

CLEARED
For Open Publication

2
Apr 29, 2026

Department of War
OFFICE OF PREPUBLICATION AND SECURITY REVIEW

06 April 2026

Contract No. HQ0034-24-D-0023

Task Order No. 048

NLP Recognition and Categorization of Acquisition Text and Documents

Jose E. Ramirez- Marquez, Akram Amer, Joshua Gorman, Douglas Buettner
Stevens Institute of Technology
1 Castle Point Terrace, Hoboken, NJ 07030

Brian Mayer, Nathan Self, Harith Laxman, Naren Ramakrishnan
Virginia Tech
300 Turner St NW, Suite 4200, Blacksburg, VA 24061

Sponsor: Ms. Kimberly Ziegler, Director for Defense Acquisition Regulations System (Acting)
Ms. Nickee Abbott, Strategic Advisor for Mission-Based Acquisition & Data Analytics
Office of the Under Secretary of War for Acquisition and Sustainment

Department of War
U.S. Department of Defense

Cleared for Public Release.

The views, findings, conclusions, and recommendations expressed in this report are solely those of the authors and do not necessarily reflect the views or positions of the United States Government (including the War Department [Department of Defense (DoD)] and any government personnel), the Stevens Institute of Technology, or Virginia Tech.

Final Technical Report AIRC-2026-TR-005
WRT-2507.11B

Copyright © 2026 Stevens Institute of Technology and Virginia Tech. All rights reserved.

The Acquisition Innovation Research Center (AIRC) is a multi-university partnership led and managed by the Stevens Institute of Technology and sponsored by the U.S. Department of Defense (DoD) through the Systems Engineering Research Center (SERC)—a DoD University-Affiliated Research Center (UARC).

This material is based upon work supported, in whole or in part, by the U.S. Department of Defense through the Office of the Under Secretary of War for Acquisition and Sustainment (OUSW(A&S)) and the Office of the Under Secretary of War for Research and Engineering (OUSW(R&E)) under Contract HQ0034-24-D-0023, Task Order (TO) 048.

No Warranty. This Material is furnished on an “as-is” basis. The Stevens Institute of Technology and Virginia Tech make no warranties of any kind—either expressed or implied—as to any matter, including (but not limited to) warranty of fitness for purpose or merchantability, exclusivity, or results obtained from use of the material. The Stevens Institute of Technology and Virginia Tech do not make any warranty of any kind with respect to freedom from patent, trademark, copyright, or other intellectual property infringement.

Abstract

This research builds upon prior efforts by the authors to streamline Defense Federal Acquisition Regulation Supplement (DFARS) rule development from National Defense Authorization Act (NDAA) text using artificial intelligence (AI)-based tools. The current research effort focuses on developing a unified, web-based interface that integrates previously developed prototypes and incorporates advanced Natural Language Processing (NLP), Large Language Models (LLMs), and Machine-Based Reasoning (MBR) techniques to improve automation in identifying, extracting, and recommending regulatory language changes.

The proposed unified system connects modules for document ingestion, keyword and context identification, text summarization, clustering, and visualization—through an integrated backend and user interface, which will enable Defense Pricing, Contracting, and Acquisition Policy (DPCAP) staff to move seamlessly from NDAA review to DFARS draft generation. The tool also proposes novel MBR techniques that leverage LLM models for updating proposed rule language and summarizing public comments from Regulations.gov. The combined framework is being deployed on a secure, sponsor-accessible server and will be evaluated against real DFARS updates in collaboration with DPCAP subject matter experts (SMEs). This work aims to significantly reduce manual analysis time, enhance the traceability of regulatory updates, and strengthen the Department's capacity to apply AI responsibly in acquisition policy modernization.

Executive Summary

This research demonstrates the feasibility and value of integrating Natural Language Processing (NLP), Machine-Based Reasoning (MBR), and Large Language Models (LLMs) into a unified framework to support the Defense Federal Acquisition Regulation Supplement (DFARS) rule development process. This effort evaluated, refined, and consolidated previously independent artificial intelligence (AI)-driven capabilities into a single, end-to-end system available to the customer for testing. The result of the research is a more streamlined process for National Defense Authorization Act (NDAA) analysis and DFARS rule drafting/revision.

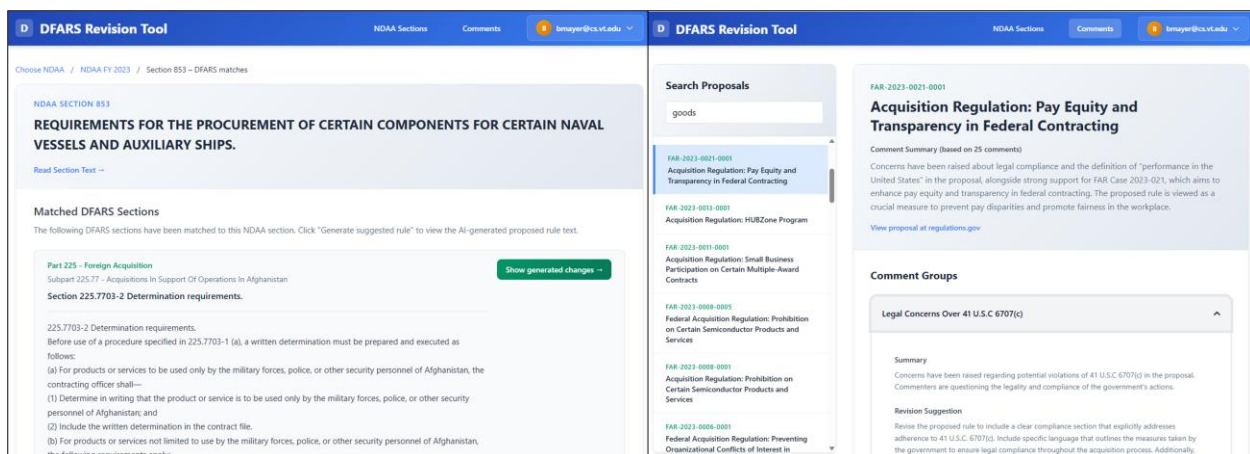
The DFARS text generation process was improved by using a system of LLM agents that helped validate the match, provided additional context, and generated a better draft DFARS section. The multi-agent process forces the agents to perform tasks with more reasoning, which consequently provides additional explainability to the user. This explainability could be further leveraged by the rest of the system and interface for further improvements and adoption.

Evaluation of the entire system found that:

- More than 90% of the sections of interest in each NDAA are extracted.
- Most of the matching DFARS sections (to the NDAA section) are identified, most of the time.
- DFARS draft text generated for 1:1 matches (1 NDAA section to 1 DFARS section) is representative of the actual finalized DFARS section.

The research project resulted in an integrated system for testing. The system is made available through a publicly accessible interface behind a secure login. This also enables a feedback loop to identify improvements.

Figure 1. Screenshots of the System Interface



NOTE: Left: DFARS Draft Generation. Right: Comment Summaries.

