ambiguities requires domain expertise or negotiation context that the system cannot fully capture. Likewise, the dynamic nature of legal interpretation—where judicial decisions and novel contractual clauses continuously reshape the domain—demands ongoing maintenance of the knowledge base (Moetesum et al. 2020). Implementing robust change management processes is thus essential to ensure the system remains current.

Issues of explainability and reliability also intersect with ethical concerns. An automated system that flags or fails to flag a crucial contractual risk could significantly impact negotiations or dispute outcomes. Ensuring that human legal professionals remain in the loop, with final authority to interpret or override system outputs, is generally considered best practice (Migliorini et al. 2020). From a policy perspective, guidelines are emerging to address AI explainability, fairness, and accountability. Given the heightened sensitivity of legal matters, stakeholders may expect rigorous audits of the system's knowledge representation and inference mechanisms.

One of the emerging frontiers is the integration of machine learning techniques with rule-based systems. Hybrid approaches might allow systems to learn from historical data how certain contractual clauses lead to disputes or how negotiations typically converge, refining the weighting parameters used in risk assessments (Lourenço et al. 2020). Nevertheless, the interpretability of purely data-driven models often pales in comparison to the clarity of rule-based logic, reinforcing the value of combining symbolic and statistical methods.

From an architectural standpoint, the modular design of rule-based expert systems offers a pathway to incremental upgrades. Organizations can begin with a modest set of rules addressing a high-value area (e.g., liability or confidentiality) and subsequently expand the rule base. Each module can remain relatively self-contained, making it easier to track changes, add new knowledge, and diagnose logical conflicts. (Rana, Rathi, and Ganguly 2020)

The effectiveness of a legal expert system also depends on the user's trust in the technology. Through consistent performance, transparent reasoning steps, and a user-centric interface, system credibility can be built over time. Legal professionals who see value in automated checks for standard provisions or systematic compliance verifications can incorporate the tool into their daily workflows, ultimately leading to more accurate and efficient contract management.

6. Conclusion

Rule-based expert systems offer a structured and transparent approach to automating legal reasoning, particularly in the realm of contract analysis (Li and Sugumaran 2018). By representing clauses, obligations, and exceptions through logical rules, these systems facilitate rapid identification of compliance issues, risk exposures, and potential breaches. They also provide an auditable trail of logic, ensuring that legal conclusions can be justified to both internal and external stakeholders.

The effectiveness of such systems depends on robust knowledge representation strategies that accurately model the intricacies of the legal domain. Logical formalisms must be chosen with care, balancing the expressive requirements of legal nuances against computational feasibility (P. Yang et al. 2021). Hybrid representations, integrating frame-based structures, semantic networks, and symbolic logic, can capture complex relationships among clauses and statutes. Equally important is the inference mechanism—be it forward or backward chaining—tailored to the system's principal use cases.

Nevertheless, the success of rule-based expert systems is not solely a technical matter. Ongoing maintenance of the knowledge base is critical, as legal standards and precedents evolve (Lachal et al. 2019). Human oversight remains essential for interpreting ambiguous language and ensuring that the system's recommendations align with strategic business considerations. Moving forward, the convergence of rule-based paradigms with data-driven methods holds promise for even more adaptive and robust legal expert systems. This synergy may broaden the scope of automated reasoning, enabling systems to learn from historical legal outcomes while preserving the clarity and rigor of symbolic logic.

In sum, rule-based expert systems demonstrate immense potential to transform how contracts are analyzed, reducing both the time and uncertainty inherent in conventional legal processes. Through careful design, meticulous representation of legal structures, and thoughtful integration into professional workflows, they stand poised to become indispensable tools in the modern legal landscape. (Vitek et al. 2020)

References

Abhishek, A, and Anupam Basu. 2005. A framework for disambiguation in ambiguous iconic environments. In *Ai 2004: advances in artificial intelligence: 17th australian joint conference on artificial intelligence, cairns, australia, december 4-6, 2004. proceedings 17,* 1135–1140. Springer.

