SYM-AM-26-137

Excerpt from the Proceedings

of the

Twenty-Third Annual Acquisition Research Symposium and Innovation Summit

Volume III

"Accelerating Warfighting Capabilities"

How to Improve Contract Clause Logic Software Automation

Published: April 30, 2026

Approved for public release; distribution is unlimited.

Prepared for the Naval Postgraduate School, Monterey, CA 93943.

Naval Postgraduate School Foundation

The research presented in this report was supported by the Acquisition Research Program, Graduate School of Defense Management at the Naval Postgraduate School.

To request defense acquisition research, please contact:

Acquisition Research Program
Department of Defense Management
Naval Postgraduate School
E: arp@nps.edu
www.acquisitionresearch.net

How to Improve Contract Clause Logic Software Automation

David Gill—is a procurement technology specialist in the Research office at the Department of the Treasury, Internal Revenue Service. David leads automation and artificial intelligence projects designed to reduce the friction in the procurement process. His previous papers discussed topics such as procurement lead times, contract spending data quality, and vendor risk predictive modeling. David holds degrees in government and business from University of Maryland schools and completed the AI Policy Precepts program at Princeton University. Former positions include Contracting Officer and Supervisory IT Specialist roles. David holds a Federal Acquisition Certification in Contracting, Level III.

Abstract

This study explores emerging software approaches for ensuring the correct inclusion of contract clauses in federal acquisitions—an area where accuracy directly affects legal compliance, policy implementation, and contract performance. Clause selection remains a complex, high-risk task driven by numerous intersecting variables (e.g., contract type, dollar thresholds, place of performance, and statutory triggers), and is still largely manual and error-prone across the acquisition workforce. Using empirical data from the IRS Contract Clause Review Tool and a review of government and commercial clause automation solutions, this paper compares traditional rule-based systems, questionnaire-driven tools, and emerging AI-enabled approaches. Findings show that template-only and manual methods are inefficient and prone to error, while unconstrained generative AI introduces risks in version control and legal fidelity. A hybrid approach—combining automated document analysis, rules-based logic, and targeted AI support—achieves the best balance of precision, scalability, and usability. For acquisition leaders, the key implication is clear: clause compliance can be significantly improved by shifting from user-driven selection to system-assisted validation, reducing errors, saving time, and strengthening the defensibility of contract outcomes at scale.

Problem Statement

Identifying the applicable clauses and provisions is highly complex, largely manual, and an error prone task that must be performed by the acquisition workforce. Law professor Alex Ritchie described the typical inadequate approach to contract boilerplate explaining that "[w]hether because of time pressures, habit, or carelessness, review and drafting of contract boilerplate often is relegated to untrained lawyers or contract personnel, performed at the wee hours of the night, or simply skipped altogether." Clauses are prescribed in the Federal Acquisition Regulations (FAR) and agency supplements. Clause prescriptions depend on intersecting variables: place of performance (United States/overseas), acquisition method (simplified/negotiated), contract type, commerciality, inflation adjusted dollar thresholds, construction vs. services, data/intellectual property deliverables, privacy, security/supply-chain prohibitions, and agency-specific deviations. Alternate clause versions further multiply paths; a single omitted fact (e.g., awardee is a university) can flip an alternate.

Contracting Officers have some of the longest job descriptions in government with responsibilities extending far beyond the foundational contract management triad of price, quality, and delivery. Contracting Officers must also navigate a myriad of complex areas such as intellectual property, Buy American Act / international trade agreements, energy efficiency, policing conflicts of interest, Service Contract Labor Standards, and other responsibilities described in multiple layers of laws, regulations, and policies applicable to contracts.

Failure to include a clause can impair Government interests and impact contract performance and legal compliance. For example, a bid protest case discussed an agency's omission of FAR Clause 52.219-14, Limitations on Subcontracting. Failure to include this clause in a set-aside contract hindered the agency's ability to ensure that contract work would be primarily performed by small businesses—a legal requirement of small business contracting programs. Another, example illustrating the importance of clause inclusion is a Board of Contract Appeals decision explaining the Government has the burden of proving an option to extend a contract is properly exercised. If an option clause is omitted and other issues are present (e.g., lack of required advance notice to contractor, failure to evaluate option pricing at time of award) then the Government cannot properly exercise an option to extend a contract.

Research Results

Better clause software can help reduce buying friction and speed up one aspect of the acquisition process. Template-only tools oversimplify complex prescriptions; unconstrainted artificial intelligence (AI) large language model (LLM) first drafting risks departure from verbatim official text and latency delays. Hybrid approaches combining AI and hard coded business rules with automation achieve high accuracy, preserve verbatim text, and save time for users. We conclude that surgical use of AI, maximizing automation, and maintaining immutable clause offers the best balance of accuracy, automation, and usability.

Some Traditional contract writing information systems require users to manually select each clause that is applicable to a solicitation or contract. Alternatively, user questionnaires can be used by a system to generate applicable clauses. Other clause logic software provided clause templates based on the FAR Matrix (e.g., clauses list for fixed price service contract). Contracting Officers often resort to copying and pasting many pages of clauses from an old contract. The former method was time consuming. Pulling clauses from an old contract failed to account for unique characteristics of each contract and missed new clauses mandated by law or regulation.

