pAtChWoRK
Patching the Pieces of Public Procurement Documents

Lorena Calvo-Bartolomé[C] **Saúl Blanco-Fortes**[C]
Erick Abdiel Cedeño Guerra[P] **Jerónimo Arenas-García**[C]
[C]Universidad Carlos III de Madrid, Spain [P]Universidad Politécnica de Madrid, Spain

Abstract

Public procurement data is legally open, yet practically locked inside thousands of unstructured PDFs and inconsistent portal metadata. pAtChWoRK starts with these fragmented, unstructured sources then leverages a hybrid pipeline (traditional NLP with LLM-based technologies) to restructure this information into a navigable knowledge base. Specifically, pAtChWoRK corrects manual classification errors, extracts complex unstructured fields such as award and solvency criteria and tenders' objectives, and assists users in easily navigating the tender landscape. This unified process enables more effective handling of the transparency bottlenecks that hinder competition and oversight in public administration. A user study with practitioners across different procurement roles demonstrates the system's effectiveness in supporting daily procurement tasks.

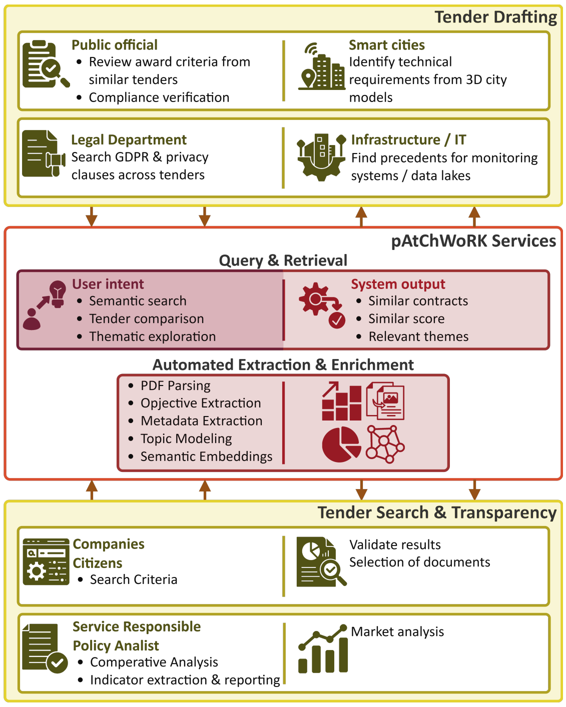

Figure 1: pAtChWoRK serves heterogeneous user profiles. Extraction of solvency and award criteria from similar tenders can aid the preparation of new tenders (top). Search utilities help companies and individuals identify relevant tenders, while aggregated indicators support comparative analysis and transparency.

1 Introduction

Imagine you are a procurement officer drafting a contract for 3D city visualizations. You need immediate access to technical requirements defined in analogous projects across different administrations, yet the only tools at your disposal are basic metadata-based filters and thousands of unstructured PDFs. This is far from an isolated case of practitioners being unable to effectively access and reuse public procurement information; *the root cause is structural:* **public procurement data is technically open but practically inaccessible.** In Spain, as in numerous jurisdictions (European Commission, 2023; OECD Publishing, 2025), individual administrations maintain their own procurement portals, but all are legally required to aggregate data into a central national platform. In practice, however, this aggregation comes at a significant cost to data quality. *First,* aggregated records are systematically less complete, as portals are only required to publish a subset of metadata fields.[1] *Second,* missing data is widespread even in those fields that are mandatory. *Third,* while PDFs contain the richest procurement information, they are authored manually by contracting officers and are therefore prone to errors and inconsistencies (Soylu et al., 2022b). *Fourth,* and perhaps most critically, contract objectives provided as metadata (the field for understanding what is actually being procured) are typically reduced to a bare title, offering little granular insight beyond what a Common Procure-

[1]For example, the technical specifications for Spanish public procurement datasets show that different metadata fields are mandatory across the three main datasets.

Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations), pages 556–565
July 2-7, 2026 ©2026 Association for Computational Linguistics

ment Vocabulary (CPV) code already provides, as the statistics of the Spanish procurement data in Table 2 before applying pAtChWoRK reveal.

These shortcomings have real consequences for all actors involved. Companies face serious difficulties identifying relevant tenders, which reduces market competition and overall concurrence. Public administrations waste resources drafting contracts for which precedents already exist, often without the technical expertise to do so effectively, and struggle to identify suitable companies for their needs. Society bears the cost of inefficient public spending: the lack of transparency makes it difficult to monitor procurement expenditure, while fragmented, heterogeneous, and unreliable procurement metadata hinder benchmarking practices and the systematic comparison of spending patterns across administrations (Fazekas and Blum, 2021).