Ahmad, Aakash, Claus Pahl, Ahmed B. Altamimi, and Abdulrahman Alreshidi. 2018. Mining patterns from change logs to support reuse-driven evolution of software architectures. *Journal of Computer Science and Technology* 33, no. 6 (November 19, 2018): 1278–1306. https://doi.org/10.1007/s11390-018-1887-3.

Alani, Mohammed M. 2021. Big data in cybersecurity: a survey of applications and future trends. *Journal of Reliable Intelligent Environments* 7, no. 2 (January 6, 2021): 85–114. https://doi.org/10.1007/s40860-020-00120-3.

Alizadehsani, Roohallah, Mohamad Roshanzamir, Sadiq Hussain, Abbas Khosravi, Afsaneh Koohestani, Mohammad Hossein Zangooei, Moloud Abdar, et al. 2021. Handling of uncertainty in medical data using machine learning and probability theory techniques: a review of 30 years (1991–2020). *Annals of operations research* 339, no. 3 (March 21, 2021): 1–42. https://doi.org/10.1007/s10479-021-04006-2.

Almustafa, Khaled Mohamad. 2020. Prediction of heart disease and classifiers' sensitivity analysis. *BMC bioinformatics* 21, no. 1 (July 2, 2020): 278–. https://doi.org/10.1186/s12859-020-03626-y.

Basu, Anupam, et al. 2006. Iconic interfaces for assistive communication. In *Encyclopedia of human computer interaction,* 295–302. IGI Global.

Bock, Estuardo Alpirez, Joppe W. Bos, Chris Brzuska, Charles Hubain, Wil Michiels, Cristofaro Mune, Eloi Sanfelix Gonzalez, Philippe Teuwen, and Alexander Treff. 2019. White-box cryptography: don't forget about grey-box attacks. *Journal of Cryptology* 32, no. 4 (February 12, 2019): 1095–1143. https://doi.org/10.1007/s00145-019-09315-1.

Böhm, Fabian, Florian Menges, and Günther Pernul. 2018. Graph-based visual analytics for cyber threat intelligence. *Cybersecurity* 1, no. 1 (December 28, 2018): 1–19. https://doi.org/10.1186/s42400-018-0017-4.

Bozsahin, Cem. 2018. Computers aren't syntax all the way down or content all the way up. *Minds and Machines* 28, no. 3 (July 18, 2018): 543–567. https://doi.org/10.1007/s11023-018-9469-2.

Büyükkaramikli, Nasuh, Maureen P.M.H. Rutten-van Mölken, Johan L. Severens, and Maiwenn Al. 2019. Tech-ver: a verification checklist to reduce errors in models and improve their credibility. *PharmacoEconomics* 37, no. 11 (November 8, 2019): 1391–1408. https://doi.org/10.1007/s40273-019-00844-y.

Critical care canada forum 2019 abstracts. 2019. *Canadian journal of anaesthesia = Journal canadien d'anesthesie* 67, no. Suppl 1 (December 10, 2019): 1–162. https://doi.org/10.1007/s12630-019-01552-z.

Forbus, Kenneth D, Christopher Riesbeck, Lawrence Birnbaum, Kevin Livingston, Abhishek B Sharma, and Leo C Ureel. 2007. A prototype system that learns by reading simplified texts. In *Aaai spring symposium: machine reading,* 49–54.

Funke, Isabel, Sören Torge Mees, Jürgen Weitz, and Stefanie Speidel. 2019. Video-based surgical skill assessment using 3d convolutional neural networks. *International journal of computer assisted radiology and surgery* 14, no. 7 (May 18, 2019): 1217–1225. https://doi.org/10.1007/s11548-019-01995-1.

Galdi, Paola, Michele Fratello, Francesca Trojsi, Antonio Dello Russo, Gioacchino Tedeschi, Roberto Tagliaferri, and Fabrizio Esposito. 2019. Stochastic rank aggregation for the identification of functional neuromarkers. *Neuroinformatics* 17, no. 4 (January 2, 2019): 479–496. https://doi.org/10.1007/s12021-018-9412-y.