Clause Logic Software at Federal Agencies

The Department of Defense's (DoD) Clause Logic Service (CLS) implements a centralized, rules-driven engine that asks contracting personnel a structured set of questions and then returns the applicable provisions/clauses (including alternates) based on clause prescription text from regulations. The CLS is exposed through the Procurement Integrated Enterprise Environment (PIEE) and supports both human users and system-to-system calls, emphasizing consistency and auditability across the DoD. Its use for new DoD contract-writing systems has been mandated by policy memoranda, with guidance and user manuals describing how updates propagate and how the Q&A workflow operates.

Appian's Government Clause Automation solution takes a platform approach. It combines a centralized repository of FAR/DFARS and agency-supplement language with guided, configurable decision support to help users assemble clause packages. Public documentation describes a low-code application that agencies tailor to their supplements and workflows, providing a structured experience to generate clause supplements and speed reviews; marketing materials position this as "intelligent clause automation."

Unison PRISM embeds clause assistance inside a widely deployed federal contract-writing system. The PRISM 7.6 release adds generative-AI clause guidance intended to recommend applicable clauses in context, alongside no-code configurability and Microsoft 365 integration; agency materials note PRISM's broad use as an end-to-end acquisition system. While the AI feature promises faster selection, agencies still need governance to ensure version/deviation fidelity.

The Defense Logistics Agency (DLA) selected Icertis Contract Intelligence (ICI) as an enterprise contract management platform to modernize contracting and create a single, scalable workflow that integrates with DLA's SAP Procurement for Public Sector environment; with centralized templates and automated controls (including clause/template standardization), claimed benefits a reduction in contract cycle times, while also reducing contract errors through standardized templates and improving reporting speed through AI-powered dashboards.

Contracting Officers at federal agencies aim to include all applicable clauses in a solicitation or contract. Clauses are prescribed by law, regulation, and agency policy. Generally official clause titles, effective dates, and text must be used, but occasionally there is some discretion allowing tailoring or customizing of clause language. Due to frequent policy changes, An AI model trained on labeled historical data (contracts) is not a prudent approach in federal acquisition. Errors with clause inclusion are commonplace in individual contract documents.

Even if historical contracts were labeled accurately new policies and clauses would make this old data irrelevant. When Congress and agencies decide it's important to change a policy via a contract clause—there is normally an implementation lag time in new contracts. The Clause Tool measures the frequency of the Wrong [Clause] Date being observed in contract documents. For example, the chart below shows compliance metrics for FAR 52.204-25 clause inclusion and illustrates the typical lag in implementing an updated clause. The orange color on the stacked column chart shows failure to include the newly updated telecommunications security language mandated by National Defense Authorization Act.

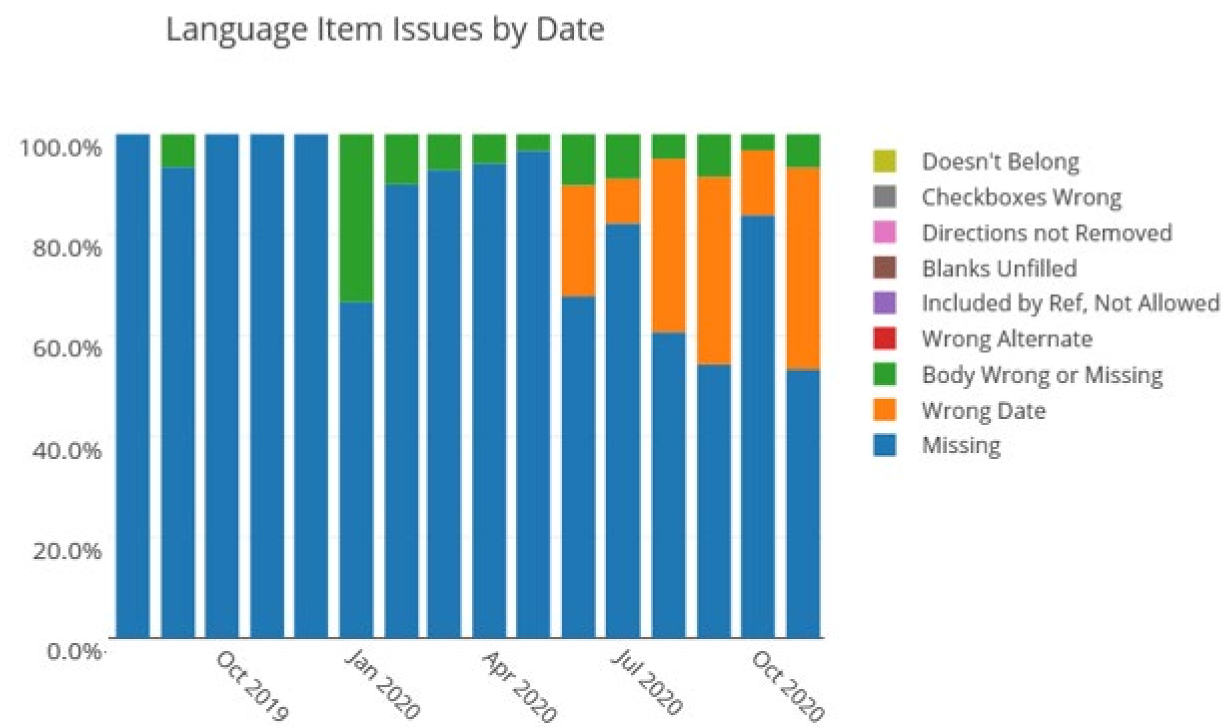

Figure 1. Language Item Issues by Date

Innovating at the IRS with the Contract Clause Review Tool

In 2017, the author initiated a project to rethink and improve clause logic software for Contracting Officers at the Department of the Treasury, Internal Revenue Service. The vision was for software to automatically examine the words in a draft contract and output the correct legal clauses. Software can quickly apply simple rules—like an old, obsolete clause shouldn't be in a new contract. Clauses have publication dates and procurement policy experts have lists of valid clauses and obsolete clauses. Further many different traits can impact whether a clause is applicable and must be included in a contract. For example, applicability criteria includes contract dollar value, contract type, geographic location, intellectual property, small business status of the contractor, etc.