The tools to address this gap exist in pieces (§5): CPV prediction (Siciliani et al., 2024b; Navas-Loro et al., 2022), topic modeling (Rejeb et al., 2023; Hott et al., 2023) to identify thematic clusters and trends, anomaly detection (Soylu et al., 2022b), but no existing system stitches them into a coherent end-to-end solution spanning the full procurement lifecycle (Siciliani et al., 2024a). We present pAtChWoRK, a pipeline feeding a web tool that jointly leverages traditional NLP techniques and LLMs to **(i)** extract and enrich structured metadata from raw procurement PDFs, reducing missing data and correcting classification errors; **(ii)** enable semantic and thematic search over the enriched corpora; and **(iii)** provide targeted services for companies, public administrations, and civil society, including a tender recommendation system and a procurement quality indicators dashboard (see Fig. 2). While the data processed here covers Spanish procurement (PLACE[2]), the proposed tools are readily adaptable to other jurisdictions and frameworks. In sum, our contributions are:

1. We develop pAtChWoRK, a system integrating all the above into a user-friendly interface with dedicated panels (§2) for: uploading procurement PDFs and automatically extracting key metadata; exploring pre-enriched documents with structured metadata; discovering topics and semantic links across tenders; performing semantic and thematic similarity search; and visualizing indicators for transparency and performance assessment. The source code is available under the MIT license[3] and the web interface is publicly accessible[4] upon login.
2. We conduct a user study with domain experts spanning different procurement roles (§4), demonstrating pAtChWoRK's viability and shedding light on future work directions.

2 pAtChWoRK in Action

A video demonstration of pAtChWoRK is available online.[5] Here we describe the three main functional layers of the system, illustrated in Fig. 2.

Text Extraction and Metadata Augmentation Users can upload raw procurement PDFs (either administrative, PCAP, or technical, PCPT) for automated processing. The system extracts and enriches key metadata fields: contract objectives, CPV codes, award criteria, solvency requirements, and special conditions. Processed documents are immediately available for exploration (Fig. 2a).

Keyword, Thematic, and Semantic Tender Exploration & Retrieval. The corpus search tool supports three complementary retrieval modes over the enriched repository. Returning to our opening example, a procurement officer looking for precedents for a 3D city visualization contract can query the system using a syntactic query (e.g., *"3D"*), or a natural language description of their needs (e.g., *"generation of 3D visualizations of the city"*) for thematic- or semantic-based retrieval. In **thematic retrieval**, the user selects a collection, a CPV code, and a granularity level. The system builds two topic models per CPV, providing a hierarchical thematic representation of the corpus. Results include a topic distribution chart situating the query within the thematic landscape, alongside ranked documents with their predicted CPVs and generated objectives (Fig. 2c). Users can drill down into any cluster to inspect its keywords, statistics, and representative documents (Fig. 2b). In **semantic retrieval**, results follow the same ranked format but without the topic distribution chart, as similarity is computed directly over dense embeddings. Users can further refine results via syntactic filters.

Analytical Oversight. The indicators dashboard (Fig. 2d) currently visualizes four procurement metrics, filterable by year and quarter, and are broken

[2] https://contrataciondelestado.es/wps/portal/plataforma

[3] https://github.com/beauseant/patchwork
[4] https://patchwork.uc3m.es/
[5] https://youtu.be/IyWAeOWHpvE