Goh, Kenneth C. W., Raymond B. C. Ng, Yoke-Keong Wong, Nicholas J. H. Ho, and Matthew Chin Heng Chua. 2021. Aerial filming with synchronized drones using reinforcement learning. *Multimedia Tools and Applications* 80, no. 12 (February 15, 2021): 18125–18150. https://doi.org/10.1007/s11042-020-10388-5.

Gudkov, Aleksei. 2020. On fiduciary relationship with artificial intelligence systems. *Liverpool Law Review* 41, no. 3 (June 13, 2020): 251–273. https://doi.org/10.1007/s10991-020-09248-x.

He, Jiaye, and Jan Huisken. 2020. Image quality guided smart rotation improves coverage in microscopy. *Nature communications* 11, no. 1 (January 9, 2020): 150–150. https://doi.org/10.1038/s41467-019-13821-y.

Islam, Saiful, Aurpan Dash, Ashek Seum, Amir Hossain Raj, Tonmoy Hossain, and Faisal Muhammad Shah. 2021. Exploring video captioning techniques: a comprehensive survey on deep learning methods. *SN Computer Science* 2, no. 2 (February 27, 2021): 1–28. https://doi.org/10.1007/s42979-021-00487-x.

30th annual computational neuroscience meeting: cns*2021-meeting abstracts. 2021. *Journal of computational neuroscience* 49, no. Suppl 1 (December 21, 2021): 3–208. https://doi.org/10.1007/s10827-021-00801-9.

Kehren, Christian, Pascal Ortlieb, Daniel Henaux, and Tolga Uhlmann. 2021. Friction estimation for dedicated hybrid engines. *MTZ worldwide* 83, no. 1 (December 10, 2021): 16–23. https://doi.org/10.1007/s38313-021-0740-1.

Khatri, Sahil, Hrishikesh Vachhani, Shalin Shah, Jitendra Bhatia, Manish Chaturvedi, Sudeep Tanwar, and Neeraj Kumar. 2020. Machine learning models and techniques for vanet based traffic management: implementation issues and challenges. *Peer-to-Peer Networking and Applications* 14, no. 3 (September 17, 2020): 1778–1805. https://doi.org/10.1007/s12083-020-00993-4.

Kottner, Sören, Sarah Schaerli, Martin Fürst, Wolfgang Ptacek, Michael J. Thali, and Dominic Gascho. 2019. Virtoscan-on-rails – an automated 3d imaging system for fast post-mortem whole-body surface documentation at autopsy tables. *Forensic science, medicine, and pathology* 15, no. 2 (March 9, 2019): 198–212. https://doi.org/10.1007/s12024-019-00095-5.

Kumar, Binod, Masahiro Fujita, and Virendra Singh. 2019. Sat-based silicon debug of electrical errors under restricted observability enhancement. *Journal of Electronic Testing* 35, no. 5 (November 19, 2019): 655–678. https://doi.org/10.1007/s10836-019-05830-y.

Kwon, Soonjo, Hyunoh Lee, and Duhwan Mun. 2020. Semantics-aware adaptive simplification for lightweighting diverse 3d cad models in industrial plants. *Journal of Mechanical Science and Technology* 34, no. 3 (March 3, 2020): 1289–1300. https://doi.org/10.1007/s12206-020-0228-y.

Lachal, Romain, Margrit Ascher, Francisco Gallego Durán, Fanny Pradalier, C. Jourdan, Kévin Chalard, Flora Djanikian, et al. 2019. Proceedings of réanimation 2019, the french intensive care society international congress. *Annals of intensive care* 9, no. Suppl 1 (March 29, 2019): 1–153. https://doi.org/10.1186/s13613-018-0474-7.

Laranjeira, Matheus, Aurélien Arnaubec, Lorenzo Brignone, Claire Dune, and Jan Opderbecke. 2020. 3d perception and augmented reality developments in underwater robotics for ocean sciences. *Current Robotics Reports* 1, no. 3 (July 10, 2020): 123–130. https://doi.org/10.1007/s43154-020-00014-5.

Lee, Gil-Yong, Mincheol Kim, Ying-Jun Quan, Min-Sik Kim, Thomas J. Y. Kim, Hae-Sung Yoon, Sangkee Min, et al. 2018. Machine health management in smart factory: a review. *Journal of Mechanical Science and Technology* 32, no. 3 (April 13, 2018): 987–1009. https://doi.org/10.1007/s12206-018-0201-1.

