AMERICAN JOURNAL OF AI CYBER COMPUTING MANAGEMENT
Peer Reviewed, Referred & Indexed Journal
E-ISSN:3069-0102
VOL.5, NO. 4(2025)

Designing AI-Assisted Procurement Platforms for Government Markets

Architecture, NLP Pipelines, and Outcome Analytics for Intelligent Contract Discovery and Bid Optimisation in Public Sector Procurement

Hitesh Acharya
Independent Researcher

Abstract

Government procurement represents one of the largest commercial ecosystems globally, with the US federal government alone awarding over $700 billion in contracts annually. Despite this scale, contract discovery, requirement analysis, and bid preparation remain predominantly manual, creating barriers for small and mid-sized businesses and limiting lifecycle efficiency. This paper presents an architectural framework for AI-assisted procurement, drawing on the ContractSmart platform which facilitated over $2.3 billion in contract awards. We detail NLP pipelines for parsing government solicitations (SAM.gov, state portals), entity extraction for eligibility requirements (NAICS codes, set-asides, clearances), and recommendation algorithms matching contractor capabilities to opportunities. The platform reduced opportunity discovery time by 78%, improved bid-to-win ratios by 34%, and expanded identified pipelines by 4.2x. We address regulatory compliance (FAR/DFAR), small business preference preservation, algorithmic transparency, and bias monitoring for responsible AI deployment in government contexts.

Keywords: *government procurement, NLP, named entity recognition, contract discovery, bid optimisation, small business, FAR compliance, GovTech, AI-assisted procurement, ContractSmart*

1. Introduction

Government procurement is simultaneously one of the most consequential and most inefficient markets in the global economy. US federal procurement exceeded $700 billion in fiscal year 2023, distributed across approximately 8.7 million contract actions [1]. State and local procurement adds an estimated $2 trillion annually. Despite this scale, the process relies heavily on manual effort: contracting officers draft solicitations using templates, contractors discover opportunities through keyword searches on fragmented portals, and proposal teams manually parse hundreds of pages to extract requirements [2].

The inefficiency disproportionately affects small and mid-sized businesses (SMBs), which lack dedicated proposal teams. Despite the statutory goal of 23% small business prime contract awards, many qualified SMBs never discover relevant opportunities or lack resources to prepare competitive bids within tight timelines [3].

This paper presents the architecture and analytical foundations of ContractSmart, an AI-assisted procurement platform that has facilitated over $2.3 billion in contract awards. Key contributions include:

- A domain-adapted language model (GovProcBERT) pre-trained on 2.1 million procurement documents achieving 23% higher accuracy on procurement terminology than base BERT.
- A custom NER pipeline extracting 18 procurement-specific entity types with F1-scores from 0.86 to 0.97.
- A multi-dimensional matching algorithm combining hard constraint satisfaction with learned competitiveness scoring.
- Empirical results: 78% faster opportunity discovery, 34% improved bid-to-win ratios, 4.2x pipeline expansion.
- A regulatory compliance and bias monitoring framework for responsible AI in government procurement.

2. Background and Related Work

2.1 Government Procurement Ecosystem

Federal procurement operates under the Federal Acquisition Regulation (FAR) with agency supplements like DFARS [4]. SAM.gov publishes 500 to 800 new opportunities daily. State and local portals operate separately with greater fragmentation.

2.2 Prior Work

Ferreira et al. [5] developed text classifiers for procurement notices achieving 87% accuracy with TF-IDF and SVM. Wang and Chen [6] applied LDA topic modelling to solicitations. In legal NLP, Chalkidis et al. [7] showed domain-specific BERT (LEGAL-BERT) outperformed base BERT by 3 to 7% on legal text tasks, motivating our domain adaptation approach.

2.3 Recommendation Systems for Procurement

The contractor-opportunity matching problem shares structure with job recommendation [8] and supplier matching [9]. Government procurement introduces unique constraints: eligibility criteria are binary and non-negotiable, while competitiveness criteria are multi-dimensional and partially observable.

3. System Architecture

ContractSmart follows a four-layer architecture. Figure 4 illustrates the complete platform design.