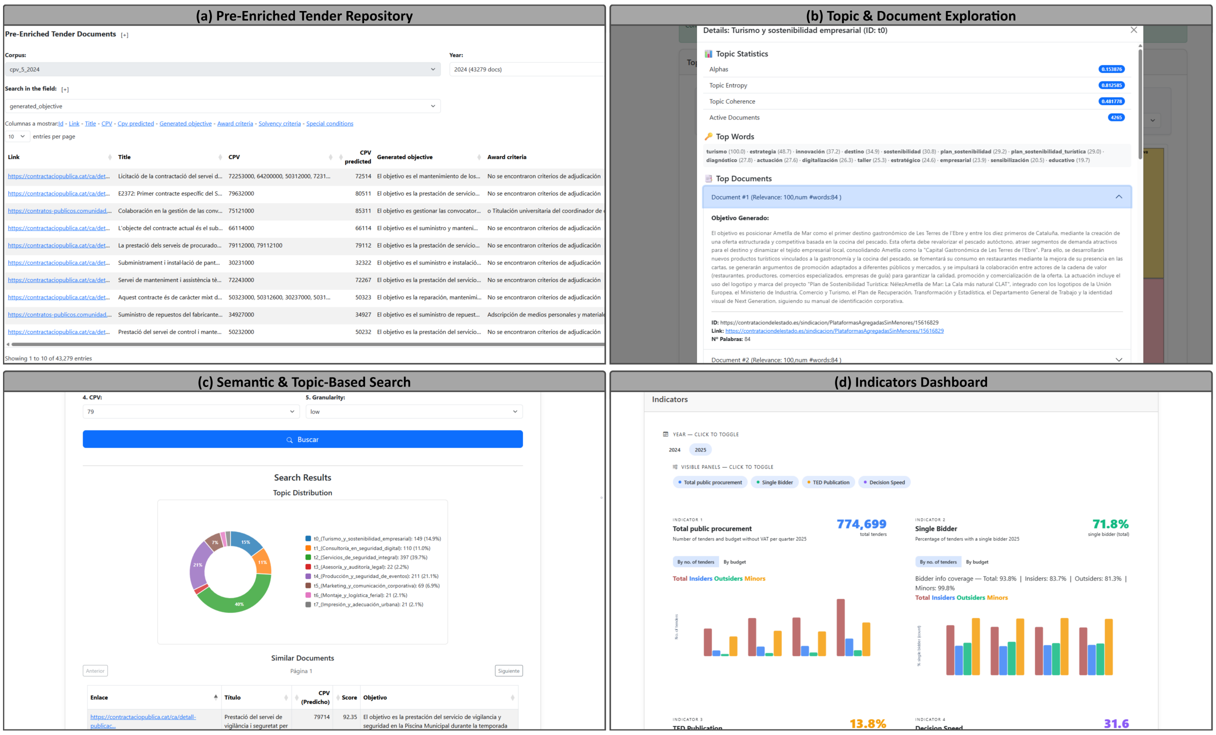

Figure 2: pAtChWoRK's main functional panels. (a) Pre-enriched tender repository with LLM-generated contract objectives, award criteria, and solvency requirements; (b) topic exploration panel showing statistics, top words, and representative documents for a given thematic cluster; (c) corpus search tool with topic distribution chart and ranking of similar documents; (d) Procurement indicators dashboard for transparency and performance monitoring.

down by three procurement datasets: *insiders*, *outsiders* and *minors*.[6] The selected indicators are taken from those defined at the Public Procurement Data Space (PPDS) from the European Commission,[7] but are calculated for the whole Spanish public procurement, whereas PPDS refer only to tenders that have been published in Tenders Electronic Daily (TED), the European procurement portal.

The **Total Public Procurement** indicator tracks aggregate procurement activity, computed separately by number of tenders and by budget volume. Also under these two dimensions, the **Single Bidder** indicator shows the proportion of tenders awarded to a single bidder, serving as a proxy for market competition. The **TED Publication** indicator measures the percentage of tenders and associated budget that have been published in TED, reflecting openness and cross-border accessibility. Finally, the **Decision Speed** indicator captures the average number of days between the bid deadline and the award decision, offering a measure of administrative efficiency in the procurement process.

3 System architecture

pAtChWoRK follows a three-tier architecture (Fig. 5 in §C). The backend hosts data processing pipelines and NLP modules exposed via a REST API (Fig. 3), with the frontend and each backend component deployed as independent Docker containers. We next describe each backend component and its role in enabling the capabilities demonstrated in §2.

3.1 Data Preprocessing

This module transforms raw procurement data into analysis-ready text through three coordinated stages. PDFs are parsed with PyMuPDF, cleaned, and combined with raw portal metadata before passing through an NLP pipeline[8] for topic modeling.

3.2 Metadata Extraction and Document Representation

Beyond raw procurement metadata, the system enriches tenders through automated extraction and

[6] Three different datasets are published as open feeds: contracts with a maximum budget less than 15 K€ (*minors*), contracts published in the state procurement portal excluding *minors* (*insiders*), and contracts published in procurement portals from other administrations (*outsiders*).

[7] https://www.public-procurement-data-space.europa.eu/en/dashboards.

[8] https://github.com/lcalvobartolome/NLPipe

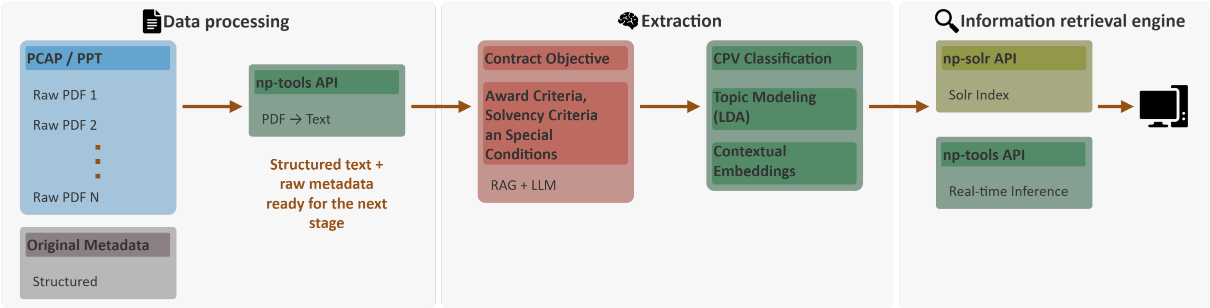

Figure 3: Backend pipeline: (i) raw metadata and PDF text are unified into a structured format; (ii) hidden metadata is extracted via document-level RAG + LLM and enriched with CPV codes, topics, and embeddings; (iii) enriched data is indexed into Solr via `np-solr`, which serves core queries, while `np-tools` provides real-time inference.

multiple document representations to enable flexible retrieval and exploration.

3.2.1 Metadata Extraction

To address gaps in procurement data, the system implements three extraction modules that can be activated independently: tender objectives, award criteria, solvency requirements and special conditions, and CPV codes.

