

DART: A Structured Dataset of Regulatory Drug Documents in Italian for Clinical NLP

Mariano Barone^{1,2}, Antonio Laudante¹, Giuseppe Riccio^{1,2,*}, Antonio Romano^{1,2}, Marco Postiglione³ and Vincenzo Moscato^{1,2}

¹University of Naples Federico II, Department of Electrical Engineering and Information Technology (DIETI), Via Claudio, 21 - 80125 - Naples, Italy

²Consorzio Interuniversitario Nazionale per l'Informatica (CINI) - ITEM National Lab, Complesso Universitario Monte S. Angelo, Naples, Italy

³Northwestern University, Department of Computer Science, McCormick School of Engineering and Applied Science, 2233 Tech Dr, Evanston, IL 60208, United States

Abstract

The extraction of pharmacological knowledge from regulatory documents has become a key focus in biomedical natural language processing, with applications ranging from adverse event monitoring to AI-assisted clinical decision support. However, research in this field has predominantly relied on English-language corpora such as DrugBank, leaving a significant gap in resources tailored to other healthcare systems. To address this limitation, we introduce DART (Drug Annotation from Regulatory Texts), the first structured corpus of Italian Summaries of Product Characteristics derived from the official repository of the Italian Medicines Agency (AIFA). The dataset was built through a reproducible pipeline encompassing web-scale document retrieval, semantic segmentation of regulatory sections, and clinical summarization using a few-shot-tuned large language model with low-temperature decoding. DART provides structured information on key pharmacological domains such as indications, adverse drug reactions, and drug-drug interactions. To validate its utility, we implemented an LLM-based drug interaction checker that leverages the dataset to infer clinically meaningful interactions. Experimental results show that instruction-tuned LLMs can accurately infer potential interactions and their clinical implications when grounded in the structured textual fields of DART. We publicly release our code on GitHub: <https://github.com/PRAISELab-PicusLab/DART>.

Keywords

Pharmacological Text Mining, Adverse Drug Reactions, Drug-Drug Interactions, Italian Biomedical NLP

1. Introduction

In recent times, extracting and organizing pharmacological information from regulatory documents has taken on a pivotal role in the domain of biomedical natural language processing (NLP). This research goal is focused on automating the assimilation of clinical and regulatory data into decision-oriented processes [1], thereby supporting applications like prescription aid systems and pharmacovigilance instruments. Among these regulatory resources, the Summary of Product Characteristics (SmPC) — referred to in Italy as the Riassunto delle Caratteristiche del Prodotto (RCP) — is notably distinguished as a comprehensive and reliable document published by the Italian Medicines Agency¹ (AIFA). Designed for healthcare professionals, the RCP serves as the 'identity card' of a medicinal product, providing standardized and regularly updated information on efficacy, safety, therapeutic use, contraindications, adverse drug reactions (ADR), drug-drug interactions (DDI), and other essential clinical characteristics.

ITADATA2025: The 4th Italian Conference on Big Data and Data Science, September 9–11, 2025, Turin, Italy

*Corresponding author.

✉ mariano.barone@unina.it (M. Barone); antonio.laudante@unina.it (A. Laudante); giuseppe.riccio3@unina.it (G. Riccio); antonio.romano5@unina.it (A. Romano); marco.postiglione@northwestern.edu (M. Postiglione); vincenzo.moscato@unina.it (V. Moscato)

🌐 <https://github.com/csmariano> (M. Barone); <https://github.com/giuseppericcio> (G. Riccio); <https://github.com/LaErre9> (A. Romano); <http://wpagina.unina.it/vmoscato/> (V. Moscato)

🆔 0009-0004-0744-2386 (M. Barone); 0009-0002-8613-1126 (G. Riccio); 0009-0000-5377-5051 (A. Romano); 0000-0001-6092-940X (M. Postiglione); 0000-0002-0754-7696 (V. Moscato)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

¹<https://www.aifa.gov.it/>

Despite its importance, such texts remain underrepresented in the literature, with most prior work focusing exclusively on English-language corpora and overlooking the linguistic and structural particularities of national regulatory frameworks. In the Italian context, the absence of tailored resources hampers the development of clinically grounded AI (Artificial Intelligence) systems that align with local healthcare practices and regulatory standards. To address this gap, we present DART (Drug Annotation from Regulatory Texts), a structured corpus of RCPs in Italian, developed through a scalable and reproducible pipeline. The dataset is built by automatically retrieving documents from AIFA, extracting and semantically segmenting their contents, and organizing the information into structured fields that correspond to standard regulatory sections. Additionally, DART is enhanced with clinical summaries generated using large language models (LLMs) via few-shot learning and low-temperature decoding strategies. These summaries are intended to support downstream applications such as interaction checking, knowledge graph construction, and automated risk profiling. With more than 16,000 processed RCPs and over 95 million tokens, DART represents a high-value asset for the Italian clinical NLP community and the broader healthcare data science ecosystem. It provides a robust foundation for the training, evaluation, and deployment of large-scale language models in both regulatory and clinical contexts. Furthermore, DART contributes significantly to the healthcare Big Data ecosystem by offering a high-resolution corpus of regulatory texts that supports the training of LLMs, the development of interpretable knowledge graphs, and the implementation of AI-driven clinical decision-making tools.