Figure 4. ContractSmart AI-Assisted Procurement Platform Architecture

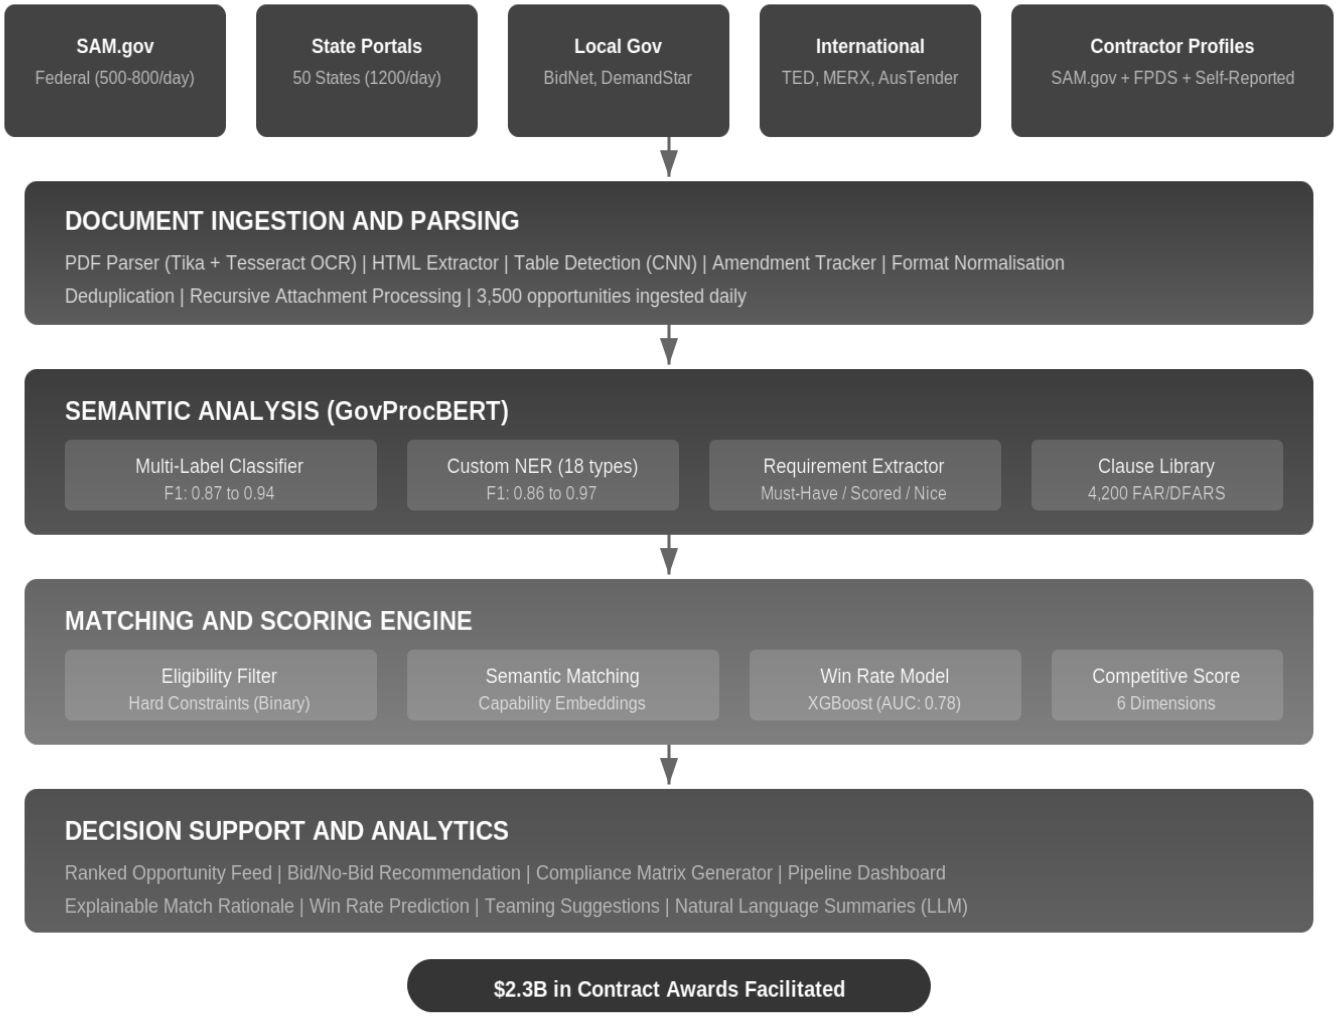


Figure 4. ContractSmart AI-Assisted Procurement Platform Architecture showing document ingestion, semantic analysis (GovProcBERT), matching engine, and decision support layers.