The first two follow a document-level RAG approach: relevant content is retrieved from the raw administrative and technical documents and processed by an LLM[9] to produce structured metadata.

CPV classification is formulated as multi-label text classification over the generated tender objective, as tenders may be associated with multiple CPV codes. The pre-trained Spanish BERT encoder from Cañete et al. (2020) (`dccuchile/bert-base-spanish-wwm-cased`) is fine-tuned at two granularities (5 and 8 digits-depth) using binary cross-entropy with class-imbalance re-weighting. At inference, a validation-calibrated threshold is applied with top-k selection as fallback, typically yielding a single CPV code per tender.

3.2.2 Document representation

Our system relies on two complementary document representations (dense embeddings and topic distributions) which, together with Solr's keyword-based capabilities, support three retrieval modalities: keyword, semantic, and topic-based search.

Topic Modeling We employ LDA via MALLET to uncover latent themes to complement the CPV taxonomy. Based on domain expert feedback, we train one topic model per CPV code to yield focused, relevant topics. Models incorporate topic re-ranking and automatic labeling via LLM prompting and calculation of relevant documents based on the criterion from Calvo-Bartolomé et al. (2025b).

Contextualized Embeddings We generate dense vector representations using the SBERT model `paraphrase-multilingual-mpnet-base-v2` (Reimers and Gurevych, 2019), enabling semantic search beyond keyword matching.

3.3 Information Retrieval Engine

The system exposes two APIs for serving enriched procurement data: *np-solr-api* and *np-tools*. *np-solr-api* extends the CASE framework (Calvo-Bartolomé et al., 2025a) in two ways: by adding dense vector fields for semantic retrieval via `DenseVectorField`, and by maintaining separate collections per CPV category for domain-specific topic and semantic search. It indexes raw and enriched procurement metadata, including document representations and temporal information, and serves all queries that power pAtChWoRK's functionalities. *np-tools* further extends CASE by exposing all the extraction modules as a real-time API, with CPV-specific topic models and fallback to a default model trained on the full corpus.

4 User Study

To evaluate pAtChWoRK's effectiveness and collect feedback for further improvement, we conduct a user study with 5 procurement officials from different Spanish administrations. Participants interact freely with the system for approximately 20 minutes and then complete a survey covering their professional background, system usage, and overall satisfaction (see §A for some of the user cases explored and §B for the complete evaluation survey).

For the pre-indexed collection, we use all *outsider* and *insider* tenders from 2024 with a depth-5

[9] We use `qwen3:32b` with `temperature=0` (see §D).

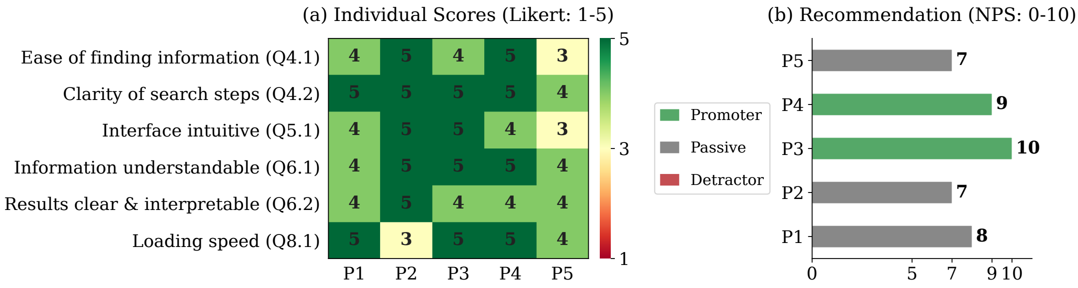

Figure 4: User study results. (a) Likert scores (1–5) per participant (P1–P5) across all evaluated dimensions. (b) Recommendation scores (0–10) with NPS scale: promoters (≥9), passives (7–8), and detractors (≤6).

CPV available, yielding 43,279 tenders for the topic modeling collection.[10] We train two topic models at different granularity levels for CPV-79, which is the depth-2 CPV with the highest number of assigned tenders. For the indicators panels, we consider all 2024-2025 tenders: 719 731 (2024) and 774 699 (2025).

Participant Profiles and Prior Workflows All participants reported consulting procurement information 1-3 times per month, and all had prior experience with existing platforms (most commonly PLACE), with some also using TED or other procurement portals. Their backgrounds span three profiles: *procurement officers* who draft tenders and seek reference contracts (e.g., award criteria, modifications, budget adjustments); *researchers and academics* who monitor open calls relevant to their groups; and *data and AI specialists* who analyse procurement data in an institutional context. Prior to the session, their workflows relied on keyword search, CPV codes, or browsing by contracting authority. Four out of five participants logged in with a clear search goal in mind.

Feature Adoption All participants explored the full tool during the session, visiting every available section. Search terms reflected their professional backgrounds, ranging from domain-specific queries (e.g., *"eventos", events*; *"inteligencia artificial", artificial intelligence*; *"data warehouse"*) to document uploads combined with short-text search. Regarding filters, most participants found them useful and comprehensible, though one noted initially struggling to find results when searching by ID rather than title.