Li, Li, and Vijayan Sugumaran. 2018. A cognitive-based aes model towards learning written english. *Journal of Ambient Intelligence and Humanized Computing* 10, no. 5 (March 1, 2018): 1811–1820. https://doi.org/10.1007/s12652-018-0743-1.

Lourenço, Maicon Pierre, Alexandre dos Santos Anastácio, Andreia Luisa da Rosa, Thomas Frauenheim, and Mauricio Chagas da Silva. 2020. An adaptive design approach for defects distribution modeling in materials from first-principle calculations. *Journal of molecular modeling* 26, no. 7 (July 1, 2020): 1–12. https://doi.org/10.1007/s00894-020-04438-w.

Malarvizhi, N., P. Selvarani, and Pethuru Raj. 2019. Adaptive fuzzy genetic algorithm for multi biometric authentication. *Multimedia Tools and Applications* 79, no. 13 (March 15, 2019): 9131–9144. https://doi.org/10.1007/s11042-019-7436-4.

Malavasi, S., Domenico Barone, G. Gavelli, and A. Bevilacqua. 2018. Can perfusion heterogeneity in ct perfusion maps of nsclc flaw clinical considerations based on global mean blood flow values. *Insights into imaging* 9, no. Suppl 1 (February 15, 2018): 312–312. https://doi.org/10.1007/s13244-018-0603-8.

Mastoi, Qurat-ul-ain, Muhammad Suleman Memon, Abdullah Lakhan, Mazin Abed Mohammed, Mumtaz Qabulio, Fadi Al-Turjman, and Karrar Hameed Abdulkareem. 2021. Machine learning-data mining integrated approach for premature ventricular contraction prediction. *Neural Computing and Applications* 33, no. 18 (March 14, 2021): 11703–11719. https://doi.org/10.1007/s00521-021-05820-2.

Migliorini, Matteo, Riccardo Castellotti, Luca Canali, and Marco Zanetti. 2020. Machine learning pipelines with modern big data tools for high energy physics. *Computing and Software for Big Science* 4, no. 1 (June 15, 2020): 1–12. https://doi.org/10.1007/s41781-020-00040-0.

Miloslavskaya, Natalia, and Alexander Tolstoy. 2018. Internet of things: information security challenges and solutions. *Cluster Computing* 22, no. 1 (July 23, 2018): 103–119. https://doi.org/10.1007/s10586-018-2823-6.

Moetesum, Momina, Imran Siddiqi, Shoaib Ehsan, and Nicole Vincent. 2020. Deformation modeling and classification using deep convolutional neural networks for computerized analysis of neuropsychological drawings. *Neural Computing and Applications* 32, no. 16 (January 20, 2020): 12909–12933. https://doi.org/10.1007/s00521-020-04735-8.

Nagasawa, Kazuki, Junki Yoshii, Shoji Yamamoto, Wataru Arai, Satoshi Kaneko, Keita Hirai, and Norimichi Tsumura. 2021. Prediction of the layered ink layout for 3d printers considering a desired skin color and line spread function. *Optical Review* 28, no. 4 (July 8, 2021): 449–461. https://doi.org/10.1007/s10043-021-00679-z.

Neha, null, Harish Kumar Sardana, R Kanwade, and Suman Tewary. 2021. Arrhythmia detection and classification using ecg and ppg techniques: a review. *Physical and engineering sciences in medicine* 44, no. 4 (November 2, 2021): 1–22. https://doi.org/10.1007/s13246-021-01072-5.

Patalas-Maliszewska, Justyna, Adam Dudek, and Sławomir Kłos. 2019. The automated acquisition of expert knowledge using a service department as an example, 119–126. United States: Springer International Publishing, April 26, 2019. https://doi.org/10.1007/978-3-030-18715-6_10.

Portelli, Raechel A. 2020. Don't throw the baby out with the bathwater: reappreciating the dynamic relationship between humans, machines, and landscape images. *Landscape Ecology* 35, no. 4 (March 19, 2020): 815–822. https://doi.org/10.1007/s10980-020-00992-z.