2. Related Work

The extraction of pharmacological knowledge from regulatory texts—such as Summary of Product Characteristics (SmPC)—is a growing area in biomedical NLP. These documents offer authoritative information on adverse drug reactions (ADRs), drug–drug interactions (DDIs), contraindications, and indications, and form the normative basis for safe prescribing. However, most existing work has focused on English-language corpora, leaving national regulatory texts, especially Italian RCPs, underrepresented. Early ADR extraction relied on classical machine learning models, including ensemble methods and multilayer perceptrons [2, 3]. The adoption of transformer-based architectures such as BERT, BioBERT, and PubMedBERT significantly improved performance [4], though non-English texts still require costly adaptation and fine-tuning [5]. More recently, large language models (LLMs) like GPT-4 have shown strong zero- and few-shot performance in biomedical tasks, including ADR detection, outperforming traditional baselines and enhancing interpretability in pharmacovigilance pipelines [6, 7, 8]. Retrieval-augmented generation and agent-based simulation approaches have further demonstrated the value of context-aware models [9, 10]. DDI prediction has similarly evolved toward hybrid and graph-based architectures. Recent studies integrate knowledge graphs (KGs) with LLMs to produce accurate and explainable predictions [11, 12], while in-context learning techniques have improved interaction detection [13]. Medication recommendation systems now incorporate regulatory text and clinical narratives, outperforming structured-code-based methods, especially in multilingual and safety-aware settings [14, 15, 16]. Ongoing work also explores explainability in recommendations [17] and the combination of symbolic and generative approaches in medical summarization [18]. Despite this progress, Italian regulatory documents remain largely unexplored. Resources like DrugBank [19] include Italian drug names but abstract away regulatory phrasing and section structure. Challenges such as DIMMI [20] and aggregation efforts [21] highlight the need for domain-specific resources aligned with the Italian context. Real-world data sources like the National Pharmacovigilance Network (RNF) [22, 23], VALORE [24], and regional datasets [25] offer complementary insights but are often incomplete, misaligned with regulatory language, or dependent on patient self-reporting [26]. In contrast, RCPs provide standardized, high-quality knowledge that enables direct modeling of pharmacological phenomena [27]. To fill this gap, we present DART, a structured dataset derived from full-text Italian RCPs, designed to support the development of LLM-based systems grounded in official regulatory content. DART enables validation of LLM outputs against both normative sources and observational datasets, fostering a bidirectional loop between automated pharmacological reasoning and real-world

clinical safety evidence.

3. Materials and Methods

3.1. Dataset Construction

The dataset DART was constructed through a three-step pipeline designed for reproducibility and scalability. Specifically: (i) automated retrieval of URLs for the Summary of Product Characteristics (RCPs) from the AIFA portal, (ii) semantic parsing and segmentation of the extracted RCP text, and (iii) data structuring, filtering, and validation. All modules were implemented in Python using open-source libraries. The complete construction workflow is illustrated in Figure 1.

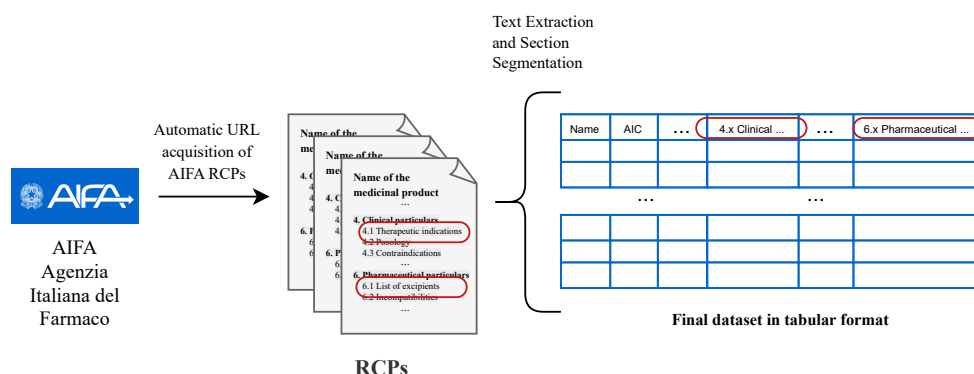


Figure 1: End-to-end pipeline for the construction of the DART dataset. The diagram outlines the three main components of the workflow: automated retrieval of RCP documents from AIFA, text extraction and segmentation, and the final structuring and validation of the resulting dataset.

3.1.1. Automated Retrieval of RCP URLs

The first phase of the pipeline involved the programmatic retrieval of RCP PDFs by interrogating undocumented but publicly accessible RESTful APIs exposed by the AIFA web portal. Due to the SPA-based (Single Page Application) architecture of the website—built using frameworks such as Angular—static DOM scraping was ineffective. Instead, a detailed network traffic analysis was conducted via browser developer tools (*DevTools*, “Network” tab), which led to the identification of two critical endpoints. The first is a search endpoint, which requires a zero-padded AIC code (e.g., 123456 becomes 00123456) and returns a JSON payload containing metadata for each drug, including the keys `CodiceSis` and `aic6`. These values are then used to query a second endpoint that provides a direct URL to the corresponding RCP PDF. This two-step API interaction is illustrated in Figure 2. Although these APIs are unofficial and subject to change without notice, they provided the only viable and scalable access to RCP documents at the time of this study (June 2025). A web spider was implemented in Python using the `requests` library and seeded with a list of valid AIC codes, sourced from public datasets or inferred from known numerical intervals, in compliance with applicable ethical and legal constraints. For each AIC code, the system executed: querying the search endpoint, parsing the JSON response, constructing the PDF URL, and downloading the file. Failures and exceptions were handled using structured logging.