Ease of Use and Navigation As shown in Fig. 4a, participants rated the tool positively across ease-of-use dimensions. The clarity of search steps (Q4.2) received the highest scores (mean 4.8), while ease of finding information (Q4.1) and interface intuitiveness (Q5.1) both averaged 4.2. All participants reported no difficulties during navigation.

Clarity of Results Participants found the information presented largely understandable (Q6.1, mean 4.6) and the results clear and interpretable (Q6.2, mean 4.2). Several participants suggested presentation improvements: making extracted condition fields collapsible, offering alternative visualizations for topics, and returning groupings of similar tenders as an additional result view.

System Performance Loading speed (Q8.1) was rated positively overall (mean 4.4), with three out of five participants awarding the maximum score. Two participants highlighted response agility as one of the tool's main strengths, noting that it could save significant time compared to browsing PLACE.

Satisfaction and Recommendation All participants reported being very satisfied with the tool, with recommendation scores (Fig. 4b) ranged from 7 to 10 (mean 8.2), yielding an NPS of +60 (3 promoters, 2 passives, 0 detractors).

Constructive Critiques Participants identified several areas for improvement. On the data side, they requested broader document coverage beyond a single year, more CPV categories for topic modeling, and the inclusion of adjudicates information and award amounts. On the functionality side, the main requests were filtering by award procedure type, and more compact presentation of award criteria, solvency requirements, and special conditions. One participant suggested a separation between two user profiles (procurement analysts seeking aggregate views, and contracting bodies or companies searching for individual tenders), with a more conversational, RAG-based interface for the latter.

[10] Only 24.82% and 39.00% of all 2024 *outsider* and *insider* tenders have a CPV at this depth.

System or line of work	Raw PDF processing	Metadata extraction	CPV enrichment	Topic modeling	Semantic retrieval	Indicators	Unified operational system
TheyBuyForYou (Soylu et al., 2022a)	No	Partial	No	No	Yes	Partial	Partial
OIE4PA (Siciliani et al., 2023, 2024a)	Partial	Yes	No	No	Partial	No	No
CPV classification approaches (Navas-Loro et al., 2022; Siciliani et al., 2024b; Moiraghi et al., 2024)	No	No	Yes	No	No	No	No
Topic modeling approaches (Rejeb et al., 2023; Hott et al., 2023)	No	No	No	Yes	Partial	No	No
Commercial platforms (Tendios; Sophia Anphis; Gobierto; Mercell)	Partial	Partial	Partial	No	Yes	Partial	Yes
pAtChWoRK	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table 1: Component-level comparison with existing procurement-related systems, tools, and research lines. The table highlights that prior approaches typically cover individual parts of the procurement intelligence pipeline, while pAtChWoRK stitches raw PDF processing, metadata extraction, CPV enrichment, topic modeling, semantic retrieval, indicator-based oversight, and a unified operational interface into a single system.

Open Feedback The features most valued were the speed of keyword-based filtering, the clarity of the interface, the free-text search capability, and the Topic Viewer design. Proposed improvements focused on information layout (e.g., expanding truncated titles), better onboarding for the document upload flow, integration of a PDF viewer across all sections, and linking the tool to its underlying academic documentation to make the NLP methodology more transparent to technically curious users.

5 Related Work

Public procurement aims to ensure transparency and efficiency in public spending (Thai, 2017; OECD, 2024). Despite EU-driven harmonization and digitalization efforts (OECD, 2025; OECD Publishing, 2025), procurement data remain fragmented across TED and national platforms, limiting interoperability and large-scale analytics (Publications Office of the European Union, 2024; European Commission, 2023; Soylu et al., 2022b).

Prior work has addressed this fragmentation from different angles. Research systems such as TheyBuyForYou and OIE4PA apply knowledge graphs and open information extraction to enhance semantic access to procurement data (Soylu et al., 2022a; Siciliani et al., 2023, 2024a). Other lines of work focus on specific analytical components, including topic modeling for procurement analysis (Rejeb et al., 2023; Hott et al., 2023), bias detection (Torres-Berru et al., 2023), and automatic CPV assignment (SIM, 2008). The latter has evolved from traditional classifiers such as SVM (Marzagão, 2015; Cedeño et al., 2024) to multi-label transformer-based models and hierarchical or zero-shot approaches (Navas-Loro et al., 2022; Siciliani et al., 2024b; Moiraghi et al., 2024).

Commercial platforms have also attempted to address procurement needs. In Spain, Tendios and Sophia provide AI-assisted tender discovery, summarization, and sector-based categorization tools (Tendios; Sophia Anphis), while Gobierto improves access to and navigation of CPV hierarchies within Spanish procurement data (Gobierto). At the European level, Mercell offers large-scale tender aggregation and structured indexing services across multiple markets (Mercell). In parallel, the eProcurement Ontology and the PPDS promote semantic standardization and harmonized access to EU procurement data (Publications Office of the European Union; European Commission, 2023).

