OntoDrug: Enhancing Brazilian Health System Interoperability with a National Medication Ontology

Nelson Miranda
nelson.miranda@usp.br
Institute of Mathematical and
Computer Sciences, USP
São Carlos, São Paulo, Brazil

Matheus Matos Machado
matheusmatos@usp.br
Institute of Mathematical and
Computer Sciences, USP
São Carlos, São Paulo, Brazil

Dilvan A. Moreira
dilvan@usp.br
Institute of Mathematical and
Computer Sciences, USP
São Carlos, São Paulo, Brazil

ABSTRACT

This paper presents OntoDrug, an ontology designed to enhance medicine management in Brazil by integrating regulatory frameworks and standardizing terminologies. OntoDrug improves patient safety and treatment efficacy by accurately identifying and classifying medications and supporting interoperability with health information systems. A proof-of-concept application integrated into the Hospital das Clínicas de Marília's hospital EHR system demonstrated OntoDrug's utility, achieving high precision and recall. An experimental study using large language models grounded on the ontology achieved, using GPT-4 turbo, 0.97 precision, 1.0 recall and an F1-score of 0.99. We also evaluated open-source models llama3-8b, llama3-70b, and gemma-7b-it. Their performance was close to GPT-4's. The significant effectiveness is primarily due to the utilization of large language models (LLMs). While using these large language models enhanced performance, challenges related to cost, privacy, and service availability were identified. OntoDrug represents a significant advancement in Brazil's medication information standardization and optimization.

KEYWORDS

Medication Ontologies, Drug Management, Semantic Interoperability, Health Informatics

1 INTRODUCTION

Medication management is a critical component of healthcare, encompassing the entire lifecycle of a patient's medications, from procurement to discontinuation [33]. This cycle includes accurate transcription of prescriptions, tracking of medication lots and expiration dates, precise dispensation, and detailed administration. Additionally, continuous monitoring and adherence checks are conducted to optimize therapeutic outcomes. Regular reviews are performed to adjust medication plans according to changing patient needs, leading to appropriate discontinuation of medications. This systematic approach is essential for enhancing patient safety, treatment efficacy, and optimizing healthcare resources [17].

Despite its importance, the medication management process is fraught with challenges. One significant issue is the lack of standardization and controlled vocabularies, leading to ambiguities in the nomenclature used throughout the cycle. Such ambiguities can result in communication errors between prescribers, dispensers,

and patients, ultimately compromising patient safety and treatment effectiveness. Therefore, a controlled vocabulary is crucial in normalizing terms for pharmaceutical forms, administration routes, packaging, and measurement units, ensuring uniform communication and reducing errors [13]. Such harmonization can reduce errors and ambiguity in medication prescription and administration, fostering patient safety and effective treatment. [43] Furthermore, the system supports interoperability with other health information systems, making data management and exchange activities more effective.

In Brazil, managing medications is challenging due to the presence of different regulatory frameworks and lists such as the DCB (Brazilian Common Denominations), CMED (Chamber of Medicines Market Regulation), and RENAME (National List of Essential Medicines). Prescribing drugs from RENAME is vital as it guarantees patients access to treatments that are both effective and affordable, supporting rational drug use in the healthcare system[11].These lists standardize drug names, regulate the pharmaceutical market, and detail essential medications for public health, respectively [21]. However, integrating these regulatory frameworks into a cohesive system remains a challenge, often resulting in inconsistent medication identification and classification.

The characteristics of medications and substances dictate their classification and prescription restrictions. Categories include generic drugs, similar drugs, biological drugs, and biosimilars, each with specific prescription levels indicated by colored stripes on their packaging [2]. Additionally, regulatory lists categorize controlled substances such as narcotics, psychotropic drugs, anorexigenic substances, immunosuppressants, antiretrovirals, anabolic steroids, and precursors used in drug manufacturing. These classifications ensure responsible prescription and dispensation practices, adhering to health and legal standards to manage risks associated with their use [12].

The primary objectives of this research were to develop and implement the OntoDrug ontology to enhance medication management within the Brazilian healthcare system.

The OntoDrug ontology aims to address these challenges by providing a framework that integrates Brazil's regulatory lists into a standardized, computer-readable format. By leveraging controlled vocabulary and coding systems like the CAS-Anatomical Therapeutic Chemical Classification System and RxNorm, OntoDrug ensures accurate identification and disambiguation of medications [26]. This ontology enhances the organization and accessibility of pharmacological knowledge and also supports interoperability with other health information systems, improving data management and exchange activities [9].

Nelson Miranda, Matheus Matos Machado, and Dilvan A. Moreira

To demonstrate the utility of OntoDrug, a proof-of-concept application was developed and integrated into the EHR of the Hospital das Clínicas de Marília hospital. Whenever a prescription is created in the EHR system, this application reads it, identifies the medications, and warns physicians if some medication is not listed in the Rename, suggesting alternatives whenever possible. The ontology is used to ground the LLMs doing the medicine recognition and to suggest replacements for the ones not in Rename. When recognizing medicines, its best results (when using the GPT-4 turbo LLM) were 1.0 recall, 0.974 precision, and 0.987 F1-score. The physicians were also happy with the replacement suggestions.

In the following sections, we discuss the development of the OntoDrug ontology, its implementation in a clinical setting, and its potential applications in improving medication management practices.