The IRS held a competition and received six vendor proposals to build clause software. A key design principle was producing software that employees would want to use. Once a contract was awarded, the availably of FedRAMP certified cloud infrastructure enabled the small business vendor to rapidly deploy a web application for a modest cost. Employees could simply upload solicitation or contract documents to the web application and receive instant clause recommendations. Since deployment in 2017, more than 200,000 clause errors have been identified and corrected.

The IRS's use of clause review technology illustrates an "inspection/QA" approach. Fercent's DocScout is positioned to scan solicitations/contracts, flag potentially applicable or missing clauses, and check whether included text is current and properly completed; recent reporting indicates the IRS awarded a contract specifically for a "DocScout Contract Clause Review Tool."

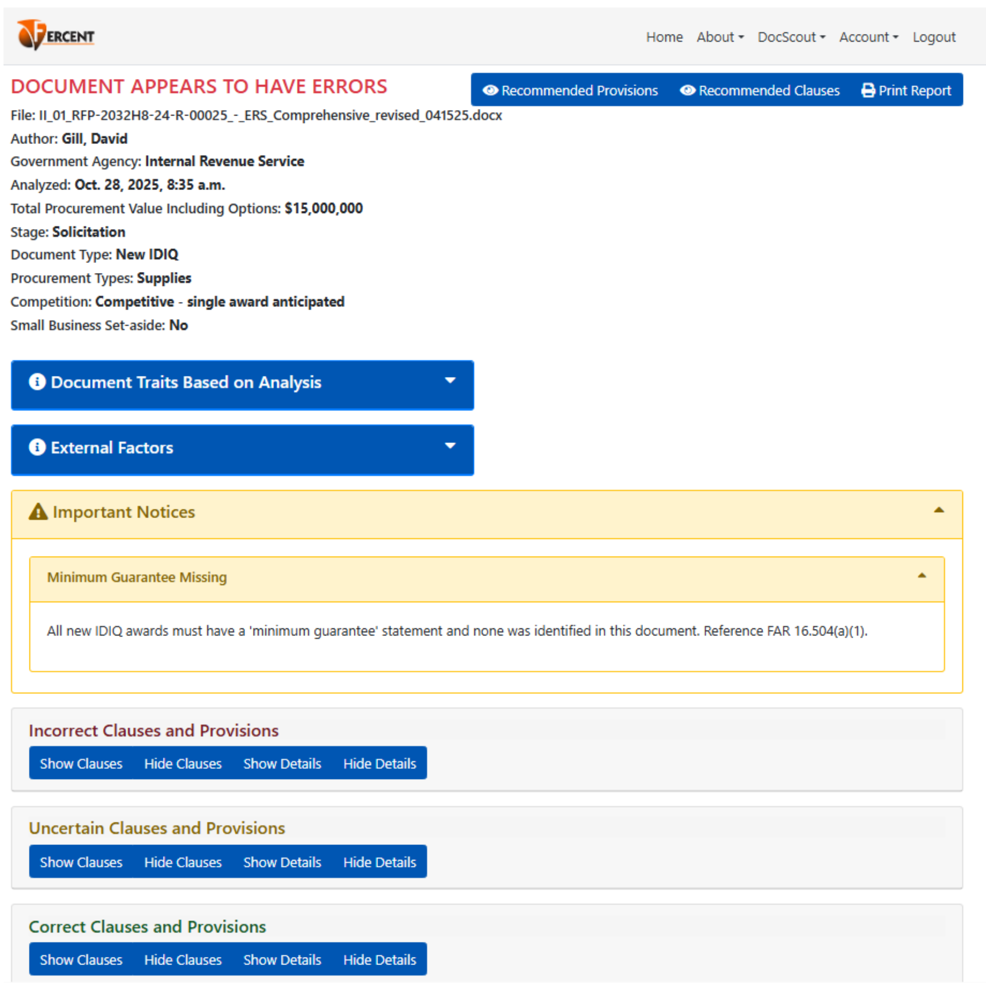

Figure 2. Screenshot of DocScout Contract Clause Review Tool

Detailed analytics are produced quantifying clause errors in uploaded contract documents. A variety of error types are shown such as wrong date / version, doesn't belong / not applicable, and inclusion by reference not allowed because full text of clause must be inserted. The below chart shows error rates for a specific clause.