Table 1 provides a component-level comparison with these research systems, analytical approaches, and commercial tools, making explicit the gap addressed by pAtChWoRK across several dimensions of the procurement intelligence pipeline: raw PDF processing, metadata extraction, CPV enrichment, topic modeling, semantic retrieval, procurement indicators, and unified operational support. Existing work typically covers only a subset of these dimensions, and these capabilities are often developed, evaluated, or deployed as separate components.

pAtChWoRK differs from this prior work in that, while existing systems have advanced procurement analytics along individual dimensions (whether through knowledge-graph representations, standalone CPV classifiers, topic modeling studies, or commercial discovery platforms) they invariably assume structured data as input or remain isolated from operational workflows. pAtChWoRK instead in-

Field	Before	After
Objective = Title	98.96%	0%
Avg. Obj. Length (chars)	173.51	841.86
Award criteria	70.33%	86.73%
Solvency requirements	51.08%	85.59%
Special conditions	44.26%	84.36%
CPV5 code (fine-grained)	85.30%	100%
CPV8 code (fine-grained)	14.69%	100%
Thematic description	✗	✓
Semantic embeddings	✗	✓

Table 2: Quality of procurement metadata before and after pAtChWoRK enrichment. Percentage values represent metadata availability or extraction success rates.

tegrates automated PDF parsing, LLM-based metadata enrichment, multi-label CPV correction, topic modeling, dense semantic indexing, and indicator-based oversight into a unified, research-oriented platform that starts from raw procurement PDFs and bridges unstructured tender documents and structured decision-support services in a workflow tailored to real public procurement practice.

6 Does PAtChWoRK Turn Fragmented Data into a Useful Resource?

The results suggest it largely does. On the data side, Table 2 shows that pAtChWoRK successfully enriches the large majority of tenders with previously unavailable or incomplete metadata: contract objectives are extracted in 96% of cases, award criteria and solvency requirements in over 85%, and CPV codes at fine-grained levels in 100% of cases. On the user side, four out of five participants reported being very satisfied with the tool (NPS = +60), and all highlighted the speed of keyword filtering and free-text search as a major improvement over existing platforms such as PLACE. The main weaknesses were not with the core functionality but with presentation (e.g., truncated titles, incomplete metadata display). These are engineering limitations we plan to address in future work, rather than fundamental shortcomings of the approach. Beyond presentation, we plan to integrate a RAG-based tool to assist users in framing tenders, and to extend the indicators dashboard with topic information following the PPDS[7] model.

7 Limitations

Several limitations should be acknowledged. First, the user study involved only five participants, which limits the generalisability of the satisfaction results; a larger-scale evaluation with a more diverse sample would strengthen these findings. Second, the current document corpus covers a single year and a single CPV category for topic modeling, though both will be expanded and kept up to date as the system is adopted. Third, one participant reported that semantic similarity results were not sufficiently accurate, suggesting that retrieval quality requires further tuning. Finally, the system currently targets Spain, though its modular design allows adaptation to other jurisdictions by adjusting the data pipelines and metadata extraction prompts to different regulatory frameworks.

Acknowledgements

We are grateful to Rubén Notivol Bezares for his help defining use cases and framing the evaluation, and to all participants for their time and feedback. This project has received funding from Grant PID2023-146684NB-I00, funded by MICIU/AEI/10.13039/501100011033 and by ERDF/UE. The work of Erick Cedeño is supported by the Predoctoral Grant PIPF-2024/COM-35132 of the Consejería de Educación, Ciencia y Universidades de la Comunidad de Madrid, Spain.

References

2008. Common procurement vocabulary (cpv). SIMAP/TED portal. CPV 2008 (Regulation (EC) No 213/2008) – catalogue, guides and explanatory notes.

Lorena Calvo-Bartolomé, Jerónimo Arenas-García, and David Pérez-Fernández. 2025a. Case: Large scale topic exploitation for decision support systems. In *Proceedings of the 31st International Conference on Computational Linguistics: System Demonstrations*, pages 151–162.

Lorena Calvo-Bartolomé, Diamantis Tziotzios, Saúl Blanco-Fortes, David Pérez-Fernández, and Jerónimo Arenas-García. 2025b. Topic models for decision-support systems. part i: Training, representation and exploitation of topic models. *Engineering Applications of Artificial Intelligence*, 159:111424.

José Cañete, Gabriel Chaperon, Rodrigo Fuentes, Jou-Hui Ho, Hojin Kang, and Jorge Pérez. 2020. Spanish pre-trained bert model and evaluation data. In *PML4DC at ICLR 2020*.

Erick Cedeño, Maria Navas-Loro, and Oscar Corcho. 2024. Asignación automática del código cpv a la contratación pública. Master's thesis, Universidad Politécnica de Madrid, E.T.S. de Ingenieros Informáticos, Madrid, Spain.

European Commission. 2023. Public procurement: A data space to improve public spending, boost data-driven policy-making and improve access to tenders for smes. Official Journal of the European Union, C 98 I, 16.3.2023, pp. 1–11. Communication from the Commission; C/2023/1696; CELEX: 52023XC0316(02).