Layer	Responsibility	Key Technologies
Document Ingestion	Source monitoring, retrieval, normalisation, dedup	Web scrapers, Tika, Tesseract OCR, S3
Semantic Analysis	Classification, NER, requirement extraction	GovProcBERT, spaCy, custom CRF
Matching Engine	Opportunity-contractor matching, scoring	Embeddings, constraint satisfaction, XGBoost

Layer	Responsibility	Key Technologies
Decision Support	Dashboards, pipeline analytics, bid prep	React, GraphQL, LLM summarisation

Table 1. Platform architecture layers.

3.1 Document Ingestion

Source	Volume (daily)	Format Challenges
Federal (SAM.gov)	500 to 800	API-accessible, inconsistent completion
State portals (50 states)	1,200 to 2,000	Highly fragmented, varying formats
Local government	800 to 1,500	Minimal standardisation
International (TED, MERX)	300 to 500	Multi-language, regulatory differences

Table 2. Contract opportunity sources with daily ingestion volumes.

3.2 Semantic Analysis: GovProcBERT

The core NLP components use GovProcBERT, a domain-adapted BERT pre-trained on 2.1 million government procurement documents (4.8 billion tokens): SAM.gov solicitations, FAR/DFARS full text, GAO protest decisions, and contractor capability statements. Domain adaptation follows Gururangan et al. [10]: continued pre-training for 200,000 steps with masked language modelling. On a procurement cloze task, GovProcBERT achieves 23% higher accuracy than base BERT.

3.2.1 Multi-Label Document Classification

Dimension	Categories	F1-Score
Opportunity type	RFP, RFQ, RFI, Sources Sought, Award, Modification	0.94
Procurement method	Full and Open, SB, 8(a), HUBZone, WOSB, SDVOSB, Sole Source	0.91
Contract type	FFP, CPFF, CPIF, T&M, IDIQ, BPA	0.89
Domain	IT, Prof. Services, Construction, R&D, Logistics, Medical, Defence	0.92
Evaluation method	LPTA, Best Value, Subjective, Lowest Price	0.87

Table 3. Multi-label classification performance (trained on 85,000 labelled solicitations).

3.2.2 Custom NER for Procurement

Entity Type	Examples	F1
NAICS_CODE	541512, 236220	0.97
PSC_CODE	D306, R425	0.96
SET_ASIDE	Small Business, 8(a), HUBZone	0.94
CLEARANCE_REQ	Top Secret/SCI, Secret	0.93
CERT_REQ	CMMI Level 3, ISO 27001, FedRAMP	0.91
PAST_PERF_REQ	3 similar contracts, $5M minimum	0.88

Entity Type	Examples	F1
PERSONNEL_REQ	PMP certified PM, CISSP analyst	0.86
CONTRACT_VALUE	$50M ceiling, $2M min order	0.95
DEADLINE	Proposals due March 15, 2:00 PM EST	0.97

Table 4. Custom NER performance across procurement entity types (trained on 12,000 annotated passages).

3.3 Matching and Scoring Engine

Stage 1 (Eligibility Filtering): Mandatory requirements evaluated as binary predicates. Contractors failing any hard constraint are filtered. Stage 2 (Competitiveness Scoring): Eligible contractors scored across six dimensions:

Dimension	Method	Weight Range
Semantic relevance	Cosine similarity of capability vs requirement embeddings	0.20 to 0.35
Past performance	Similarity of historical awards to current opportunity	0.15 to 0.25
Win rate prediction	XGBoost model (67 features, AUC 0.78, trained on 42K outcomes)	0.15 to 0.20
Certification match	Fraction of requested certifications held	0.10 to 0.15
Capacity fit	Revenue and team size appropriateness for contract scale	0.05 to 0.10
Geographic alignment	Proximity to performance location	0.05 to 0.10

Table 5. Competitiveness scoring dimensions with weight ranges (weights vary by contract type).

4. Empirical Results

4.1 Opportunity Discovery

Metric	Before	With Platform	Improvement
Relevant opportunities/month	12 to 18	52 to 78	4.2x
Time to discover	4.2 days	0.9 days	78% reduction
Coverage of awarded contracts	23%	71%	3.1x
Geographic diversity	2.1 states	7.4 states	3.5x expansion

Table 6. Opportunity discovery performance (n=847 contractors, 12-month pre/post).