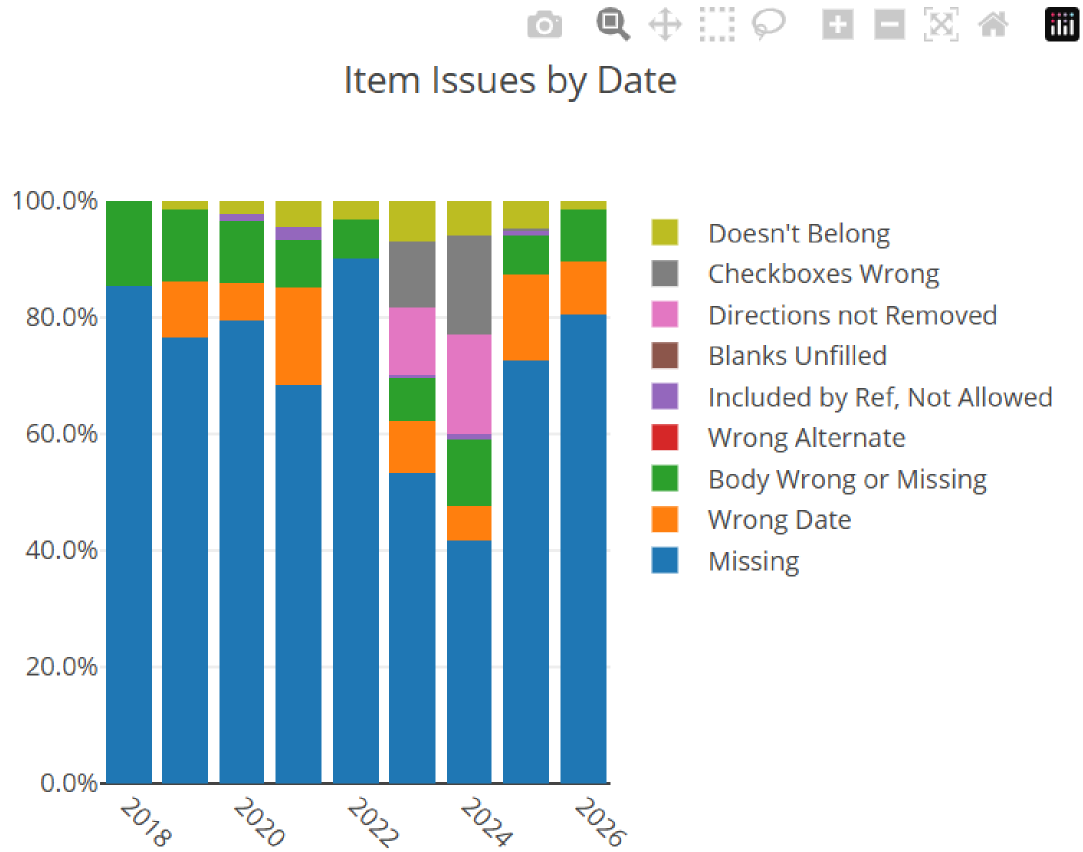

Figure 3. Item Issues by Date

The numerous prescriptions contained in the FAR regulations can be thought of as variables that trigger clause applicability. Some clause logic software uses a question-and-answer wizard approach to elicit the traits of an acquisition from the user. Questions include contract dollar value, commercial vs. noncommercial, type of contractor entity, etc. Asking questions burdens the user's time there are a large number of acquisition traits (moreo than 100 variables) that can trigger clauses. The FAR Smart Matrix offers a web-based tool for identifying clauses. While helpful, the FAR Matrix only accounts for a limited number of acquisition traits (e.g., Fixed-Price Supply). Further the FAR Matrix has many checkboxes and drop-down menus that can require a significant amount of time for a Contracting Officer to fill-out. In contrast, robust clause automation can produce more complete and tailored clause recommendations in a shorter amount of time.

Currently the Clause Tool automatically detects 114 acquisition traits by analyzing the words in contract documents. The Clause Tool has been used to inspect thousands of contract documents and the automated classification has avoided asking nearly 250,000 questions of users. The result was significant time savings for Contracting Officers with approximately 100 questions automatically answered without user input each time a contract document is inspected by the software.

Table 1. Feature Matrix—Software Currently Used by Federal Agencies

Feature	Traditional Contract Writing System	Contract Clause Review Tool	Commercial-Off-The-Shelf AI Chat (e.g., ChatGPT, Llama)
Immutable Clause Text	Yes	Yes	No (unauthorized deviations from official clause text)
Application programming interface (API) with Official Federal Regulations (acquisition.gov or ecfr.gov)	Yes	Yes	Can integrate with APIs, but interfaces are often disabled by federal agencies.
Can Recommend Clauses Based on Automatic Analysis of Words in Solicitation or Contract Document	No	Yes (Note: currently only matches exact keywords)	Yes (Note: flexible and not limited by pre-defined keywords)
Can Generate All Applicable Clauses for an Acquisition (e.g., 50 pages of text)	Yes, but significant manual user effort required	Yes, automated	No (struggles to produce lengthy documents due to token / word limits, typically only generates a few pages at a time).
Version Control	Yes	Yes	No (difficulty discerning the correct version of a clause? For example, standard FAR clause versus FAR Overhaul deviation version).

Clause Software Focused on Private Sector Contracts

Innovative software approaches are being used to negotiate clause language in private sector contracts. Private sector contracting parties have different goals and rules of engagement when compared with federal acquisition. Companies have greater latitude to negotiate clause language. For example the allocation of liability between the parties. Liability allocation in the private sector is driven by market leverage, risk tolerance, and business priorities. Federal contractors are prohibited from granting open-ended indemnification. It’s quite different for a Contracting Officer who is normally required to insert official clause language verbatim—without edits or revisions (* except for some limited CO discretion—e.g., clause tailoring or choosing between different official clauses). Despite the differences, contract boilerplate software approaches from private sector focused vendors can serve as inspiration for improving clause logic automation in federal acquisition.