Mihály Fazekas and Jurgen Rene Blum. 2021. Improving public procurement outcomes: Review of tools and the state of the evidence base. Technical Report 9690, World Bank.

Gobierto. Visualización y análisis de contratación pública. https://contratos.gobierto.es.

Henrique R Hott, Mariana O Silva, Gabriel P Oliveira, Michele A Brandão, Anisio Lacerda, and Gisele Pappa. 2023. Evaluating contextualized embeddings for topic modeling in public bidding domain. In *Brazilian Conference on Intelligent Systems*, pages 410–426. Springer.

Thiago Marzagão. 2015. Using SVM to pre-classify government purchases. *arXiv preprint arXiv:1601.02680*.

Mercell. Mercell – public procurement & tendering platform. https://www.mercell.com.

Federico Moiraghi, Matteo Palmonari, Davide Allavena, and Federico Morando. 2024. Zero-shot hierarchical classification on the common procurement vocabulary taxonomy. *arXiv preprint arXiv:2405.09983*.

Maria Navas-Loro, Daniel Garijo, and Oscar Corcho. 2022. Multi-label text classification for public procurement in spanish. *Procesamiento del Lenguaje Natural*, 69:73–82.

OECD. 2024. Public procurement. https://www.oecd.org/en/topics/policy-issues/public-procurement.html.

OECD. 2025. *Implementing the OECD Recommendation on Public Procurement in OECD and Partner Countries: 2020–2024 Report*. OECD Public Governance Reviews. OECD Publishing, Paris.

OECD Publishing. 2025. Digital transformation of public procurement: Good practice report. OECD Public Governance Policy Papers 77, Organisation for Economic Co-operation and Development, Paris.

Publications Office of the European Union. eprocurement ontology. https://github.com/OP-TED/ePO.

Publications Office of the European Union. 2024. European public procurement. https://ted.europa.eu/en/simap/european-public-procurement. TED: Tenders Electronic Daily (SIMAP).

Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.

Abderahman Rejeb, Karim Rejeb, Andrea Appolloni, Yasanur Kayikci, and Mohammad Iranmanesh. 2023. The landscape of public procurement research: a bibliometric analysis and topic modelling based on scopus. *Journal of Public Procurement*, 23(2):145–178.

Lucia Siciliani, Eleonora Ghizzota, Pierpaolo Basile, and Pasquale Lops. 2024a. Oie4pa: open information extraction for the public administration. *Journal of Intelligent Information Systems*, 62(1):273–294.

Lucia Siciliani, Vincenzo Taccardi, Pierpaolo Basile, Marco Di Ciano, and Pasquale Lops. 2023. Ai-based decision support system for public procurement. *Information systems*, 119:102284.

Lucia Siciliani, Emanuele Tanzi, Pierpaolo Basile, and Pasquale Lops. 2024b. Automatic generation of common procurement vocabulary. In *Proceedings of the 9th Italian Conference on Computational Linguistics CLiC-it 2023: Venice, Italy, November 30-December 2, 2023*, page 397. Accademia University Press.

Sophia Anphis. Plataforma de inteligencia artificial para licitaciones. https://www.sophia-anphis.es.

Ahmet Soylu, Oscar Corcho, Brian Elvesæter, Carlos Badenes-Olmedo, Tom Blount, Francisco Yedro Martínez, Matej Kovacic, Matej Posinkovic, Ian Makgill, Chris Taggart, et al. 2022a. Theybuyforyou platform and knowledge graph: Expanding horizons in public procurement with open linked data. *Semantic Web*, 13(2):265–291.

Ahmet Soylu, Óscar Corcho, Brian Elvesæter, Carlos Badenes-Olmedo, Francisco Yedro-Martínez, Matej Kovacic, Matej Posinkovic, Mitja Medvešček, Ian Makgill, Chris Taggart, et al. 2022b. Data quality barriers for transparency in public procurement. *Information*, 13(2):99.

Tendios. Plataforma de búsqueda y análisis de licitaciones públicas. https://tendios.com.

Khi V Thai. 2017. *International handbook of public procurement*. Routledge.

Yeferson Torres-Berru, Vivian F Lopez-Batista, and Lorena Conde Zhingre. 2023. A data mining approach to detecting bias and favoritism in public procurement. *Intelligent Automation & Soft Computing*, 36(3).

Appendix

A User Cases

The following scenarios illustrate representative use cases for pAtChWoRK, drawn from real administrative contexts; the first and third were used by one of the evaluators during the user study:

1. The Networks and Systems Service seeks reference tenders for an advisory contract to upgrade its data warehouse.

2. The Digital Transformation Office wants to identify tenders related to citizen kiosks for administrative procedures.

3. The Transparency and Open Government Office is drafting procurement documents for a 3D city visualization service and wants to find similar tenders from other administrations.

4. A user wants to find procurement processes similar to the "Automatic Web Accessibility Analysis Tool" tender from a given City Council, including award amounts from comparable contracts as a pricing reference.