Prakash, Nikhil, Andrea Manconi, and Simon Loew. 2021. A new strategy to map landslides with a generalized convolutional neural network. *Scientific reports* 11, no. 1 (May 6, 2021): 9722–9722. https://doi.org/10.1038/s41598-021-89015-8.

Raj, Shyla, D. S. Vinod, B. S. Mahanand, and Nagaraj Murthy. 2020. Intuitionistic fuzzy c means clustering for lung segmentation in diffuse lung diseases. *Sensing and Imaging* 21, no. 1 (July 16, 2020): 1–16. https://doi.org/10.1007/s11220-020-00300-8.

Rana, Rashmi, Vaishnavi Rathi, and Nirmal Kumar Ganguly. 2020. A comprehensive overview of proteomics approach for covid 19: new perspectives in target therapy strategies. *Journal of proteins and proteomics* 11, no. 4 (November 2, 2020): 1–10. https://doi.org/10.1007/s42485-020-00052-9.

Saheb, Tahereh, and Tayebeh Saheb. 2020. Understanding the development trends of big data technologies: an analysis of patents and the cited scholarly works. *Journal of Big Data* 7, no. 1 (February 22, 2020): 1–26. https://doi.org/10.1186/s40537-020-00287-9.

Sayyouh, M., L. M. Hadjiiyski, H Chan, and P. Agarwal. 2019. Lung nodule: imaging features and evaluation in the age of machine learning. *Current Pulmonology Reports* 8, no. 3 (July 22, 2019): 86–95. https://doi.org/10.1007/s13665-019-00229-8.

Sedaghatbaf, Ali, and Mohammad Abdollahi Azgomi. 2018. Sqme: a framework for modeling and evaluation of software architecture quality attributes. *Software & Systems Modeling* 18, no. 4 (May 31, 2018): 2609–2632. https://doi.org/10.1007/s10270-018-0684-3.

Sharma, Abhishek. 2011. Structural and network-based methods for knowledge-based systems. PhD diss., Northwestern University.

Sharma, Abhishek, and Kenneth Forbus. 2012. Modeling the evolution of knowledge in learning systems. In *Proceedings of the aaai conference on artificial intelligence,* 26:669–675. 1.

Sharma, Abhishek, Michael Witbrock, and Keith Goolsbey. 2016. Controlling search in very large commonsense knowledge bases: a machine learning approach. *arXiv preprint arXiv:1603.04402.*

Shemberko, L. V., and T. B. Uvarova. 2019. Information and analytical support of archaeological research on the basis of digital technologies. *Scientific and Technical Information Processing* 46, no. 2 (August 14, 2019): 133–142. https://doi.org/10.3103/s0147688219020151.

Singh, Jasvinder Pal, Sanjeev Jain, Sakshi Arora, and Uday Pratap Singh. 2019. A survey of behavioral biometric gait recognition: current success and future perspectives. *Archives of Computational Methods in Engineering* 28, no. 1 (November 20, 2019): 107–148. https://doi.org/10.1007/s11831-019-09375-3.

Singh-Bains, Malvindar K., Nasim F. Mehrabi, Adelie Y. S. Tan, Richard L.M. Faull, and Mike Dragunow. 2021. Preparation, construction and high-throughput automated analysis of human brain tissue microarrays for neurodegenerative disease drug development. *Nature protocols* 16, no. 4 (March 19, 2021): 2308–2343. https://doi.org/10.1038/s41596-021-00503-7.

Sonkusare, Saurabh, Michael Breakspear, Tianji Pang, Vinh T. Nguyen, Sascha Frydman, Christine C. Guo, and Matthew J. Aburn. 2021. Data-driven analysis of facial thermal responses and multimodal physiological consistency among subjects. *Scientific reports* 11, no. 1 (June 8, 2021): 12059–12059. https://doi.org/10.1038/s41598-021-91578-5.

Soomro, Toufique Ahmed, Lihong Zheng, Ahmed J. Afifi, Ahmed Hamza H. Ali, Ming Yin, and Junbin Gao. 2021. Artificial intelligence (ai) for medical imaging to combat coronavirus disease (covid-19): a detailed review with direction for future research. *Artificial intelligence review* 55, no. 2 (April 15, 2021): 1–31. https://doi.org/10.1007/s10462-021-09985-z.