TermScout positions itself as a contract-benchmarking and certification platform that analyzes agreements as market data rather than as isolated redlines: its engine extracts several hundred standardized datapoints per contract, compares clause-level positions across large curated datasets, and produces favorability ratings and a TrustMark/Certify badge when agreements meet predefined fairness thresholds and contain no designated "deal-breaker" clauses; the output is a market-aligned, third-party signal aimed at accelerating sales and procurement decisions by telling contracting officers whether a vendor's paper is market-standard, customer-favorable, or requires focused negotiation rather than providing line-by-line redlines.

Ironclad embeds AI across a full contract life cycle management (CLM) stack—clause detection, Smart Import, AI Playbooks, drafting and redlining assistance, automated routing/approvals, and analytics—providing out-of-the-box AI clause detectors plus customer-trainable custom clauses and properties so the system's models and playbooks refine over time with organization-specific data; the platform's goal for procurement and legal teams is to centralize negotiation workflows, enforce clause logic close to operational approvals, and surface repository intelligence to shorten review cycles and standardize risk handling across sales, procurement, and legal.

BlackBoiler focuses on automated redlining in Word that "learns from your playbook and prior edits," applying preferred insertions/deletions rather than randomly AI generating fresh prose; the company emphasizes an ensemble approach to avoid hallucination and preserve verbatim language fidelity. In a blog post BlackBoiler explains their approach and reasons for avoiding unconstrained generative AI for clause text:

"For many organizations, the promise of AI in contract review has been both exciting and frustrating. Generative AI tools can summarize documents or draft clauses. But when it comes to editing contracts in a way that reflects a company's unique risk posture, generic AI models fall short. They guess. They drift. They require constant prompt engineering. And they rarely deliver repeatable results at scale. The reality is simple: contract editing is only valuable when it reflects your standards, not the Internet's.

That's where BlackBoiler's approach to rule-driven contract automation changes the game. Instead of relying on generic AI to figure out what "good" looks like, BlackBoiler applies your rules, your negotiation patterns, and your historical edits to every contract. The result is predictable, compliant, and consistent redlines produced in minutes, without requiring teams to overhaul their workflows or adopt new contracting processes."

Emerging AI Research on Clause Automation

Practitioners can improve their clause logic applications by leveraging insights from a growing global body of research. Across these studies, a clear progression emerges in how AI supports contract analysis for acquisition: earlier approaches focus on identifying and extracting structured elements—such as clause types, pricing, and deliverables—from unstructured contract text, essentially acting as an automated reader; newer approaches move beyond identification to interpreting meaning, determining whether a contract satisfies specific legal conditions and pointing to supporting language, which aligns more closely with how Contracting Officers actually review contracts; parallel improve consistency, retrieval, and drafting, enabling systems to link obligations, find precedent clauses, and produce auditable outputs tied to source text; however, studies on Legal LLMs highlight that general-purpose models still lack reliability without domain-specific training and human oversight. Finally, implementation-focused research shows two distinct deployment models: the Air Force Institute of Technology approach uses LLM prompting to directly extract contract data and automate analysis at scale, while the NPS approach improves human decision-making by clarifying questions within systems like CLS. Taken together, the literature shows that AI in contracting is evolving from basic text extraction to reasoning, compliance checking, and decision support, with the most practical near-term use being a hybrid model where AI accelerates review and highlights risks, while Contracting Officers retain final judgment.

Two studies illustrate how Natural Language Processing (NLP) can be applied to extract meaningful traits from contract documents in ways that are relevant to acquisition professionals. The Contract Understanding Atticus Dataset (CUAD) study develops AI models trained on contracts that have been carefully labeled by legal experts, enabling the system to recognize common clause types—such as termination, indemnification, or governing law—and highlight the exact text where those clauses appear. This approach teaches the system to reliably identify standardized "contract features" based on prior examples. In contrast, *Enhancing Contract Management through Natural Language Processing (NLP): A Case Study of Three African Countries* applies NLP in a more operational setting by collecting procurement contracts from government websites, translating them when necessary, and using text analysis techniques to extract key attributes like pricing, locations, and deliverables. For a Contracting Officer, both studies demonstrate that NLP can function as an automated reader of contract documents—either by learning from labeled clause examples (CUAD) or by scanning real-world contracts to pull out important data elements—an approach that could be used to see if something in a document triggers applicability of a FAR clause prescription.

A study on construction contracts pairs LLM with a domain knowledge graph—a simple, structured map of the important concepts in construction contracts (nodes like parties, obligations, milestones, penalties) and the labeled links between them (who owes what, when, and what happens if they fail). The graph is used to guide the LLM: it helps the model find and label contract elements consistently, checks that extracted items fit the expected categories, and records where each finding came from. For contracting officers this produces clearer, auditable outputs (fewer missed obligations or vague risk flags), machine-friendly summaries that plug into procurement checklists, and traceable links back to the contract language for defensible decisions. When tested on construction contracts, the combined approach outperformed the LLM alone at accurately pulling out complex, multi-party clauses and made fewer invented or irrelevant results; however, its usefulness depends on keeping the knowledge graph up to date with regulations and contract practices and retaining human review for novel or high-risk provisions.