3.1.2. Text Extraction and Section Segmentation

Once the documents were collected, the pipeline proceeded with text extraction and semantic segmentation. Text was extracted using the `PyMuPDF` library, selected for its robustness in handling complex layouts, preserving reading order, and maintaining basic spatial formatting where feasible. This approach proved effective in most cases, except for PDFs consisting solely of rasterized images (i.e.,

Example API Call for RCP PDF Retrieval

Step 1: Query the Search Endpoint

`https://api.aifa.gov.it/aifa-bdf-eif-be/1.0.0/formadosaggio/ricerca?
query={AIC_code}&spellingCorrection=true&page=0`

Note: The input {AIC_code} must be a zero-padded version of the original AIC code (e.g., 123456 → 00123456).

Response: JSON object containing:

- CodiceSis (e.g., 10004290)
- aic6 (e.g., 123456)

Step 2: Construct the PDF Download URL

`https://api.aifa.gov.it/aifa-bdf-eif-be/1.0.0/organizzazione/{CodiceSis}/
farmaci/{aic6}/stampati?ts=RCP`

Example:

`https://api.aifa.gov.it/.../organizzazione/10004290/farmaci/123456/stampati?ts=RCP`

Output: Direct link to the corresponding RCP PDF.

Figure 2: Two-step RESTful API interaction for programmatic retrieval of RCP PDFs from AIFA's infrastructure.

scanned documents), which lack an embedded text layer. Approximately 4.1% of collected PDFs were excluded due to the absence of an embedded text layer, making them incompatible with text-based parsing. These cases are flagged for future integration through OCR modules, which are currently under development. The text structuring phase was based on the automatic identification of section headers, which follow well-defined regulatory conventions in RCPs (e.g., "04.1 Therapeutic Indications", "04.8 Undesirable Effects"). A robust regular expression was designed to recognize both the numerical and textual components of the headers, accounting for typographic variability (e.g., spacing, punctuation, capitalization). This enabled segmentation of each document into blocks corresponding to individual sections, each assigned a standardized label. Sections not detected were marked as "N/A" in the resulting dataset, preserving the structural consistency of the data model.

Tabular Content Handling Special attention was devoted to Section 04.8 ("Undesirable Effects") often includes tabular structures. While full table parsing was out of scope in this version, raw text within tables was preserved using PyMuPDF's line-by-line reading mechanism, which retains spatial alignment. Although columnar relationships are not explicitly modeled, the output allows partial semantic interpretation. Future iterations will integrate table extraction tools such as pdfplumber, camelot, or layout-aware parsing models.

3.1.3. Data Structuring and Validation

Extracted data were finally mapped into a tabular dataset, where each row corresponds to an RCP document and each column to a specific regulatory section. Final validation included completeness checks (e.g., verifying the presence of expected sections), spot comparisons between raw documents and extracted text, and analysis of error logs produced by the spider. The combined application of these methods enabled the construction of a coherent, scalable dataset suitable for downstream analyses in pharmacological, linguistic, and computational research contexts. Validation steps included logging and error tracking using loguru and structured output reports. On a random sample of 300 documents, over 97% of expected sections were correctly identified and segmented. Remaining errors were mostly due to non-standard formatting or OCR failures.

3.2. Preprocessing and Filtering

Following initial structuring, DART contains 21,502 drugs, subject to a preprocessing step was applied to improve the consistency and correctness of the dataset. This phase involved regex-based cleaning to standardize formatting, eliminating excess whitespace, resolving punctuation inconsistencies, and uniforming typographic variances in section headers. Documents with empty or flawed text were identified and excluded. The "05.0 Pharmacological Properties" section was removed entirely due to a high rate of missing or unusable data, impacting content density and usability. This led to a dataset with strong structural consistency and semantic integrity, appropriate for various clinical NLP applications. Ultimately, 16,029 documents (74.55%) were correctly segmented into at least 5 mandatory regulatory sections, while the remaining 25.45% were removed due to structural problems or incomplete content, often originating from OCR errors or missing data.

3.3. RCP Summarization

To improve the usability of the dataset for both human users and NLP systems, a summarization phase was introduced to condense long and heterogeneous regulatory sections into standardized clinical summaries. This step facilitates tasks such as text classification, knowledge extraction, semantic search, and decision support, while also enabling rapid inspection by clinicians and analysts. Summaries were generated using LLaMA 3.1-405B[28] through Nvidia NIM API², a state-of-the-art large language model, with a low-temperature setting (0.2) to ensure high consistency and minimal hallucination. Each summary was limited to 450 words and aimed to capture key information on drug interactions, adverse events, contraindications, warnings, and pregnancy-related considerations. To guide the generation, we employed a structured prompt combined with a few-shot learning strategy. Handcrafted examples were prepended to the prompt to ensure alignment with regulatory tone, content structure, and domain terminology. Input text was extracted from seven RCP sections (04.1, 04.2, 04.3, 04.4, 04.5, 04.6, and 04.8) and dynamically inserted into the prompt. The resulting summaries were integrated into the dataset as an additional field, enhancing its value for downstream applications and enabling comparisons with real-world DDI/ADR evidence. A manual review of 100 generated summaries showed a high degree of factual consistency (95%) and minimal hallucination. Most deviations involved stylistic variation or omission of low-priority details. An expert-based validation protocol is currently under development.

3.4. Dataset Analysis

DART consists of 16,029 documents, spanning multiple therapeutic areas and regulatory reimbursement classes. The dataset was last updated in May 2025. The corpus comprises over 95 million tokens, with a vocabulary of 102,749 unique terms. Document lengths are generally compact: the mean token count per document is 177.5 (median: 168.3), with a maximum of 9,512 tokens.