Future opportunities for improvement were identified based on the current limitations of the system. Further fine-tuning of the matching process will ensure the model is not over-fit and that it will continue to provide worthwhile results. Modifications to the pipeline may also enable more sophisticated methods of processing changes that have multiple NDAA sections or multiple DFARS sections in a single “case”. The Federal Acquisition Regulation (FAR) Drafting Guide [1] will be used to refine the format of the generated text.

Overall, this work highlights the transformative potential of AI-assisted systems in regulatory environments, offering a scalable, semi-automated approach that augments human expertise while maintaining the rigor required for defense acquisition policy development.

Table of Contents

Abstract.....	iii
Executive Summary	iv
Table of Contents.....	vi
List of Figures	vii
List of Tables	vii
Acknowledgements.....	viii
Research Team	viii
Acronyms and Abbreviations.....	ix
1. Introduction	1
Organization of the Report	1
2. DPCAP FAR/DFARS Change Process	3
Step 1: NDAA Section Extraction and Storage	4
Step 2: NDAA ↔ FAR/DFARS Matching & Ranking.....	5
Step 3: FAR/DFARS Text Changes	5
Step 4: Publish Proposed Text.....	6
Step 5: Public Comment Review	6
Step 6: Use comments to make changes to the proposed text.....	6
3. Evaluation of Methods.....	7
NDAA Section Extraction	7
NDAA ↔ FAR/DFARS Matching	8
Sections with Published FAR/DFARS Mappings	8
Section-Level Coverage	9
Sections with Less Than 50% Mapping Coverage.....	9
Quality Index (%).....	9
Micro-Averaged Quality Index	9
Interpretation.....	9
Evaluation of FAR/DFARS Text Changes	10
4. Proposed Unified System.....	12
5. Improvements to Generation of FAR/DFARS Text Changes	18
6. Conclusions, Recommendations and Next Steps.....	20

Conclusions	20
Suggested Future Research and Next Steps	20
Final Thoughts	20
References	21

List of Figures

Figure 1. Screenshots of the System Interface.....	iv
Figure 2. Task Flow Diagram	3
Figure 3. The Rule-making Process: How It Works	4
Figure 4. Selecting or uploading an NDAA.....	13
Figure 5. Reviewing the NDAA and identified sections.....	13
Figure 6. Selecting an NDAA section to find DFARS matches for.	14
Figure 7. Reviewing DFARS matches and selecting a match to generate a modification for.	14
Figure 8. Reviewing generated text for the modification.....	15
Figure 9. Selecting a DFARS proposal and reviewing the summary of public comments.	16
Figure 10. Reviewing the summary of a comment group, the proposed revision based on the comments, and viewing the raw text from all the grouped comments.	17
Figure 11. Process for generating changes to a DFARS section given an old DFARS section and an NDAA section.....	18

List of Tables

Table 1. Quality Index for expected NDAA sections versus extracted NDAA sections.....	7
Table 2. Quality Index for NDAA section to FAR/DFAR section(s) mappings.	10
Table 3. LLM judged quality of system generated DFARS text.....	11

Acknowledgements

We would like to thank Ms. Dawn Messer and Ms. Kimberly Ziegler, Defense Pricing, Contracting and Acquisition Policy (DPCAP), for their sponsorship and support of this analytical effort.

The views, findings, and conclusions in this document are solely those of the authors and do not necessarily reflect the views or positions of the United States Government (including the War Department [DoD] or any other government personnel), the Stevens Institute of Technology, or Virginia Tech.

Research Team

Name	Organization	Labor Category
Jose Ramirez-Marquez	Stevens Institute of Technology	Principal Investigator (PI)
Brian Mayer	Virginia Tech	Co-PI
Douglas Buettner	Stevens Institute of Technology	Co-PI
Nathan Self	Virginia Tech	Co-PI
Naren Ramakrishnan	Virginia Tech	Co-PI
Joshua Gorman	Stevens Institute of Technology	Researcher
Akram Amer	Stevens Institute of Technology	Graduate Research Assistant
Harith Laxman	Virginia Tech	Graduate Research Assistant

Acronyms and Abbreviations

AI	Artificial Intelligence
AIRC	Acquisition Innovation Research Center
BLEU	Bilingual Evaluation Understudy
CoT	Chain of Thought
CP	Contract Policy
DFARS	Defense Federal Acquisition Regulation Supplement
DoD	Department of Defense
DoW	Department of War
DPCAP	Defense Pricing, Contracting, and Acquisition Policy
EO	Executive Order
FAR	Federal Acquisition Regulation
FY	Fiscal Year
GPT	Generative Pre-trained Transformer
LLM	Large Language Model
MBR	Machine-Based Reasoning
NDAA	National Defense Authorization Act
NIPR	Non-classified Internet Protocol Router
NLP	Natural Language Processing
OSD	Office of the Secretary of Defense
OUSW(A&S)	Office of the Under Secretary of War for Acquisition and Sustainment
OUSW(R&E)	Office of the Under Secretary of War for Research and Engineering
PDF	Portable Document Format
ROUGE	Recall-Oriented Understudy for Gisting Evaluation
SBA	Small Business Administration
SERC	Systems Engineering Research Center
SME	Subject Matter Expert
USDOL	U.S. Department of Labor

1. Introduction

The Defense Pricing, Contracting, and Acquisition Policy (DPCAP), Defense Acquisition Regulations System (DARS) Directorate in the Office of the Assistant Secretary of Defense (OSD) is responsible for implementing statutes, executive orders, and policies through the timely update of the Federal Acquisition Regulation (FAR) and Defense FAR Supplement (DFARS) and issuance of memoranda and guidance. Fundamentally, they enable operations through business systems and standards.

DPCAP is responsible for yearly updates to the DFARS based on changes in the National Defense Authorization Act (NDAA), Small Business Administration (SBA) rule changes, U.S. Department of Labor (USDOL) rule changes or through executive orders (EOs). Reading through the changes made necessary by these multiple sources to complete required DFARS updates is labor intensive for DPCAP staff. It also requires knowledge of all the rules in the FAR/DFARS to ensure that changes are made appropriately, and references are made to the correct sections of the FAR/DFARS.

Artificial Intelligence (AI) is a powerful tool that can automate many tasks and support humans with others, but it has its limitations. Therefore, system development has to be a deliberate process that should be guided by policy and end use. Limitations can include bias, explainability, and trustworthiness (i.e. the well-known Large Language Model (LLM) hallucination problem). [2-4] Proper policy and implementation can limit bias, increase accuracy, and improve human effectiveness. When implementing AI solutions, it is important to understand these limitations and to create environments where AI systems and humans work in tandem to obtain the best results possible. For many tasks, the critical importance of human judgment means AI should serve as a complementary tool to improve human efficiency rather than as a standalone solution.

This project resulted in a tool that provides a cost-effective, scalable, semi-automated capability for managing regulatory updates, while ensuring long-term efficiency and adaptability. This paper shares the process of evaluating, integrating, and improving AI methods developed to assist DPCAP subject matter experts (SMEs) in the FAR/DFARS change process into a unified system. SMEs are currently executing the tasks manually.