5. The Public Infrastructure Service wants to tender a street element survey to assess their conservation status, and seeks procurement documents from other administrations to identify existing capture technologies available on the market.

6. The Environment Service wants to find tenders related to noise map development.

7. The Mobility Service wants to explore tenders on regulated parking, including special conditions and procurement documents.

8. The Transparency and Open Government Office is drafting an interoperability audit contract and wants to know which award criteria are typically used.

B User Study Survey

Q1. User Profile

Q1.1 How often do you search for information about public tenders?

Q1.2 What type of information do you usually look for?

Q1.3 Briefly describe your professional background and the context in which you consult procurement data.

Q1.4 Have you previously used other tools or platforms to search for tenders? (Yes / No). If yes, which ones?

Q2. Prior Expectations

Q2.1 What is your usual workflow for finding relevant tenders?

Q2.1 Before using the tool, did you have a clear idea of the information you were looking for? (Yes / No)

Q3. System Usage

Q3.1 What terms or criteria did you use in your searches?

Q3.2 Which sections of the tool did you visit during the session?

Q3.3 What filters did you apply? Were they useful and easy to understand?

Q4. Ease of Finding Information

Q4.1 How easy was it to find the information you needed? *(5-point Likert scale: 1 = Very difficult, 5 = Very easy)*

Q4.2 Did you clearly understand the steps required to perform a search and consult the results? *(5-point Likert scale: 1 = Very difficult, 5 = Very easy)*

Q5. Intuitiveness and Navigation

Q5.1 Do you consider the interface intuitive and easy to navigate? *(5-point Likert scale: 1 = Not intuitive at all, 5 = Very intuitive)*

Q5.2 Did you encounter any difficulties moving between the different sections of the tool? (Yes / No). If yes, please describe them.

Q6. Clarity of Presented Information

Q6.1 Did you have any difficulty understanding the information presented? *(5-point Likert scale: 1 = Very difficult to understand, 5 = Very clear)*

Q6.2 Were the displayed results clear and interpretable? *(5-point Likert scale: 1 = Very difficult to understand, 5 = Very clear)*

Q6.3 Would you prefer alternative ways of visualizing or presenting the results? (Yes / No). If yes, please describe them briefly.

Q7. Additional Functionalities

Q7.1 Does the tool need any more functionalities?

Q7.2 What features do you think are missing or need improvement?

Q8. Loading Speed

Q8.1 How would you rate the loading speed of the tool? *(5-point Likert scale: 1 = Not intuitive at all, 5 = Very intuitive)*

Q9. Overall Satisfaction

Q9.1 How satisfied are you overall with the tool? *(5-point Likert scale: 1 = Very unsatisfied, 5 = Very satisfied)*

Q9.2 Would you recommend this tool to other professionals searching for public procurement information? *(NPS scale: 0–10)*

Q10. Open Feedback

Q10.1 What did you like most about the interface and its functionality?

Q10.2 What improvements or changes would you propose to enhance usability / functionality?

Q10.3 (Optional) Any additional comments about the tool or your experience?

C pAtChWoRK's software architecture

Fig. 5 provides a schematic view of the software architecture of pAtChWoRK. The frontend provides an interactive web interface built with HTML5, Bootstrap 5, and JavaScript libraries (jQuery for AJAX, D3.js for visualizations, DataTables for document browsing) and communicates via HTTP/AJAX with a PHP middleware layer that handles CORS and request routing, which in turn exposes all backend NLP services (topic inference, semantic similarity, CPV prediction, and metadata extraction) through a REST API.

D Deployment infrastructure

The system is deployed across two servers. The first (64 GB RAM, 5 cores, 100 GB storage) hosts the Apache web server in a Docker container. The second (128 GB RAM, 10 cores, 400 GB storage) hosts each API in a separate Docker container, the Solr database as an independent container, and Ollama[11] with three NVIDIA GeForce RTX 4090 GPUs for inference.

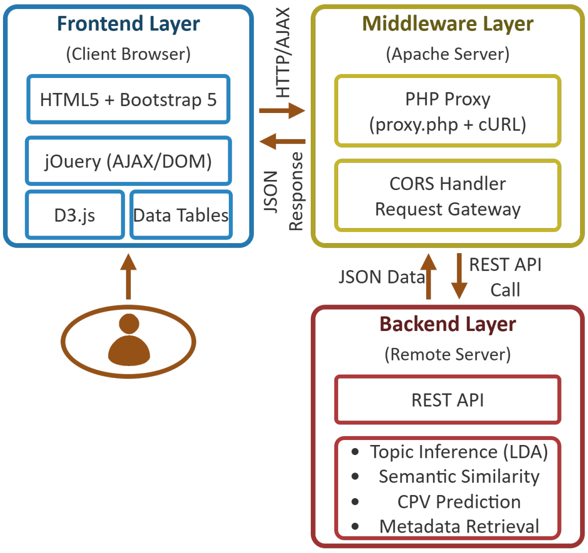

Figure 5: Three-tier architecture of pAtChWoRK.

[11] https://ollama.com/