Table 1
Overview statistics of the DART dataset

Statistic	Value
# Number of Medicines	16,029
# Number of Therapeutic Classes	6
Last Update	May 2025
Total Tokens	95,760,718
Unique Vocabulary	102,749
Avg. Tokens per Document	177.47

Table 2
Distribution by Reimbursement Class (AIFA Classification)

Class Code	Frequency
C	5,406
A	5,156
N	1,842
C-nn	1,724
H	1,599
C-bis	942

The reimbursement classes reflect regulatory categories distinct from therapeutic classes (including six ATC Level 1 codes). Classes C and A dominate the dataset, representing non-reimbursed and

²<https://build.nvidia.com/>

essential reimbursed drugs, respectively. The subclasses C-nn and C-bis are nested under C, which explains the sum exceeding the total document count.

Section Coverage and Quality Metrics. We evaluated the presence of key regulatory sections across documents to assess completeness and usability for NLP tasks. Table 3 reports the coverage of selected sections critical for pharmacological information extraction. The results indicate a high degree of consistency, with most sections present in over 90% of the RCPs, ensuring reliable availability of therapeutic indications, dosage information, contraindications, interactions, and adverse effects for computational analysis. Only section 04.6 Pregnancy/Lactation has a slightly lower percentage of coverage (89.6%).

Table 3
Coverage of key sections of the RCPs in the dataset DART

Section Code	Section Name	Coverage (%)
04.1	Therapeutic Indications	97.4
04.2	Posology and Use	95.9
04.3	Contraindications	94.1
04.5	Interactions	92.7
04.6	Pregnancy/Lactation	89.6
04.8	Undesirable Effects	93.2

Lexical and Semantic Insights. The vocabulary size ($\approx 103k$ unique terms) includes a significant portion of technical jargon, multi-token entities, and standard pharmaceutical terminology. Key pharmacological terms (e.g., “interactions”, “pregnancy”, “contraindications”) occur with high frequency across classes, supporting targeted NLP extraction. Further lexical analysis is ongoing to quantify the proportion of domain-specific terms and evaluate term distribution across reimbursement and therapeutic classes.

4. Applications

The DART dataset derived from RCPs supports high-precision tasks in computational pharmacovigilance, structured biomedical information extraction, and explainable clinical decision support systems. The dataset applies a semantic structuring pipeline, which categorizes regulatory text into standardized groups (e.g., interactions, contraindications, adverse effects), allowing traceable links to source sections, aligning with regulatory demands, and enhancing interpretability, notably during clinical or legal audits. This structured, authoritative data anchors AI-driven systems, fostering the development of reliable, explainable tools for clinicians, pharmacists, and health IT systems. Key applications of this dataset are outlined below.

4.1. LLM-based Drug-Drug Interaction Checker

To assess the effectiveness of the DART dataset in the context of automated processing of regulatory information, we designed and implemented an advanced system for the identification of drug-drug interactions (DDIs), leveraging the capabilities of LLMs. The system takes as input a set of drugs $F = (F_1, F_2, \dots, F_N)$ each represented through its active ingredient and the structured sections of the RCP, extracted directly from the DART dataset. RCPs, being rich and complex technical documents, contain relevant information for DDI detection dispersed across heterogeneous sections such as “Warnings and Precautions”, “Interactions”, or “Pharmacokinetic Properties”. However, direct analysis of the full text proves suboptimal for LLMs due to both input length limitations and high semantic dispersion. In order to tackle these issues, as outlined in Section 3.3, we implemented a regulatory summarization

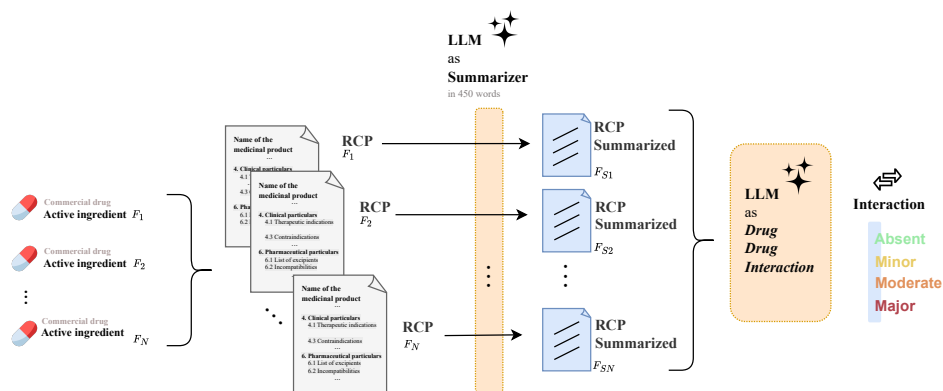


Figure 3: System pipeline for automatic detection of drug–drug interactions (DDIs). Each commercial drug, associated with its active ingredient (F_1, F_2, \dots, F_N) is linked to its respective Summary of RCPs. From each RCP, the system extracts clinically and pharmacologically relevant sections. These sections are then processed by the Summarizer component, it is an LLM which generates a structured summary (max. 450 words). The resulting summaries are fed into a LLM, which analyzes all drug pairs and classifies the interaction severity as absent, minor, moderate, or major with prompt engineering.