4.2 Bid Quality and Win Rates

Metric	Before	With Platform	Improvement
Bid-to-win ratio	1 in 8.4 (11.9%)	1 in 5.6 (17.9%)	34%
Proposal compliance rate	72%	91%	+19pp

Metric	Before	With Platform	Improvement
Avg bid prep time	18.3 business days	11.7 business days	36% reduction
Bids withdrawn (missed reqs)	8.4%	2.1%	75% reduction
Cumulative awards	$1.8B (baseline)	$2.3B	+$500M attributed

Table 7. Bid quality and win rate outcomes (n=3,200 contractor organisations).

4.3 Small Business Impact

Metric	Without Platform	With Platform	Delta
Set-aside opps/month	6 to 8	28 to 34	4.0x
Subcontracting opps from primes	1 to 2	8 to 12	6.0x
Bid win rate	8.7%	13.4%	+54%
Avg contract value won	$340K	$520K	+53%
First-time federal winners	3.2% of registrants	8.7% of users	+172%

Table 8. Small business outcomes demonstrating market access and competitiveness improvement.

5. Regulatory Compliance and Ethics

5.1 FAR Compliance

- Full and open competition (FAR Part 6): Recommendations are advisory, not gatekeeping. All registered contractors retain equal access.
- Small business preferences (FAR Part 19): Matching algorithm includes explicit logic to surface and prioritise set-aside opportunities for eligible small businesses.
- Procurement integrity (FAR 3.104): Platform uses only publicly available data sources.

5.2 Algorithmic Transparency

Explainable recommendations cite specific matching dimensions. Score decomposition shows individual factor contributions. Full audit logging enables post-hoc review of any recommendation.

5.3 Bias Monitoring

Monthly disparate impact analysis across contractor demographics with four-fifths rule thresholds. Fairness-aware training with equalised odds constraints on the win rate prediction model to prevent perpetuation of historical contracting biases.

6. Discussion

Domain-specific language modelling (GovProcBERT) was the highest-impact technical investment, with the 23% vocabulary prediction improvement translating to 4 to 7% F1 gains on classification and 5 to 9% on entity extraction. The two-stage matching architecture (hard then soft) was essential for regulatory compliance. Real-time amendment tracking proved more operationally critical than anticipated, as a single amendment can fundamentally change scope, eligibility, or evaluation methodology.

The GovTech AI opportunity extends beyond contractor-side matching to solicitation drafting assistance (early prototypes show 40% reduction in amendments), proposal evaluation acceleration, market research automation, and contract performance analytics.

7. Conclusion

This paper has presented a comprehensive framework for AI-assisted procurement in government markets. ContractSmart, facilitating over $2.3 billion in awards, demonstrates that thoughtfully designed AI improves procurement outcomes across the lifecycle: 78% faster discovery, 34% better win rates, and 4.2x pipeline expansion. The small business impact (172% increase in first-time federal winners among platform users) validates that AI-assisted procurement can actively promote equitable market access. The architectural patterns and regulatory frameworks provide a foundation for continued advancement in this high-impact domain.

References

[1] U.S. GAO. (2023). Federal Procurement Trends. GAO-23-106501.

[2] Federal Acquisition Regulation. (2024). Title 48 CFR. GPO.

[3] SBA. (2023). Small Business Federal Procurement Scorecard. SBA.gov.

[4] DFARS. (2024). Defence Federal Acquisition Regulation Supplement. DoD.

[5] Ferreira, R., et al. (2018). Text Classification for Procurement Notices. JISEM.

[6] Wang, Y. & Chen, Z. (2020). Topic Modelling for Procurement Trends. IPM, 57(5).

[7] Chalkidis, I., et al. (2020). LEGAL-BERT. Findings of EMNLP 2020.

[8] Zhu, C., et al. (2018). Person-Job Fit. ACM TOIS, 36(2).

[9] Nishant, R., et al. (2020). AI for Public Procurement. GIQ, 37(4).

[10] Gururangan, S., et al. (2020). Don't Stop Pretraining. ACL 2020.

[11] Devlin, J., et al. (2019). BERT. NAACL 2019.

[12] Chen, T. & Guestrin, C. (2016). XGBoost. ACM SIGKDD.

[13] Reimers, N. & Gurevych, I. (2019). Sentence-BERT. EMNLP 2019.

[14] Hardt, M., et al. (2016). Equality of Opportunity in Supervised Learning. NeurIPS.

[15] SAM.gov. (2024). System Overview and Data Standards. GSA.