The ContractNLI paper introduces a practical way for AI to "reason" about contracts—not just find clauses, but determine what the contract actually means. In simple terms, the system is given a contract and a plain-language statement (for example, "Some obligations continue after termination"), and it must decide whether that statement is supported by the contract, contradicted by it, or not addressed at all, while also pointing to the exact sentences that justify the answer. This is called natural language inference (NLI). For a Contracting Officer, this is similar to asking: "Does this contract include a limitation on subcontracting?"—and having the system both answer and show the supporting clause. Unlike earlier tools (like CUAD) that only identify where a clause exists, ContractNLI focuses on interpreting the legal effect of the clause across the whole document, which is closer to how contract reviews are actually performed. The study also shows that this task is difficult for AI because contracts contain complex language like exceptions and conditions, but it demonstrates that models can be trained to improve this reasoning capability.

Subsequent research builds on ContractNLI by improving how systems perform this type of reasoning. For example, later work explores transfer learning, where models trained on general language tasks are adapted to contract analysis, improving accuracy when data is limited. Other follow-on uses of ContractNLI (such as LegalBench tasks) simplify the problem into binary compliance checks (e.g., “Does this clause require return of confidential information?”), making it more directly usable for automated policy checks in contracts . Across these studies, a consistent pattern emerges: modern contract AI is evolving from simply finding clauses to evaluating whether a contract satisfies specific legal requirements and providing traceable evidence, which aligns closely with real-world acquisition review tasks.

A proceedings paper at the Association for Computational Linguistics introduced the Atticus Clause Retrieval Dataset (ACORD), the first expert-annotated benchmark specifically designed for contract clause retrieval to support contract drafting tasks. ACORD focuses on complex contract clauses such as Limitation of Liability, Indemnification, Change of Control, and Most Favored Nation. ACORD is about finding the best precedent clause, not just classifying it. The benchmark has 114 queries and more than 126,000 query-clause pairs, and it focuses on high-value boilerplate like limitation of liability, indemnification, change of control, and most-favored-nation language. For a Contracting Officer, this is especially relevant because it points toward AI that can search a clause library the way an experienced negotiator does.

The paper “Large Language Models are Legal but They Are Not: Making the Case for a Powerful LegalLLM” addresses a key issue for anyone working with contracts: today’s AI tools can appear capable of legal analysis, but they are not yet reliable enough to be treated as true legal decision tools. The authors tested several well-known large language models (like ChatGPT and LLaMA) on a task similar to what a Contracting Officer does—identifying and classifying contract clauses (e.g., what type of provision a paragraph represents). They found that general-purpose models can often get the gist right, but they are significantly less accurate than smaller models specifically trained on legal data—by as much as ~20–27% . In plain terms, this means the AI may correctly recognize a clause sometimes, but will miss or misclassify enough cases to create real risk if used without oversight.

A Naval Postgraduate School study looked at whether generative AI (like ChatGPT) can help make the Clause Logic Service (CLS) easier for Contracting Officers to use by rewriting confusing questions in plain language and adding short, practical examples. The researchers found that many CLS questions are hard to interpret and rely too much on individual judgment, which can lead to inconsistent clause selection and risk of missing required FAR/DFARS clauses. By using AI to simplify wording and add “helpful hints,” most users in the study said the questions became easier to understand and apply, and their accuracy improved when answering them. The CLS approach contrasts with DocScout which is designed to minimize the number of user questions in favor of speed.

An Air Force Institute of Technology study took a different approach from the NPS work: instead of asking the user (you, the Contracting Officer) questions, it has the AI read the contract itself and answer questions using carefully designed prompts—essentially telling the model, “you are a contracts expert—pull out specific data like services, costs, triggers, and dates,” and then feeding it sections of the contract to extract those answers automatically. In simple terms, AFIT treats the contract like a document to mine for facts, using steps like “classification” (deciding if a section is relevant, similar to filtering out noise), and “information extraction” (pulling out key fields like contract line item number (CLIN)-type data into a structured format, like turning narrative text into a spreadsheet). By contrast, the NPS approach assumes the system cannot fully interpret the contract on its own, so it improves the *questions asked to the human* (e.g., clearer CLS prompts and examples) to get better inputs. Key data science terms in the AFIT study can be understood simply: “unstructured data” just means normal contract text (not organized like a table), “metadata” means key contract facts you care about (like surge requirements, cost, or triggers), “prompting” means giving the AI instructions and examples so it knows what to look for, and “hallucination” means the AI sometimes confidently gives wrong answers if the question or context is unclear. This approach replaces manual review with automation by having AI *read and extract* from contracts directly, while NPS is trying to improve how humans *interact with a rules-based system* to select the right clauses.

Sharing Innovative Procurement Technology

Adopting “modernized data analytics, and advanced technologies that allow decision making to occur in a more friction-free buying environment” is a goal per the Promoting Rigorous and Innovative Cost Efficiencies for Federal Procurement and Acquisitions (PRICE) Act. Pursuant to the PRICE Act, the Office of Management (OMB) submitted a report sharing innovative procurement technology approaches with the Senate Committee on Homeland Security and Government Affairs. In the PRICE Act Report OMB highlighted the value of automated clause software where “relevant clauses are identified primarily by having the software examine the words in an uploaded document, in addition to a user answering some questions. Missing, outdated, or otherwise erroneous clauses are rapidly identified by the software based on the clause recommendation logic.”

Clause logic can be viewed as a microservice or module that can be integrated into a larger contract writing system. Modular open system approaches to system architecture allow the government to combine various components into a single system. Adding a modular system component enables the delivery of better software functionality while integrating with an existing system. This approach reduces vendor lock-in and reliance on monopolistic procurement system contractors.