module in which each drug is paired with a corresponding summary, denoted as $F_S = (F_{S1}, F_{S2}, \dots, F_{SN})$. This feature, built upon an LLM, produces an organized summary concentrating solely on components that could be pertinent to analyzing pharmacological interactions. In this initial phase, the system markedly reduces the complexity of the regulatory text, directing the model’s focus toward clinically relevant concepts while ensuring compliance with the computational constraints of current LLMs. The resulting summary is then forwarded to the LLM-as-DDI module, which—through the application of targeted prompt engineering techniques—detects potential interactions between active pharmaceutical ingredients, elucidates the underlying pharmacological mechanisms (such as receptor synergies or enzymatic pathways), and assesses the clinical relevance of each interaction based on the patient’s profile. This process is followed by the formulation of context-specific recommendations—such as dosage adjustments or monitoring requirements—and the assignment of a severity level to each identified interaction. The full pipeline is illustrated in Figure 3 and an example end-to-end is showed in Figure 4. Interactions are categorized into four ascending levels of clinical severity—*Absent*, *Minor*, *Moderate*, and *Major*—in accordance with taxonomies commonly employed in scientific research. Specifically, *Absent* indicates the lack of any known or clinically meaningful interaction between the drugs; *Minor* denotes a pharmacological interaction of negligible clinical relevance, typically not requiring any intervention; *Moderate* refers to a clinically significant interaction that may necessitate monitoring or dosage adjustments; and *Major* implies a severe interaction, which is either contraindicated or requires substantial modifications to the therapeutic regimen. To facilitate comparison with widely used online tools such as Drugs.com³, Medscape⁴, WebMD⁵, and RxList⁶, which adopt a binary classification framework, we employed a simplified model that consolidates the *Minor*, *Moderate*, and *Major* categories into a single class, labeled *Interaction*, while retaining *Absent* as a distinct category. This adaptation ensures compatibility with systems commonly used in clinical practice. Performance was evaluated using standard metrics—Precision, Recall, F1-score, and Accuracy—on a manually annotated test set comprising 100 examples. Particular emphasis was placed on Recall, as it serves as a critical metric for assessing the system’s ability to detect all clinically relevant drug–drug interactions (DDIs). In medical contexts, achieving high Recall is essential to minimize false negatives and thereby ensure patient safety. Table 4 highlights the substantial advantages of the proposed framework. The upper

³https://www.drugs.com/drug_interactions.html

⁴<https://reference.medscape.com/drug-interactionchecker>

⁵<https://www.webmd.com/interaction-checker/default.htm>

⁶<https://www.rxlist.com/drug-interaction-checker.htm>

section of the table presents the results obtained from four established web-based tools. A range of large language models (LLMs) were tested in standalone configuration, including both closed-source models (GPT-4o, Claude, Gemini) and open-source models (LLaMA, Mistral, Gemma). Additionally, the table includes performance data for open-source models enhanced with regulatory summaries generated via the DART system. Comparative analysis reveals that certain closed-source models, such as Claude-3.5 and GPT-4o, achieve performance comparable to or exceeding that of conventional clinical tools. However, the incorporation of DART summarization emerges as a pivotal factor. For instance, the configuration LLaMA-3.1-8B + DART achieves a Recall of 0.843, substantially outperforming the same model without summarization (Recall = 0.229), and surpassing most of the evaluated web-based systems. These findings underscore the critical role of guided regulatory summarization in enhancing DDI detection capabilities without compromising precision. Overall, the results validate the effectiveness of the proposed framework: integrating the DART dataset with advanced language models enables even lightweight open-source architectures to effectively identify complex pharmacological interactions. This approach demonstrates the potential to rival high-end proprietary systems, offering an optimal balance of accuracy, coverage, and computational efficiency—factors essential for practical deployment in both clinical and regulatory domains.

Table 4

Comparative Evaluation of DDI Detection Tools: Web-based Systems, Standalone LLMs (Closed/Open Source), and LLMs Enhanced with the DART Framework. Results are reported for binary classification performance (Absent vs. Interaction).

Web-based Drug Interaction Tools				
Tool	Precision	Recall	F1-score	Accuracy
Drugs.com ⁷	0.701	0.812	0.756	0.76
Medscape ⁸	0.761	0.712	0.734	0.74
WebMD ⁹	0.682	0.673	0.675	0.675
RxList ¹⁰	0.663	0.633	0.645	0.645
LLM Standalone				
Model	Precision	Recall	F1-score	Accuracy
<i>Closed-source Models</i>				
Gemini-2.5-Flash	0.814	0.500	0.619	0.657
GPT-4o	0.688	0.786	0.733	0.737
Claude-3.5	0.730	0.771	0.750	0.750
<i>Open-source Models</i>				
Gemma-2-9b	0.750	0.214	0.333	0.554
LLaMA-3.1-8B	0.800	0.229	0.356	0.561
Mistral-7B	0.842	0.229	0.360	0.565
LLM + DART (Our Demonstration)				
Model	Precision	Recall	F1-score	Accuracy
Gemma-2-9b + DART	0.711	0.457	0.557	0.684
LLaMA-3.1-8B + DART	0.728	0.843	0.781	0.786
Mistral-7B + DART	0.778	0.300	0.433	0.640

4.2. Other Applications

Training and Fine-tuning of Multilingual NLP Models. The dataset DART serves as a natural benchmark for training and fine-tuning NLP models specialized in Named Entity Recognition (NER) and Relation Extraction (RE) in Italian, particularly for regulatory and clinical pharmacology domains. It includes entities such as active substances, administration routes, pharmacokinetic mechanisms, and

Illustrative Example of Drug–Drug Interaction Detection using the DART Framework

Input:

Drug F1: Warfarin
Active Ingredient: Warfarin

Drug F2: Ibuprofen
Active Ingredient: Ibuprofen

Step 2 – Extract Summarized RCPs:

RCP for Warfarin → Summarized RCP F1 (≈ 450 words)

RCP for Ibuprofen → Summarized RCP F2 (≈ 450 words)

Step 3 – Compare Summaries to Detect Interaction:

The LLM receives the summarized RCPs F1 and F2, then prompting them → Evaluates interaction risk, mechanism, and severity

Output:

Interaction Detected: **Major**

Drug Pair: Warfarin + Ibuprofen

Mechanism: Inhibition of CYP2C9 by ibuprofen increases the bleeding risk associated with warfarin

Recommendation: Avoid co-administration or closely monitor INR levels

Figure 4: An end-to-end example of drug–drug interaction detection using the DART framework. The system retrieves, summarizes, and compares documents to assess clinical relevance.

pregnancy risk categories. Thanks to its semantic consistency and structural regularity, the dataset supports both supervised training and distant supervision, filling a critical gap in the multilingual biomedical NLP landscape, which remains heavily English-centric.