Organization of the Report

Section 1 introduces the problem and the objective of the project.

Section 2 reviews prior work that was completed to identify AI methods and processes to automate certain tasks in the FAR/DFARS change process and assist the DPCAP team.

Section 3 discusses the evaluation of those methods that were done to determine effectiveness and impact.

Section 4 provides a description of how all processes were then incorporated into a unified system and made available to the sponsor for testing.

Section 5 discusses how evaluation results and feedback led to the improvement of the text generation process for FAR/DFARS changes.

Section 6 covers the conclusion, recommendations and next steps.

A copy of the Interface User/Tool Guide was provided as a separate attachment to this report.

2. DPCAP FAR/DFARS Change Process

DPCAP SMEs go through a 6-step process (shown in Figure 1 and 2) to identify necessary FAR/DFARS changes and make necessary updates to the FAR/DFARS.

Figure 2. Task Flow Diagram

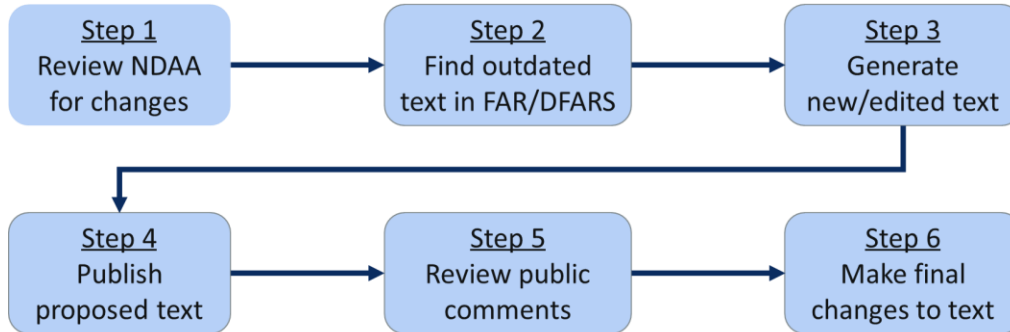
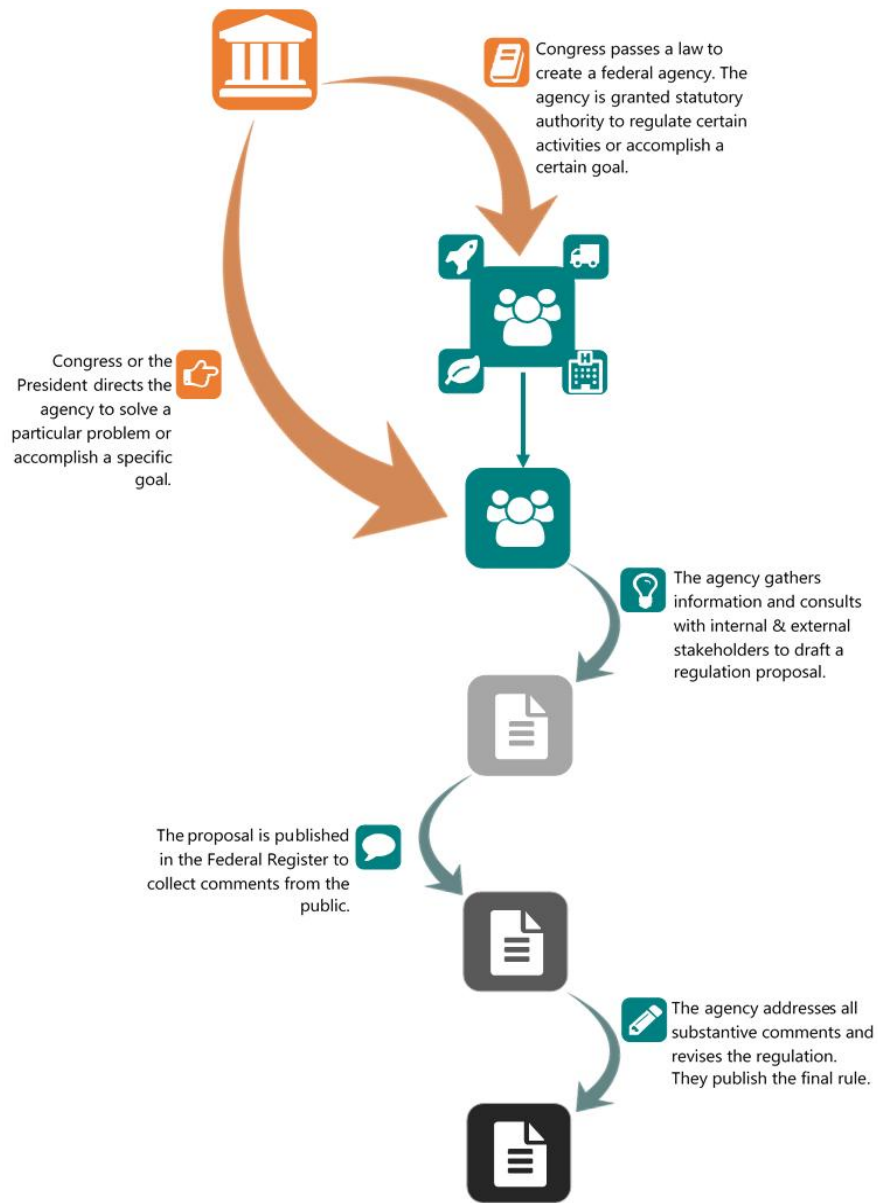


Figure 3. The Rule-making Process: How It Works



NOTE: <https://publiccommentproject.org/how-it-works>

Step 1: NDAA Section Extraction and Storage

The AI methods developed support this process by converting NDAA, FAR, and DFARS PDFs (portable document formats) into structured, searchable section records stored in the database and then identifying NDAA sections that may require changes to FAR/DFARS.

The process for accomplishing this is as follows:

1. Upload a document for a specific year (NDAA year or FAR/DFARS regulation year).
2. The system parses the PDF and splits it into individual sections.

- a. NDAA is split by section
 - b. FAR/DFARS are split by part → subpart → section
3. Each extracted section is stored in the database with key metadata, including:
- a. Section number
 - b. Title
 - c. Full text
 - d. Year
 - e. Document type (NDAA / FAR / DFARS)
 - f. For FAR/DFARS: part/subpart alignment
4. After the NDAA sections are stored, an analysis algorithm runs across all NDAA sections to identify which ones are likely to drive regulatory changes.
5. These flagged sections are saved as NDAA candidate sections, allowing the system to reuse the same candidate list without reprocessing the NDAA.

A file-type detection and parsing pipeline was developed, enabling reliable text extraction from both the modern and archived legacy document formats that were incorporated to address older PDF formats used for DFARS/FAR documents prior to 2020. These changes required substantial updates to PDF parsing algorithms.