Agencies should modernize legacy contract writing systems replaced poorly-designed, labor intensive clause logic software applications. Clauses implement critical public policies mandated by law, regulation, and agency policy. Omitted mandatory clauses, failing to incorporate new policies, and other clause errors should not be commonplace. Better clause software is needed to reduce the frequency of clause errors and create legally enforceable responsibilities for contractors. Ensuring that a complete set of applicable official clauses is included in a contract should be automated.

Appendix—Automatically Classified Acquisition Traits in DocScout

Clause automation software can detect the presences of these traits using keywords in solicitation and contract documents. Each trait triggers one or more clause prescription. The percent switched metric indicated how often users changed the classification (i.e., due to incorrect automatic classification by the software). The % Switched metric is likely an undercount of misclassification error because it only shows when an acquisition professional chose to override a misclassification in the system.

Classifier Title	% Switched
Firm Fixed Price	18.84%
Commercial Item	10.85%
Commercial Off-The-Shelf (COTS)	9.76%
Key Personnel Involved	8.08%

Appearance of Personal Services	6.24%
Purchasing IT Services, Software, and/or Hardware	6.10%
Period of Performance	6.02%
Accessibility 100% Compliance	5.71%
Government Site Performance	5.52%
Performance-based Acquisition Approach	5.42%
Government Furnished Property	5.42%
Electronic Deliverables	5.07%
F.O.B. Destination Only	4.38%
Mandatory Source Service	4.29%
Labor Hours	4.20%
Has Option Periods	3.99%
Staff-like Access	3.87%
IT Hardware other than PCs	3.58%
Subject to Commercial Supplier Agreement (CSA, EULA, TOS)	3.58%
Least Cost Selection	3.56%
Classified Information Involved	3.20%
PII Access	3.13%
SBU Information Involved	3.11%
GSA Supply Schedule Order or BPA	2.83%
Brand Name or Equal	2.73%
Separate SOW or SOO	2.69%
Cost Plus Fixed Fee	2.68%
USD Payment	2.64%
Technical Specification	2.58%
EPA Designated Product	2.35%
Best Value Selection	2.28%
Subcontracting Plan Involved	2.07%
Mandatory Source Supply	2.06%
Cost Price Evaluation	1.96%
Personal Computers Procurement	1.82%
Educational Institution Contract	1.74%
IT Development	1.69%
Performance and Delivery Outside United States	1.53%
IT Access for Contractor	1.46%
Matter of National Security	1.45%
Biobased Product Preference	1.40%
Product Quality	1.37%

Key Facilities Involved	1.32%
Time and Materials	1.28%
U.S. Citizens and Residents Only	1.18%
Textile Product	1.13%
Office Imaging Equipment	1.07%
Telework Allowed	0.87%
Cost Realism	0.87%
Cost-sharing Contract	0.70%
Travel Reimbursed	0.68%
Hazardous Materials	0.66%
Television Procurement	0.62%
Incremental Funding	0.55%
Custom Data Rights Language	0.54%
Construction Procurement	0.50%
Cost Least Important	0.46%
Construction Material Procurement	0.45%
U.S. Citizens Only	0.42%
Staff-like Access, Reduced Background	0.36%
Staff-like Access, Intermittent	0.30%
Funded by External Treasury Office	0.29%
Appliance Procurement	0.29%
Architect/Engineering Procurement	0.29%
Cost Contract	0.28%
Subcontracting Plan With Proposal	0.27%
Heating & Cooling Procurement	0.21%
Top Secret Clearance	0.18%
Specific SBA District or Region	0.18%
Park and Recreation Product	0.17%
Limit One Proposal	0.14%
Cost Most Important	0.14%
Secret Clearance	0.14%
Oral Presentation	0.14%
F.O.B. Origin	0.09%
Economic Price Adjustment	0.09%
Cost-plus-fixed-fee Contract	0.09%
Sealed Bidding	0.09%
Personal Services Contract	0.09%
Landscaping Product	0.08%

F.O.B. Origin, Freight Allowed	0.05%
F.O.B. Origin, Contractor Facility	0.05%
Mitigable Conflict of Interest	0.05%
Pre-contract Costs	0.05%
Confidential Clearance	0.05%
Funded by Bureau of Engraving and Printing	0.04%
Transportation Services Procurement	0.04%
Vehicular Product	0.04%
Minimum Guarantee Found	0.04%
Bid Guarantee	0.00%
No Currency Adjustment	0.00%
Common Carrier Communication Services	0.00%
Matter of National Interest	0.00%
Definitized Letter Contract	0.00%
Transportation Product	0.00%
Confinement Facility	0.00%
Swift System	0.00%
Cost-plus-award-fee Contract	0.00%
Cost-plus-incentive-fee Contract	0.00%
Nonprofit Blind Disabled	0.00%
Cost Equally Important	0.00%
Federal Prison Industries is Eligible	0.00%
R&D Procurement	0.00%
Alternate Proposals Allowed	0.00%
F.O.B. Origin Only	0.00%
Federal Protective Service Guard Services	0.00%
Food Provision, Service, or Sale	0.00%
Grounds Maintenance	0.00%
Unmitigable Conflict of Interest	0.00%
Lead System Integrator Involved	0.00%
VAT Exempt	0.00%
Airport Related	0.00%
Cost No Fee	0.00%
Overseas Supplier	0.00%

References

Appian. (2025). *Clause intelligence—Government clause automation*. https://appian.com/industries/public-sector/acquisition-management/government-clause-automation