Fine-tuning of Domain-specific LLMs or SLMs. The normalized corpus of RCP texts offers a unique foundation for domain-specific fine-tuning of LLMs or Small Language Models (SLMs) tailored to the Italian pharmaceutical regulatory domain. Potential downstream applications include: Automatic classification of clinical risks from free text; Assisted generation of pharmacovigilance reports; Controlled rewriting of regulatory documents (e.g., technical leaflets, RCPs). Such models could significantly enhance automation and consistency in regulatory workflows, particularly in contexts requiring traceable and explainable outputs.

Construction of Regulatory Knowledge Graphs. The extracted relational triples (e.g., active substance → causes → adverse effect, drug → interacts with → compound) can be transformed into semantic knowledge graphs (KGs). These KGs support automated inference over contraindications and interactions, and allow structured linking between regulatory sources (e.g., RCPs) and observational data (e.g., national registries such as RNF or VALORE). Moreover, KGs facilitate the generation of explainable clinical decision rules, increasing transparency and trust in AI-powered systems.

Semi-automated Population of Clinical Decision Support Systems (CDSS). The structured dataset is highly suitable for integration into next-generation Clinical Decision Support Systems that combine structured knowledge (e.g., ontologies, terminologies) with unstructured textual evidence. Data extracted from RCPs can be used to populate modules within Electronic Health Records (EHRs), generate safety alerts in hospital pharmacy systems, or support real-time prescription checks. The goal is to enhance patient safety and prescribing appropriateness, by anchoring decisions to verified, regulatory-grade information.

5. Conclusion & Future Work

This work introduces a structured and scalable method for transforming Italian Summary of Product Characteristics (RCPs) into machine-readable resources for biomedical AI. Through semantic parsing and organization, we demonstrate their applicability in multiple domains, including pharmacological interaction checking, domain-specific model tuning, knowledge graph creation, and clinical decision support systems. Despite their linguistic variability, RCPs offer a strong foundation for transparent and regulation-compliant AI systems. The resulting dataset serves both as a benchmark for multilingual biomedical NLP and as a driver for innovation in pharmacovigilance and clinical AI. Future developments will aim to extend coverage to additional regulatory document types and therapeutic areas, improve prompt and alignment strategies, introduce validation processes with domain experts, and publish reusable tools and subsets to support open research in regulatory science.

Limitations Although DART represents a relevant step for Italian biomedical NLP, some limitations apply. It includes only RCPs, thus lacking real-world clinical nuances such as patient adherence or off-label use. Not all AIFA-listed medicines are included due to technical issues like malformed or inaccessible documents, potentially underrepresenting some drug categories. Additionally, the LLM-based components, while optimized for factual consistency, may miss rare or context-specific details and remain sensitive to prompt design and model variability.

Ethical Issues Since DART relies exclusively on publicly available regulatory texts intended for healthcare professionals, it presents minimal direct ethical risks. However, caution is necessary when using generated outputs in clinical contexts, as language models may propagate inaccuracies, particularly in sensitive areas like drug safety. Human oversight remains essential, and future work should include expert review and mechanisms to flag uncertainty.

Data License and Copyright Issues All documents were sourced from the official website of the Italian Medicines Agency (AIFA) under public access policies. The DART dataset is released under the *Creative Commons Attribution 4.0 International License (CC BY 4.0)*, allowing reuse with proper attribution. However, the original RCPs remain property of AIFA, and downstream use must respect ethical and regulatory guidelines.

Acknowledgments

This work was conducted with the financial support of (1) the PNRR MUR project PE0000013-FAIR and (2) the Italian ministry of economic development, via the ICARUS (Intelligent Contract Automation for Rethinking User Services) project (CUP: B69J23000270005).

Declaration on Generative AI

During the preparation of this work, the author(s) used ChatGPT and DeepL in order to: Grammar and spelling check. After using these tool(s)/service(s), the author(s) reviewed and edited the content as needed and take(s) full responsibility for the publication's content.

References

- [1] S. Velupillai, H. Suominen, M. Liakata, A. Roberts, A. D. Shah, K. Morley, D. Osborn, J. Hayes, R. Stewart, J. Downs, et al., Using clinical natural language processing for health outcomes research: overview and actionable suggestions for future advances, *Journal of biomedical informatics* 88 (2018) 11–19.