Step 2: NDAA ↔ FAR/DFARS Matching & Ranking

The AI methods we developed support this process by identifying which FAR/DFARS sections are most likely impacted by each NDAA change. This is accomplished by taking a specific NDAA year and FAR/DFARS regulation year. Then each NDAA candidate section is compared against all FAR and DFARS sections for that regulation year. A scoring algorithm evaluates relevance and potential impact. The top-ranked FAR/DFARS sections (e.g., top 10) are selected per NDAA section. These ranked matches are saved and made available to the user. This resulted in a structured, ranked matrix connecting each NDAA candidate section to its highest-scoring FAR and DFARS provisions for the selected regulation year. For each NDAA section, the system returns a prioritized list (e.g., top 10) of potentially impacted regulatory sections, forming a decision-support dataset for regulatory impact analysis.

Step 3: FAR/DFARS Text Changes

This process takes the selected NDAA year/section and its matched DFARS section as input then generates candidate text for a modified rule and the corresponding text to be published with a proposed rule in the Federal Register. The tool uses Chain of Thought (CoT) prompting [5] and requests specific steps (defined by the research team) and outputs to be performed to accomplish the task. Documents may not be formatted to FAR/DFARS guidelines. The output of the tool (i.e., LLM) needs to be edited by an SME but can serve as a starting point. Reference and identification numbers will often have placeholders. The tool can only modify current DFARS sections. It cannot draft new sections.

Step 4: Publish Proposed Text

This step was not a part of the AI automation effort due to external requirements.

Step 5: Public Comment Review

Comments for previously published DFARS proposals are downloaded from Regulations.gov. A user can then select the proposed regulation they would like to review. Duplicate comments are removed, and similar comments (based on semantic meaning) are grouped together. The groups are then summarized by an LLM (in this case Generative Pre-trained Transformer (GPT) 4o as a representative of the available Non-classified Internet Protocol Router GPT (NIPRGPT) [6]). If there are sufficient comments, a subset of the group's comments are used for the summarization to reduce LLM token requests without drastically affecting the summary since comments have been grouped based on similarity. Finally, the tool provides an overarching summary for the entire set of comments.

Step 6: Use comments to make changes to the proposed text

For each group of comments, the tool provides edit recommendations to the DFARS section based on what the comments are suggesting. The user must choose whether to incorporate the suggestions from these comments or not.

The research team also investigated providing these suggestions to the LLM along with the full text from the proposed rule, but the resulting text was often too long and often longer than the LLM's token limits would allow. The Task 3 pipeline could instead be repurposed for generating these modifications.

3. Evaluation of Methods

NDAAs Section Extraction

The system processed NDAAAs from 2010-2025. This resulted in a normalized, year-specific repository of NDAAAs stored as structured records, along with a reproducible set of NDAA candidate sections identified through automated statutory impact analysis. These candidate sections represent provisions likely to require FAR or DFARS regulatory updates and serve as the controlled input set for downstream regulatory mapping. A corresponding quality index was also produced. The quality index was calculated for 2010–2023. Discrepancies remain, however, primarily stemming from parsing limitations with the older PDFs, which are being addressed through an ongoing refinement effort of the extraction methods. The index is obtained as:

$$Q(t) = \frac{E_{\text{ext}}(t)}{E_{\text{exp}}(t)} \quad (1)$$

- $E_{\text{ext}}(t)$ is the number of extracted items for NDAA year t ,
- $E_{\text{exp}}(t)$ is the number of expected items for NDAA year t .

For example, for Fiscal Year (FY) 2020, the candidate extraction process produced the following NDAA sections:

[2821, 1648, 1212, 1123, 890, 888, 885, 874, 873, 872, 870, 865, 864, 861, 854, 853, 852, 849, 848, 847, 826, 825, 823, 822, 818, 816, 807, 803, 372, 322, 168]

These sections constitute the automated system’s identified “golden” set for FY2020 and are passed as inputs into the regulatory matching algorithm for FAR/DFARS linkage analysis.

Table 1. Quality Index for expected NDAA sections versus extracted NDAA sections.

Fiscal Year	Extracted	Expected	Quality Index	Missed Section Numbers
2010	15	16	94%	745
2011	18	20	90%	834, 893
2012	26	28	93%	841, 834
2013	18	20	90%	802, 1651
2014	7	7	100%	
2015	15	16	94%	745
2016	25	26	96%	1214
2017	41	43	95%	873, 1212
2018	30	31	97%	1214
2019	30	31	97%	888
2020	31	31	100%	
2021	28	29	97%	832

Fiscal Year	Extracted	Expected	Quality Index	Missed Section Numbers
2022	20	20	100%	
2023	19	19	100%	

Note that ground-truth “golden” sections are derived from manual analyst identification of NDAA provisions likely to require FAR or DFARS regulatory updates. The automated extraction process is evaluated against this benchmark. “Expected” denotes the analyst-identified count, while “Extracted” reflects the number successfully identified by the system. The Quality Index represents extraction accuracy relative to the validated benchmark. “Missed Sections” lists analyst-identified sections that were not captured by the automated method.

NDAA ↔ FAR/DFARS Matching

This section defines a second set of metrics used to evaluate system performance in identifying FAR/DFARS mappings associated with NDAA sections for a given fiscal year. “Sections with Published FAR/DFARS Mappings” represents the number of NDAA sections within a given fiscal year that have validated FAR/DFARS mappings and therefore serve as ground-truth references. “Sections < 50% Mapping Coverage” identifies sections for which the system retrieved fewer than 50% of the known FAR/DFARS mappings. The associated “Quality Index (%)” reflects overall mapping coverage performance relative to validated ground-truth associations for that fiscal year.

For clarity, the metrics can be described as follows:

Let S_y denote the set of NDAA sections in fiscal year y . For each section $s \in S_y$, define:

$G(s)$: the set of validated ground-truth FAR/DFARS mappings associated with section s

$R(s)$: the set of FAR/DFARS mappings retrieved by the system for section s

Evaluation is restricted to sections with at least one validated ground-truth mapping.

SECTIONS WITH PUBLISHED FAR/DFARS MAPPINGS

The subset of sections with validated FAR/DFARS mappings is defined as:

$$S_y^G = \{s \in S_y : |G(s)| > 0\} \tag{2}$$

The total number of such sections is:

$$N_y^G = |S_y^G| \tag{3}$$

Thus, “Sections with Published FAR/DFARS Mappings” represents the number of NDAA sections in fiscal year y for which validated FAR/DFARS mappings exist and can therefore serve as ground-truth references.

SECTION-LEVEL COVERAGE

For each section, s , mapping coverage is defined as:

$$C(s) = \frac{|R(s) \cap G(s)|}{|G(s)|} \quad (4)$$

This quantity measures the proportion of validated mappings for sections that were successfully retrieved by the system.

SECTIONS WITH LESS THAN 50% MAPPING COVERAGE

The subset of sections for which system retrieval covers fewer than half of the validated mappings is defined as:

$$S_y^{<50} = \{s \in S_y^G : C(s) < 0.5\} \quad (5)$$

The number of such sections is:

$$N_y^{<50} = |S_y^{<50}| \quad (5)$$

Accordingly, “Sections < 50% Mapping Coverage” identifies the number of NDAA sections with retrieved FAR/DFARS mappings that account for less than half of the known validated mappings.