Strecker, Stefan, Ulrike Baumöl, Dimitris Karagiannis, Agnes Koschmider, Monique Snoeck, and Rüdiger Zarnekow. 2019. Five inspiring course (re-)designs: examples of innovations in teaching and learning bise. *Business & Information Systems Engineering* 61, no. 2 (January 31, 2019): 241–252. https://doi.org/10.1007/s12599-019-00584-5.

Sullivan, Devin P., Casper F. Winsnes, Lovisa Åkesson, Martin Hjelmare, Mikaela Wiking, Rutger Schutten, Linzi Campbell, et al. 2018. Deep learning is combined with massive-scale citizen science to improve large-scale image classification. *Nature biotechnology* 36, no. 9 (August 20, 2018): 820–828. https://doi.org/10.1038/nbt.4225.

Trifa, Amal, Aroua Hedhili, and Wided Lejouad Chaari. 2018. Knowledge tracing with an intelligent agent, in an e-learning platform. *Education and Information Technologies* 24, no. 1 (August 25, 2018): 711–741. https://doi.org/10.1007/s10639-018-9792-5.

Valja, Margus, Fredrik Heiding, Ulrik Franke, and Robert Lagerström. 2020. Automating threat modeling using an ontology framework. *Cybersecurity* 3, no. 1 (October 1, 2020): 1–20. https://doi.org/10.1186/s42400-020-00060-8.

Vijh, Surbhi, Deepak Gaur, and Sushil Kumar. 2019. An intelligent lung tumor diagnosis system using whale optimization algorithm and support vector machine. *International Journal of System Assurance Engineering and Management* 11, no. 2 (September 6, 2019): 374–384. https://doi.org/10.1007/s13198-019-00866-x.

Vitek, Matej, Peter Rot, Vitomir Struc, and Peter Peer. 2020. A comprehensive investigation into sclera biometrics: a novel dataset and performance study. *Neural Computing and Applications* 32, no. 24 (February 19, 2020): 17941–17955. https://doi.org/10.1007/s00521-020-04782-1.

Wang, Yang, Fuqian Yang, Jun Zhang, Huidong Wang, Xianwen Yue, and Shanshan Liu. 2021. Application of artificial intelligence based on deep learning in breast cancer screening and imaging diagnosis. *Neural Computing and Applications* 33, no. 15 (April 2, 2021): 9637–9647. https://doi.org/10.1007/s00521-021-05728-x.

Yadav, Satya Prakash. 2021. Emotion recognition model based on facial expressions. *Multimedia Tools and Applications* 80, no. 17 (May 1, 2021): 26357–26379. https://doi.org/10.1007/s11042-021-10962-5.

Yang, Ling, Yeqi Liu, Huihui Yu, Xiaomin Fang, Lihua Song, Daoliang Li, and Chen Yingyi. 2020. Computer vision models in intelligent aquaculture with emphasis on fish detection and behavior analysis: a review. *Archives of Computational Methods in Engineering* 28, no. 4 (September 5, 2020): 2785–2816. https://doi.org/10.1007/s11831-020-09486-2.

Yang, Po, Gaoshan Bi, Jun Qi, Xulong Wang, Yun Yang, and Li Da Xu. 2021. Multimodal wearable intelligence for dementia care in healthcare 4.0: a survey. *Information Systems Frontiers* (July 21, 2021): 1–18. https://doi.org/10.1007/s10796-021-10163-3.

Zheltov, S. Yu., and V. V. Kos'yanchuk. 2018. Prospects for the intellectualization of state-of-the-art aviation systems. *Herald of the Russian Academy of Sciences* 88, no. 1 (March 9, 2018): 28–36. https://doi.org/10.1134/s1019331618010136.

Zubkova, T. M. 2018. Automated industrial design based on artificial intelligence. *Russian Engineering Research* 38, no. 5 (June 30, 2018): 394–398. https://doi.org/10.3103/s1068798x18050222.