- [2] E. Hong, J. Jeon, H. U. Kim, Recent development of machine learning models for the prediction of drug-drug interactions, *Korean Journal of Chemical Engineering* 40 (2023) 276–285.
- [3] S. Abbas, G. A. R. Sampredo, M. B. Abisado, A. S. Almadhor, T. Kim, M. M. Zaidi, A novel drug-drug indicator dataset and ensemble stacking model for detection and classification of drug-drug interaction indicators, *IEEE Access* 11 (2023) 101525–101536. URL: <https://doi.org/10.1109/ACCESS.2023.3315241>. doi:10.1109/ACCESS.2023.3315241.
- [4] B. Portelli, E. Lenzi, E. Chersoni, G. Serra, E. Santus, BERT prescriptions to avoid unwanted headaches: A comparison of transformer architectures for adverse drug event detection, in: P. Merlo, J. Tiedemann, R. Tsarfaty (Eds.), *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, EACL 2021, Online, April 19 - 23, 2021, Association for Computational Linguistics, 2021, pp. 1740–1747. URL: <https://doi.org/10.18653/v1/2021.eacl-main.149>. doi:10.18653/v1/2021.EACL-MAIN.149.
- [5] A. Romano, G. Riccio, M. Postiglione, V. Moscato, Identifying cardiological disorders in spanish via data augmentation and fine-tuned language models, in: G. Faggioli, N. Ferro, P. Galuscáková, A. G. S. de Herrera (Eds.), *Working Notes of the Conference and Labs of the Evaluation Forum (CLEF 2024)*, Grenoble, France, 9-12 September, 2024, volume 3740 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2024, pp. 207–222. URL: <https://ceur-ws.org/Vol-3740/paper-19.pdf>.
- [6] M. Gope, J. Wang, Using llms to extract adverse drug reaction (ADR) from short text, in: J. Kim, R. C. Conceição, M. Yousef, A. Bhavsar, S. Pelayo, A. Fred, H. Gamboa (Eds.), *Proceedings of the 18th International Joint Conference on Biomedical Engineering Systems and Technologies, BIOSTEC 2025 - Volume 2: HEALTHINF*, Porto, Portugal, February 20-22, 2025, SCITEPRESS, 2025, pp. 548–555. URL: <https://doi.org/10.5220/0013160700003911>. doi:10.5220/0013160700003911.
- [7] B. M. J. Alshehri, N. Kraiem, H. Sakly, N. Alasbali, Enhancing medication safety with large language models: Advanced detection and prediction of drug-drug interactions, in: *7th IEEE International Conference on Advanced Technologies, Signal and Image Processing, ATSIP 2024*, Sousse, Tunisia, July 11-13, 2024, IEEE, 2024, pp. 547–552. URL: <https://doi.org/10.1109/ATSIP62566.2024.10638993>. doi:10.1109/ATSIP62566.2024.10638993.
- [8] R. J. AbuNasser, M. Z. Ali, Y. Jararweh, M. Daraghme, T. Z. Ali, Large language models in drug discovery: A comprehensive analysis of drug-target interaction prediction, in: *2nd International Conference on Foundation and Large Language Models, FLLM 2024*, Dubai, United Arab Emirates, November 26-29, 2024, IEEE, 2024, pp. 417–431. URL: <https://doi.org/10.1109/FLLM63129.2024.10852448>. doi:10.1109/FLLM63129.2024.10852448.
- [9] R. Russo, D. Russo, G. M. Orlando, A. Romano, G. Riccio, V. L. Gatta, M. Postiglione, V. Moscato, Europeanlawadvisor: an open source search engine for european laws, in: W. Ding, C. Lu, F. Wang, L. Di, K. Wu, J. Huan, R. Nambiar, J. Li, F. Ilievski, R. Baeza-Yates, X. Hu (Eds.), *IEEE International Conference on Big Data, BigData 2024*, Washington, DC, USA, December 15-18, 2024, IEEE, 2024, pp. 4751–4756. URL: <https://doi.org/10.1109/BigData62323.2024.10826025>. doi:10.1109/BIGDATA62323.2024.10826025.
- [10] A. Ferraro, A. Galli, V. L. Gatta, M. Postiglione, G. M. Orlando, D. Russo, G. Riccio, A. Romano, V. Moscato, Agent-based modelling meets generative AI in social network simulations, in: L. M. Aiello, T. Chakraborty, S. Gaito (Eds.), *Social Networks Analysis and Mining - 16th International Conference, ASONAM 2024*, Rende, Italy, September 2-5, 2024, *Proceedings, Part I*, volume 15211 of *Lecture Notes in Computer Science*, Springer, 2024, pp. 155–170. URL: https://doi.org/10.1007/978-3-031-78541-2_10. doi:10.1007/978-3-031-78541-2_10.
- [11] C. Xu, K. C. Bulusu, H. Pan, O. Elemento, Ddi-gpt: Explainable prediction of drug-drug interactions using large language models enhanced with knowledge graphs, *BioRxiv* (2024) 2024–12.
- [12] D. Russo, G. M. Orlando, A. Romano, G. Riccio, V. L. Gatta, M. Postiglione, V. Moscato, Scaling llm-based knowledge graph generation: A case study of italian geopolitical news, in: W. Ding, C. Lu, F. Wang, L. Di, K. Wu, J. Huan, R. Nambiar, J. Li, F. Ilievski, R. Baeza-Yates, X. Hu (Eds.), *IEEE International Conference on Big Data, BigData 2024*, Washington, DC, USA, December 15-18, 2024, IEEE, 2024, pp. 3494–3497. URL: <https://doi.org/10.1109/BigData62323.2024.10825937>. doi:10.1109/BIGDATA62323.2024.10825937.