QUALITY INDEX (%)

The Quality Index reflects overall mapping coverage performance relative to validated ground-truth associations for fiscal year y . Two formulations may be used depending on the intended interpretation.

MICRO-AVERAGED QUALITY INDEX

The micro-averaged version weights sections according to the number of validated mappings they contain:

$$QI_y^{\text{micro}} = \frac{\sum_{s \in S_y^G} |R(s) \cap G(s)|}{\sum_{s \in S_y^G} |G(s)|} \times 100 \quad (7)$$

This formulation measures the percentage of all validated mappings in fiscal year y that were correctly retrieved by the system.

INTERPRETATION

The metrics may be interpreted as follows:

N_y^G : number of NDAA sections in fiscal year y with published and validated FAR/DFARS mappings

$N_y^{<50}$: number of sections for which the system retrieved fewer than 50% of the known mappings

QI_y^{micro} overall percentage of validated FAR/DFARS mappings captured by the system

Table 2. Quality Index for NDAA section to FAR/DFAR section(s) mappings.

Fiscal Year	Sections w/ Published FAR/DFARS Mappings	Sections <50% Mapping Coverage	Quality Index (%)
2010	15	2	87%
2011	20	3	85%
2012	31	7	77%
2013	24	7	71%
2014	7	1	86%
2015	18	7	61%
2016	27	6	78%
2017	50	16	68%
2018	33	11	67%
2019	28	9	68%
2020	31	5	84%
2021	29	9	69%
2022	20	6	70%
2023	19	5	74%

Evaluation of FAR/DFARS Text Changes

The evaluation of the proposed text changes was performed using traditional metrics such as Bilingual Evaluation Understudy (BLEU) [7] and Recall-Oriented Understudy for Gisting Evaluation (ROUGE) [8], as well as the more recent and robust method of LLM-as-a-Judge [9]. Finalized section changes were used as ground truth samples. The model was provided with a 1:1 mapping of NDAA section to DFARS section and then asked to create the proposed modified DFARS section. The generated section was compared using the above methods to the ground truth sample.

In the case where there was a Many:1 or Many:Many relationship between NDAA section and DFARS section we simplified it by taking each 1:1 pair and passing it through the text generator and evaluating the results against the ground truth (1:Many instances are not a problem because each NDAA:DFARS relationship can be represented with a 1:1 resulting in a DFARS changed by a single NDAA). In these instances, the text generator did not have all of the context that the ground truth had so it should not be expected to be perfect, but this was the selected method for evaluating the changes without complicating the correlation of changes to sections. Additionally, the tool currently can only act (generate) on 1:1 changes so it was best to evaluate what the tool could do.

Ten (10) true 1:1 matches (ground truth was 1:1) and ten (10) matches where particularly high differences between the original and the changed version were provided to the LLM judge. The LLM judge was asked to grade the following factors (out of five):

- Change Identification
- Content Accuracy
- Completeness
- Structural Fidelity (basically if the ordering of the points are same or not)

Table 3. LLM judged quality of system generated DFARS text.

Dimension	Average	Minimum	Maximum
Change Identification	3.98	1.0	5.0
Content Accuracy	2.92	1.0	4.0
Completeness	3.02	1.0	5.0
Structural Fidelity	2.75	1.0	4.0

4. Proposed Unified System

The objective was to create a unified system (backend and user interface) which enables DPCAP staff to move seamlessly from NDAA review to DFARS draft generation.

The developed system includes a backend and a publicly accessible interface available to the sponsor. The modules described in Section 2 (document ingestion, keyword and context identification, text summarization, clustering, and visualization) are connected in an integrated backend that drives the user interface.

The system

- Ingests NDAA's, FAR, and DFARS
- Extracts NDAA sections
- Finds matching FAR/DFARS sections for each extracted NDAA section
- Generates text changes to the FAR/DFARS and a change summary for the Federal Register
- Ingests, clusters, and summarizes public comments for all proposed FAR/DFARS changes

The backend steps are completed and then the front-end imports the necessary information for the interface. This includes:

- all sections of interest from the NDAA (a link is provided to the full NDAA)
- sections of the DFARS matching the identified NDAA sections, mentioned above (a link is provided to the full DFARS)
- generated DFARS proposed rules and Federal Register text for each NDAA-DFARS pair
- comments from selected proposed DFARS rules
- groups of comments along with summaries and proposed revisions

The front-end interface displays all the information to support the user as they work through the steps in the pipeline.

1. Review NDAA for changes
2. Find outdated text in FAR/DFARS
3. Generate new/edited text
4. Publish proposed text (not in current scope)
5. Review public comments
6. Make final changes to text

Virginia Tech provisioned a virtual machine that runs the backend code and is the web server for the front-end code. The front-end is publicly hosted behind a secure login to make it accessible to the sponsor for testing.

The system was made available to DPCAP staff for review and assessment. Additionally, a video demonstration and visual instruction manual were developed to help users navigate the tool. Screenshots of using the system to go through the DFARS change process are provided in Figure 4 through Figure 10.

Figure 4. Selecting or uploading an NDAA.

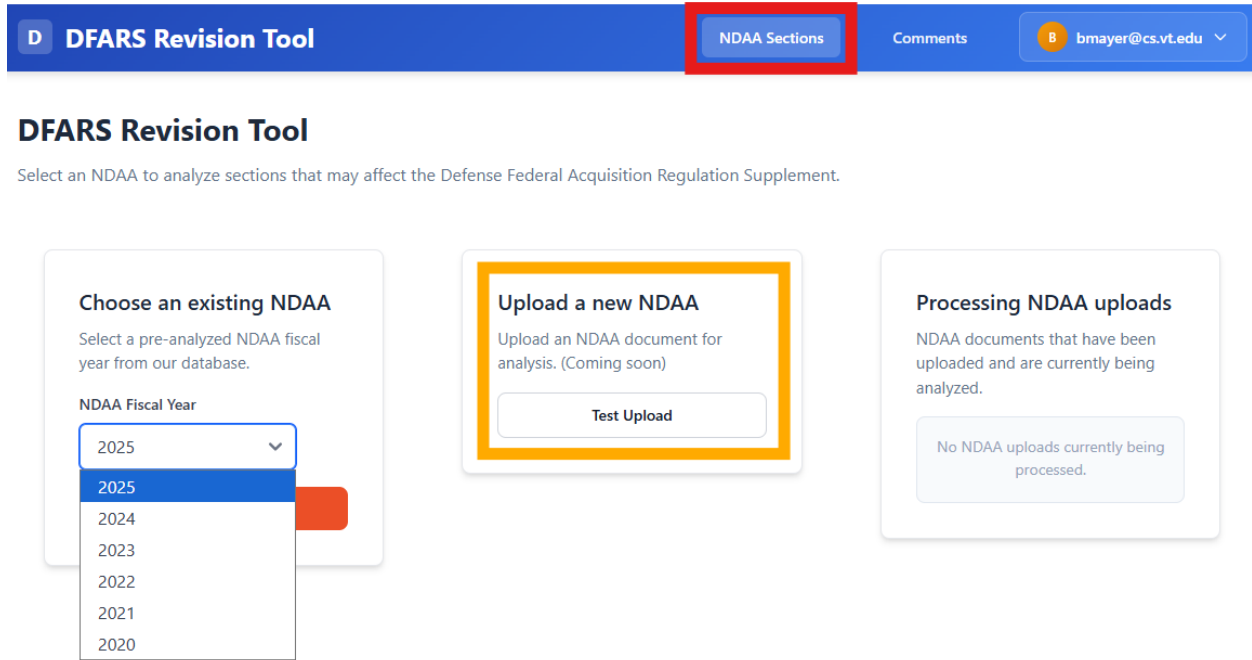


Figure 5. Reviewing the NDAA and identified sections.