Armed Services Board of Contract Appeals. (2002). *Griffin Services, Inc., ASBCA Nos. 52280, 52281* (decision). https://www.asbca.mil/Portals/143/Decisions/2002/52280_52281.pdf

Babatunde, I., Ezirim, A., & Hountondji, O. (2023). *Enhancing contract management through natural language processing (NLP): A case study of three African countries*. ResearchGate. https://www.researchgate.net/publication/373140351_Enhancing_Contract_Management_through_Natural_Language_ProcessingNLP_A_Case_Study_of_Three_African_Countries

Berry, C. M., Briggs, R. W., & Aragon, M. A. (2025). *Can generative AI improve the quality of interview questions within CLS while adhering to plain writing principles?* Naval Postgraduate School. https://calhoun.nps.edu/entities/publication/39ba922b-3c71-4612-bdca-5da08fe3e281

BlackBoiler. (2026). *Product overview*. *When legal contract editing reflects your rules and not generic AI, consistency, speed, and risk control align*. https://www.blackboiler.com/contract-editing-your-rules-vs-generic-ai/

Defense Pricing and Contracting. (2024). *Clause Logic Service user interface user manual (v6.5.3 V41)*. DoD Procurement Toolbox. https://dodprocurementtoolbox.com/uploads/Clause_Logic_System_User_Manual_6_5_3_V41_20240201_7148f82f30.pdf

General Services Administration. (2026). FAR smart matrix. Smart Matrix | Acquisition.GOV

Griffin Services, Inc., ASBCA Nos. 52280, 52281, 04-2 BCA ¶ 32,714 (2004). https://www.asbca.mil/Portals/143/Decisions/2004/52280_52281.pdf

Hazy Research. (2024). *LegalBench: ContractNLI task (return of confidential information)*. GitHub. https://github.com/HazyResearch/legalbench

Hendrycks, D., Burns, C., Chen, A., & Ball, S. (2021). *Contract Understanding Atticus Dataset (CUAD): An expert-annotated NLP dataset for legal contract review*. arXiv. https://arxiv.org/abs/2103.06268

Icertis. (2026, February 12). *U.S. Defense Logistics Agency selects Icertis Contract Intelligence to meet federal mandate for digital transformation*. https://www.icertis.com/company/news/u.s.-defense-logistics-agency-selects-icertis-contract-intelligence-to-meet-federal-mandate-for-digital-transformation

Internal Revenue Service, Acquisition Gateway. (2023, December). *Overview—Treasury IRS contract clause review tool*. https://ag-dashboard.acquisitiongateway.gov/system/files/ptai/2023-12/Overview%20Treasury%20IRS%20Contract%20Clause%20Review%20Tool.pdf

Ironclad. (2026). *AI-based contract management*. https://ironcladapp.com/product/ai-based-contract-management

Jayakumar, T., Farooqui, F., & Farooqui, L. (2023). *Large language models are legal but they are not: Making the case for a powerful LegalLLM*. arXiv. https://arxiv.org/abs/2311.08890

Koreeda, Y., & Manning, C. D. (2021). *ContractNLI: A dataset for document-level natural language inference for contracts*. arXiv. https://arxiv.org/abs/2110.01799

Office of Management and Budget. (2024). *PRICE Act report to Congress—Acquisition innovation & small business participation in federal procurement*. National Archives.

NCS/EML JV, LLC, B-412277; B-412277.2; B-412277.3 (Comp. Gen. Jan. 14, 2016) [Government Accountability Office]

Rivera, Alex. (2013). How contract boilerplate can bite: Part II. *Rocky Mountain Mineral Law Foundation Journal, 50*(2), 243–306. https://heinonline.org

Tang, Y. K. (2023). *Natural language inference transfer learning in a multi-task contract dataset: In the case of ContractNLI*. Uppsala University. https://www.diva-portal.org/smash/get/diva2:1801289/FULLTEXT01.pdf

TermScout. (2026). TrustMark™: Independent contract certification. https://www.termscout.com/trustmark

Tuggener, D., von Däniken, P., Peetz, T., & Cieliebak, M. (2020). LEDGAR: A large-scale multi-label corpus for text classification of legal provisions in contracts. In *LREC 2020* (pp. 1235–1241). https://aclanthology.org/2020.lrec-1.155.pdf

Unison. (2025, February). *Unison releases PRISM 7.6 with added AI capabilities to advance federal acquisition efficiency and modernization*. Yahoo Finance. https://finance.yahoo.com/news/unison-releases-prism-7-6-130600415.html

Wang, S. H., Zubkov, M., Fan, K., Harrell, S., Sun, Y., Chen, W., Plesner, A., & Wattenhofer, R. (2025). ACORD: An expert-annotated retrieval dataset for legal contract drafting. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics* (pp. 24739–24762). Association for Computational Linguistics. https://aclanthology.org/2025.acl-long.1206/

Xu, W., Deng, Y., Lei, W., Zhao, W., Chua, T.-S., & Lam, W. (2022). ConReader: Exploring implicit relations in contracts for contract clause extraction. In *EMNLP 2022*. https://arxiv.org/abs/2210.08697

Yae, J. H. (2024). *A staged framework for LLM-powered information extraction in government contracts* [Master's thesis, Air Force Institute of Technology]. https://scholar.afit.edu/etd/7734/

Acquisition Research Program
Department of Acquisition, Finance, and Manpower
Naval Postgraduate School
555 Dyer Road, Ingersoll Hall
Monterey, CA 93943

www.acquisitionresearch.net