- [13] H. Qi, X. Li, C. Zhang, T. Zhao, Improving drug-drug interaction prediction via in-context learning and judging with large language models, *Frontiers in Pharmacology* 16 (2025) 1589788.
- [14] Y. Lee, Enhancing medication recommendation with LLM text representation, *CoRR* abs/2407.10453 (2024). URL: <https://doi.org/10.48550/arXiv.2407.10453>. doi:10.48550/ARXIV.2407.10453. arXiv:2407.10453.
- [15] W. T. Kim, J. Shin, I.-S. Yoo, J.-W. Lee, H. J. Jeon, H.-S. Yoo, Y. Kim, J.-M. Jo, S. Hwang, W.-J. Lee, et al., Medication extraction and drug interaction chatbot: Generative pretrained transformer-powered chatbot for drug-drug interaction, *Mayo Clinic Proceedings: Digital Health* 2 (2024) 611–619.
- [16] A. Romano, G. Riccio, M. Postiglione, V. Moscato, Pie-med: Predicting, interpreting and explaining medical recommendations, in: C. Hauff, C. Macdonald, D. Jannach, G. Kazai, F. M. Nardini, F. Pinelli, F. Silvestri, N. Tonellotto (Eds.), *Advances in Information Retrieval - 47th European Conference on Information Retrieval, ECIR 2025, Lucca, Italy, April 6-10, 2025, Proceedings, Part V*, volume 15576 of *Lecture Notes in Computer Science*, Springer, 2025, pp. 6–12. URL: https://doi.org/10.1007/978-3-031-88720-8_2. doi:10.1007/978-3-031-88720-8_2.
- [17] A. Ferraro, A. Galli, V. La Gatta, M. Postiglione, D. Russo, G. M. Orlando, G. Riccio, A. Romano, V. Moscato, From explanation to exploration: Promoting diversity in recommendation systems, in: *International Workshop on Recommender Systems for Sustainability and Social Good*, Springer, 2024, pp. 135–150.
- [18] G. Riccio, A. Romano, A. Korsun, M. Cirillo, M. Postiglione, V. L. Gatta, A. Ferraro, A. Galli, V. Moscato, Healthcare data summarization via medical entity recognition and generative AI, in: N. Bena, B. D. Martino, A. Maratea, A. Sperduti, E. D. Nardo, A. Ciaramella, R. Montella, C. A. Ardagna (Eds.), *Proceedings of the 2nd Italian Conference on Big Data and Data Science (ITADATA 2023)*, Naples, Italy, September 11-13, 2023, volume 3606 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2023, pp. –. URL: <https://ceur-ws.org/Vol-3606/paper47.pdf>.
- [19] C. Knox, M. Wilson, C. M. Klinger, M. Franklin, E. Oler, A. Wilson, A. Pon, J. Cox, N. E. Chin, S. A. Strawbridge, et al., Drugbank 6.0: the drugbank knowledgebase for 2024, *Nucleic acids research* 52 (2024) D1265–D1275.
- [20] R. Manna, M. P. di Buono, L. Giordano, DIMMI - drug information mining in italian: A CALAMITA challenge, in: F. Dell’Orletta, A. Lenci, S. Montemagni, R. Sprugnoli (Eds.), *Proceedings of the Tenth Italian Conference on Computational Linguistics (CLiC-it 2024)*, Pisa, Italy, December 4-6, 2024, volume 3878 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2024, pp. –. URL: https://ceur-ws.org/Vol-3878/126_calamita_long.pdf.
- [21] S. Ayvaz, J. R. Horn, O. Hassanzadeh, Q. Zhu, J. Stan, N. P. Tatonetti, S. Vilar, M. Brochhausen, M. Samwald, M. Rastegar-Mojarad, M. Dumontier, R. D. Boyce, Toward a complete dataset of drug-drug interaction information from publicly available sources, *J. Biomed. Informatics* 55 (2015) 206–217. URL: <https://doi.org/10.1016/j.jbi.2015.04.006>. doi:10.1016/J.JBI.2015.04.006.
- [22] R. Ruggiero, F. Fraenza, C. Scavone, G. di Mauro, R. Piscitelli, A. Mascolo, C. Ferrajolo, C. Rafaniello, L. Sportiello, F. Rossi, et al., Immune checkpoint inhibitors and immune-related adverse drug reactions: Data from italian pharmacovigilance database, *Frontiers in Pharmacology* 11 (2020) 830.
- [23] L. Magro, E. Arzenton, R. Leone, M. G. Stano, M. Vezzaro, A. Rudolph, I. Castagna, U. Moretti, Identifying and characterizing serious adverse drug reactions associated with drug-drug interactions in a spontaneous reporting database, *Frontiers in Pharmacology* 11 (2021) 622862.
- [24] G. Trifirò, V. Isgrò, Y. Ingrassiotta, V. Ientile, L. L’Abbate, S. S. Foti, V. Belleudi, F. Poggi, A. Fontana, U. Moretti, et al., Large-scale postmarketing surveillance of biological drugs for immune-mediated inflammatory diseases through an italian distributed multi-database healthcare network: the valore project, *BioDrugs* 35 (2021) 749–764.
- [25] E. Galai, L. Scotti, M. Gilardetti, A. Ucciero, D. Ferrante, E. Poluzzi, A. A. Genazzani, F. Barone-Adesi, Time-trends of drug-drug interactions among elderly outpatients in the piedmont region (italy): a population-based study, *International Journal of Environmental Research and Public Health* 19 (2022) 7353.
- [26] R. Leone, L. Magro, U. Moretti, P. Cutroneo, M. Moschini, D. Motola, M. Tuccori, A. Conforti, Identifying adverse drug reactions associated with drug-drug interactions: data mining of a

spontaneous reporting database in italy, *Drug safety* 33 (2010) 667–675.

[27] Z. Shen, M. Spruit, Automatic extraction of adverse drug reactions from summary of product characteristics, *Applied Sciences* 11 (2021) 2663.

[28] A. . M. Llama Team, The llama 3 herd of models, *CoRR* abs/2407.21783 (2024). URL: <https://doi.org/10.48550/arXiv.2407.21783>. doi:10.48550/ARXIV.2407.21783. arXiv:2407.21783.