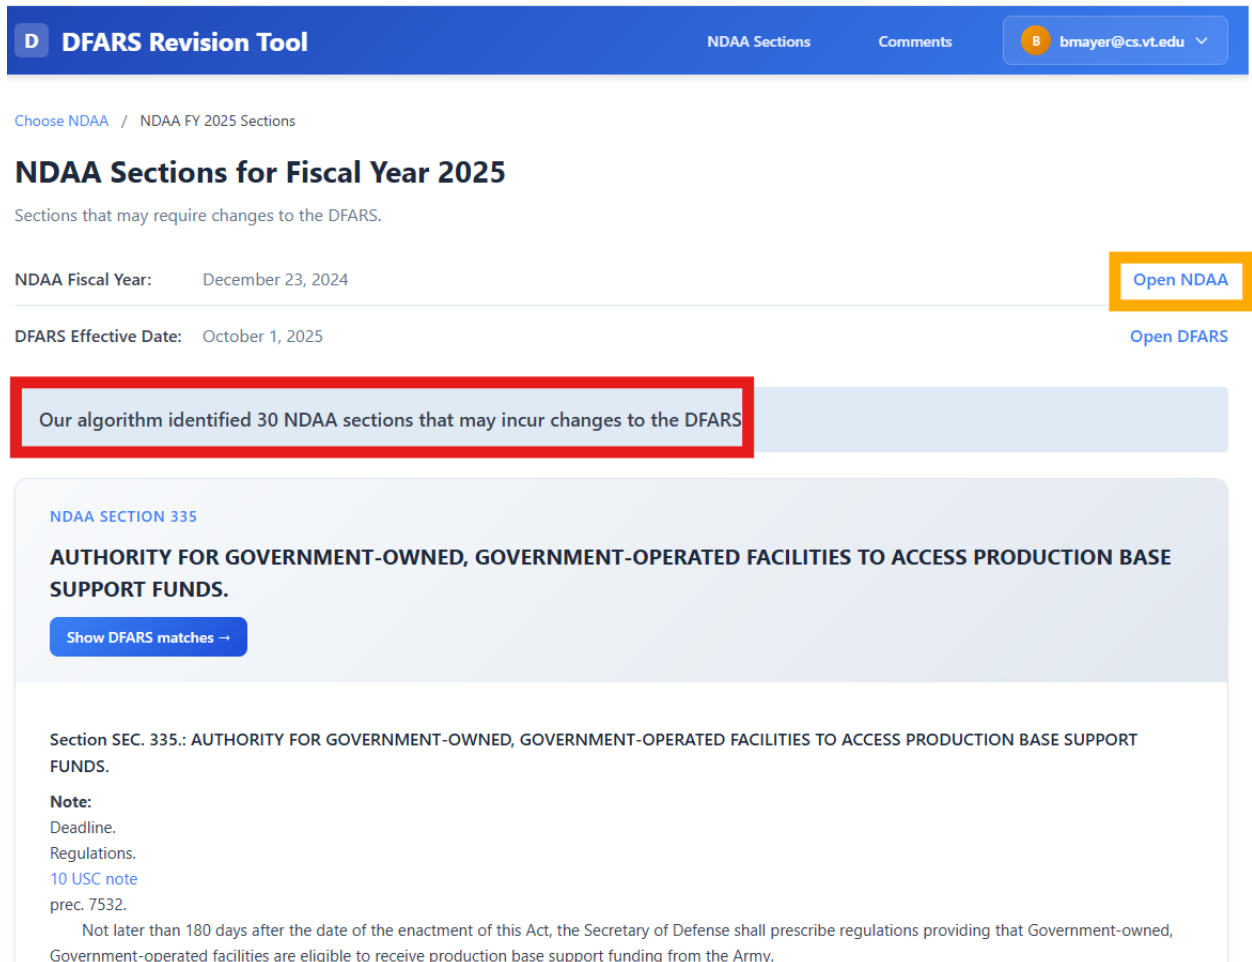


Figure 6. Selecting an NDAA section to find DFARS matches for.

The screenshot shows the 'DFARS Revision Tool' interface. At the top, there is a navigation bar with 'DFARS Revision Tool', 'NDAA Sections', 'Comments', and a user profile 'bmayer@cs.vt.edu'. Below this, the selected NDAA section is 'NDA SECTION 802' with the title 'LIMITATION ON CERTAIN OPTIONS FOR COST CONTRACTS.'. A red box highlights a button labeled 'Show DFARS matches →'. Below the button, the text of the section is displayed, including '(a) AMENDMENTS.—' and '(d) LIMITATION ON CERTAIN OPTIONS.—'.

Figure 7. Reviewing DFARS matches and selecting a match to generate a modification for.

The screenshot shows the 'DFARS Revision Tool' interface with the 'Show DFARS matches' button highlighted in red. Below the NDAA section text, there is a 'Matched DFARS Sections' section. It contains a list of matched sections: 'Part 217 - Special Contracting Methods' and 'Part 208 - Required Sources Of Supplies And Services'. A red box highlights a button labeled 'Read Section Text →'. Below the list, there are two 'Show generated changes →' buttons, one for each matched section, highlighted with a yellow box.

Figure 8. Reviewing generated text for the modification.

D **DFARS Revision Tool**

[NDAAs Sections](#)
[Comments](#)

B bmayer@cs.vt.edu

[Choose NDAAs](#) / [NDAAs FY 2025](#) / [Section 802](#) / [Suggested rule – DFARS 217.7402](#)

NDAAs Section 802 → DFARS 217.7402

Suggested Rule: 217.700 Scope of subpart.

Part 217 - Special Contracting Methods / Subpart 217.1 - Multiyear Contracting

Text of DFARS 217.7402	Proposed Change	Federal Register Summary
<p>217.7402 Limitation on Certain Options for Cost Contracts</p> <p>Scope</p> <p>This clause applies to all cost reimbursement contracts for the development of a major system, as defined in 10 U.S.C. 3041, awarded by the Department of Defense. It specifically governs contracts involving fixed-priced options for production quantities and limits the number of low-rate initial production (LRIP) lots to no more than one. This limitation applies regardless of contract value or contractor classification. Waivers to this limitation may be granted only by the service acquisition executive or, in the case of joint programs, the Secretary of Defense.</p> <p>Definitions</p> <p>For purposes of this clause:</p> <ul style="list-style-type: none"> (1) Covered contract means a cost reimbursement contract for the development of a major system. (2) Low-rate initial production (LRIP) has the meaning given in 10 U.S.C. 4231. (3) Major system has the meaning provided in 10 U.S.C. 3041. <p>Restrictions</p> <ul style="list-style-type: none"> (1) Except as provided in paragraph (2), the contract shall limit the number of low-rate initial production lots procured using fixed-priced options to not more than one. (2) A waiver to the limitation in paragraph (1) may be granted only under the following conditions: <ul style="list-style-type: none"> (i) The waiver must be formally determined by the service acquisition executive of the military department concerned or, for joint programs, by the 		

Figure 9. Selecting a DFARS proposal and reviewing the summary of public comments.

The screenshot displays the 'DFARS Revision Tool' interface. At the top, there are navigation tabs for 'NDAAs Sections' and 'Comments', and a user profile for 'bmayer@cs.vt.edu'. The main heading is 'Comment Summaries for Proposed Rules' with a count of 49. Below this is a sub-heading 'Analysis of public comment on proposed rules.'.

On the left side, there is a 'Search Proposals' section. A search input field contains 'Pay Equity', and a dropdown menu shows search results. The first result, 'FAR-2023-0021-0001 - Acquisition Regulation: Pay Equity and Transparency in Federal Contracting', is highlighted with a red box. Below this, other proposals are listed, including 'FAR-2025-0006-0001' and 'FAR-2024-0007-0001'. The selected proposal, 'FAR-2023-0021-0001', is highlighted with a yellow box.

The main content area displays the details for 'FAR-2023-0021-0001'. The title is 'Acquisition Regulation: Pay Equity and Transparency in Federal Contracting'. Below the title, it states 'Comment Summary (based on 25 comments)'. The summary text reads: 'Concerns have been raised about legal compliance and the definition of "performance in the United States" in the proposal, alongside strong support for FAR Case 2023-021, which aims to enhance pay equity and transparency in federal contracting. The proposed rule is viewed as a crucial measure to prevent pay disparities and promote fairness in the workplace.' A link 'View proposal at regulations.gov' is provided.

Below the summary, there is a 'Comment Groups' section with a list of expandable items: 'Legal Concerns Over 41 U.S.C 6707(c)', 'Clarification on USAID Contract Performance', 'File Attachment Discussion Highlights', and 'Support for FAR Case 2023-021: Pay Equity'.

Figure 10. Reviewing the summary of a comment group, the proposed revision based on the comments, and viewing the raw text from all the grouped comments.

The screenshot displays the 'DFARS Revision Tool' interface. On the left is a sidebar titled 'Search Proposals' with a search input field and a list of proposal entries. The main content area shows the details for 'Acquisition Regulation: Pay Equity and Transparency in Federal Contracting', including a comment summary and a 'Comment Groups' section. The 'Legal Concerns Over 41 U.S.C 6707(c)' group is expanded to show a 'Summary' and a 'Revision Suggestion'.

DFARS Revision Tool | NDAA Sections | Comments | bmayer@cs.vt.edu

Search Proposals

Search by title or section number...

- FAR-2025-0006-0039
Federal Acquisition Regulation: Ending Procurement and Forced Use of Paper Straws
- FAR-2025-0006-0001
Federal Acquisition Regulation: Ending Procurement and Forced Use of Paper Straws
- FAR-2024-0007-0001
Acquisition Regulation: Protests of Orders under Certain Multiple-Award Contracts
- FAR-2024-0001-0001
Acquisition Regulation: Inflation Adjustment of Acquisition-Related Thresholds
- FAR-2023-0021-0001**
Acquisition Regulation: Pay Equity and Transparency in Federal Contracting
- FAR-2023-0013-0001
Acquisition Regulation: HUBZone Program

Acquisition Regulation: Pay Equity and Transparency in Federal Contracting

Comment Summary (based on 25 comments)

Concerns have been raised about legal compliance and the definition of "performance in the United States" in the proposal, alongside strong support for FAR Case 2023-021, which aims to enhance pay equity and transparency in federal contracting. The proposed rule is viewed as a crucial measure to prevent pay disparities and promote fairness in the workplace.

[View proposal at regulations.gov](#)

Comment Groups

Legal Concerns Over 41 U.S.C 6707(c)

Summary

Concerns have been raised regarding potential violations of 41 U.S.C 6707(c) in the proposal. Commenters are questioning the legality and compliance of the government's actions.

Revision Suggestion

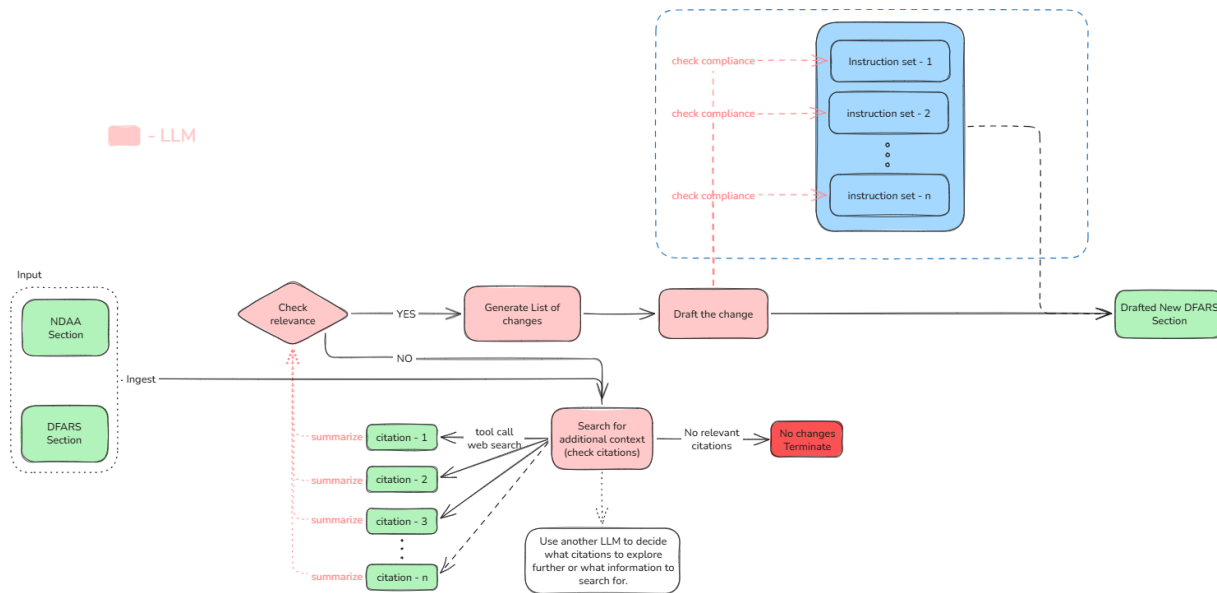
Revise the proposed rule to include a clear compliance section that explicitly addresses adherence to 41 U.S.C. 6707(c). Include specific language that outlines the measures taken by the government to ensure legal compliance throughout the acquisition process. Additionally, provide examples of actions that will be implemented to mitigate any potential violations. Ensure that the compliance section is easily accessible within the document for reference by all stakeholders.

[View All Comments](#)

5. Improvements to Generation of FAR/DFARS Text Changes

The text generation process was improved by using an LLM (GPT 4.1) agent to evaluate the need for changes then another to collect additional context which is fed into a third LLM agent to determine the necessary changes given the input sections and the identified context. Finally, a fourth LLM agent is provided with FAR drafting guidelines and the necessary changes to generate the modified section text. The process is depicted in Figure 11.

Figure 11. Process for generating changes to a DFARS section given an old DFARS section and an NDAA section



The first agent is provided the NDAA section and DFARS section. It is tasked with identifying references to other documents (e.g., US Code, NDAA, FAR, DFARS) and building up a knowledge base to support the other steps. Once this context is identified the second agent is provided with the NDAA section the DFARS section and all the identified context documents. The second agent is asked to determine if a change is necessary to the DFARS given the NDAA section and all the context.

If no change is necessary, the user is informed that no text was generated because the system could not identify an appropriate change. If a change is appropriate, the third agent is tasked with describing the necessary changes. This description is then provided to the fourth and final agent to generate the modified text.

There are a few significant limitations to this process. First, the process currently only supports modifications to sections. To account for cases where a completely new DFARS section is required the pipeline will have to be fed the entire DFARS subpart. This would need to be determined by the second agent, and provided to the third agent.

Additionally, while the first agent does identify the related citations and provide them to the other agents this tool is using generative AI and it cannot be relied upon to make accurate citations [10]. The system may generate text with citations but there is no guarantee (and there should be no expectation) that these citations will be accurate

Finally, as previously stated, the system currently only handles a 1:1 mapping of NDAA sections to DFARS sections. In many cases, multiple NDAA sections may change 1 or more DFARS sections. This would require identifying all the NDAA sections at once which is not currently a capability of the NDAA-DFARS matching process. User intervention is necessary to review proposed text and therefore they can incorporate changes from multiple NDAA sections at once using the text changes from each 1:1 mapping automatically generated from the system.

Despite some limitations, the model still provides a reasonable draft (as seen in the results in Section 6), but not without requiring SME review and potential edits.

Additionally, the multi-agent process forces the agents to perform tasks with more reasoning, which consequently provides additional explainability to the user. This explainability could be further leveraged by the rest of the system and interface for further improvements and adoption.

6. Conclusions, Recommendations and Next Steps

Conclusions

This research demonstrates the feasibility and value of integrating LLMs, Natural Language Processing (NLP), and Machine-Based Reasoning (MBR) into a unified framework to support the DFARS rule development process. By further evaluating, refining, and consolidating previously independent AI-driven capabilities into a single, end-to-end system, the project enables a more streamlined transition from NDAA analysis to DFARS rule drafting and revision.

The evaluation results indicate that automated NDAA section extraction and FAR/DFARS mapping achieve consistently high levels of accuracy, providing a reliable foundation for downstream tasks. While variability remains in mapping coverage and text generation quality, particularly in complex or many-to-many relationships, the system has the potential for meaningful reduction in manual workload and acceleration of regulatory analysis. The incorporation of multi-agent LLM pipelines further enhances contextual understanding and improves the quality of generated regulatory text, though human SME oversight remains essential to ensure correctness, compliance, and trustworthiness. Additionally, the integration of public comment analysis introduces a scalable mechanism for incorporating stakeholder feedback into rule development, further strengthening the policy lifecycle.

Despite these advances, the tool does have limitations. Current capabilities are constrained to one-to-one mappings between NDAA and DFARS sections, and the system does not yet fully support the creation of entirely new regulatory sections. Furthermore, challenges related to LLM reliability, including citation accuracy and potential hallucinations, necessitate continued emphasis on human-in-the-loop validation and robust governance frameworks. Additional fine-tuning is also imperative given the limited amount of ground truth data and risk of over-fitting.

Suggested Future Research and Next Steps

Future work should address the mentioned limitations (e.g., many:1 mappings, LLM reliability, model testing and fine-tuning).

Future work should also focus on leveraging the LLM-as-a-judge (first agent from the text generation pipeline) to act as an additional tool in the NDAA identification and DFARS matching processes. In this scenario the matching models could be tuned to increase recall, and the LLM-as-a-judge could be used to increase precision. This could result in higher recall and higher precision of the joint system. The LLM agents' output could also be further leveraged to increase traceability between regulatory documents and statutory mandates, providing further explainability to the user.

Final Thoughts

Overall, this work highlights the transformative potential of AI-assisted systems in regulatory environments, offering a scalable, semi-automated approach that augments human expertise while maintaining the rigor required for defense acquisition policy development.

References

- [1] Under Secretary of Defense for Acquisition and Sustainment. (2011, April 30) *FAR Drafting Guide*. Department of Defense.
https://www.acq.osd.mil/dpap/dars/docs/far_dfars_guide/FAR%20Drafting%20Guide--April%2030,%202011.pdf
- [2] Alansari, A. and Luqman, H. (2026). Large Language Models Hallucination: A Comprehensive Survey, [arXiv:2510.06265v3](https://arxiv.org/abs/2510.06265v3)
- [3] Anh-Hoang D, Tran V, Nguyen LM. (2025). "Survey and analysis of hallucinations in large language models: attribution to prompting strategies or model behavior." *Front Artif Intell*. doi: 10.3389/frai.2025.1622292. PMID: 41098969; PMCID: PMC12518350.
- [5] J. Wei et al., "Chain-of-Thought Prompting Elicits Reasoning in Large Language Models," Jan. 10, 2023, *arXiv*: arXiv:2201.11903. doi: 10.48550/arXiv.2201.11903.
- [4] Tauman Kalai, A., Nachum, O., Vempala, S.S. and Zhang, E. (2025). Why Language Models Hallucinate, [arXiv:2509.04664v1](https://arxiv.org/abs/2509.04664v1)
- [5] J. Wei et al., "Chain-of-Thought Prompting Elicits Reasoning in Large Language Models," Jan. 10, 2023, *arXiv*: arXiv:2201.11903. doi: 10.48550/arXiv.2201.11903.
- [6] Secretary of the Air Force Public Affairs. (2024 Jun. 10) *Department of the Air Force launches NIPRGPT*. <https://www.af.mil/News/Article-Display/Article/3800809/department-of-the-air-force-launches-niprgpt/>
- [7] Papineni, K., Roukos, S., Ward, T., & Zhu, W. J. (2002, July). Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics* (pp. 311-318).
- [8] Lin, C. Y. (2004, July). Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out* (pp. 74-81).
- [9] Gu, J., Jiang, X., Shi, Z., Tan, H., Zhai, X., Xu, C., ... & Guo, J. (2024). A survey on llm-as-a-judge. *The Innovation*.
- [10] Buchmann, J., & Gurevych, I. (2025). Citation Failure: Definition, Analysis and Efficient Mitigation. *arXiv preprint arXiv:2510.20303*.