2 RELATED WORKS

In our exploration of medication ontologies, we reviewed several key works that focus on the creation and utilization of drug ontologies. Peña (2020) [32] developed a Drug Ontology for the Mexican public health system to provide a well-structured medical knowledge base accessible to various stakeholders. Based on the "Basic Table and Catalog of Medicines" from the Secretary of Public Health, the ontology consists of 64 classes, 5 object properties, and 18 data properties. Its evaluation, focusing on model competence and quality criteria, demonstrated its potential for enhancing medical knowledge management.

A unified Brazilian drug database is meant to advance patient safety and individualized care. Queiroz (2023) [34] emphasizes the aspect of terminology standardization and semantic alignment related to medication nomenclature, underlying how this system would be advantageous for healthcare providers in securing proper, consistent, and safe pharmaceutical care for patients.

Avila (2021) [6] developed MediBot, a chatbot that provides drug information, including prices and potential substitutes. The bot works in two modes: Quick Response mode and Interactive mode. It uses a Linked Data Mashup to extract real-time information from the web. This system could be helpful for consumers and healthcare professionals seeking accurate and up-to-date information related to drugs.

DINTO[1] is a drug-drug interaction-focused ontology along with mechanisms concerning the improvement of clinical safety [19]. It efficiently classifies interactions and mechanisms, facilitating large-scale prediction of potential interactions [20]. This helps detect severe adverse reactions and safety issues mentioned under Warnings and Precautions on the labels of drugs. Moreover, DINTO supports the generation of new hypotheses on drug interaction and can be helpful during testing at the new drug's preclinical stage.

Sharp (2017) [37] and Hanna (2013) [18] contribute to the development of a comprehensive drug ontology. Sharp's work focuses on the extraction of drug-indication relations from various sources, while Hanna's work builds the Drug Ontology (DrOn)[2] based on RxNorm and other sources. Li 2019 and Herrero-Zazo 2013 further

enhance the ontology by designing a drug-repurposing-oriented Alzheimer's Disease Ontology and an ontology for drug-drug interactions, respectively. These works collectively contribute to the creation of a comprehensive drug ontology that covers various aspects of drugs, including their properties, classifications, and relationships.

The OntoDrug ontology aligns with Brazilian regulations like RENAME and CMED, distinguishing it from the broader Mexican Drug Ontology and the specialized MediBot. Unlike the Brazilian Medicines Ontology (OBM), which focuses on unifying drug nomenclature for safety, OntoDrug includes detailed drug data and regulatory information. It can enhance medication safety and be used for educational purposes. Additionally, while DINTO focuses on drug-drug interactions and DRON on drug properties, OntoDrug's specific adaptation to Brazilian standards makes it a tool for healthcare compliance and education.

3 METHOD DESCRIPTION

The OntoDrug ontology integrates Brazil's DCB, CMED, and RENAME listings into a standardized, computer-readable format. It reuses the Schema.org[3] ontology to enhance interoperability across computational systems, including medical suppliers, and to structure medical information effectively[15, 16, 23, 28, 29, 35]. The Schema.org defines a widely used formal ontology for web page semantic markup, employing a collection of interconnected schemas (classes) and properties (attributes) collaboratively developed by major search engines like Google and Microsoft to standardize the representation of web content entities, including medicines. [1]

We adopted Schema.org types such as *Thing*, *MedicalEntity*, *Drug*, and *DrugClass*, facilitating a systematic categorization and enriching the ontology with entities like *DrugCost* and *MedicalCode*. We reused the *DrugStrength* class for generic medicines but changed its name (label) to *GenericMedicine* for clarity (but kept its original ID). Reusing schema terms not only aligns OntoDrug with global standards but also supports detailed descriptions of medicines. It promotes a standardized representation compatible with semantic web content about medications on thousands of sites, enhancing the ontology's functionality within web-based health information systems.

For the development of the OntoDrug ontology, we have based our work on the structured methodology of Noy and McGuinness [29] for ontology development. We used the Web Ontology Language (OWL) [5], which is a W3C standard for the representation of ontologies, and the free, open-source Protégé ontology editor [27] [10].

Our ontology development targeted the medication domain, aiming to streamline drug identification by active ingredient or commercial name and to categorize drugs based on their listing in RENAME. We utilized OpenRefine 3.6.2, an open-source tool, for executing ETL (Extract, Transform, Load) processes on the data sourced from CMED, RENAME, and DCB tables. Data were extracted and then imported into OpenRefine for necessary transformations. OpenRefine was used to manage unstructured data, offering robust functionalities for data cleaning, transformation, and reconciliation [41].

---

[1]DINTO ontology details can be found at https://portal.bioontology.org/ontologies/DINTO

[2]DRON ontology details available at https://bioportal.bioontology.org/ontologies/DRON

[3]https://schema.org

Table 1: Details of CMED, DCB, and RENAME Lists

List	Records	Comments
CMED	26,062	Medications Marketed in Brazil
DCB	12,459	Active Ingredients
Rename	921	Essential Medications

The transformed RENAME, CMED, and DCB tables encapsulated essential drug information. The RENAME data, initially in PDF format, was manually converted into CSV to facilitate the detailing of generic names, concentrations, pharmaceutical forms, components, and ATC codes. The CMED list encompasses detailed pharmaceutical data, including active ingredients, manufacturers, registration codes, presentations (concentration, dosage form, packaging, quantity), therapeutic classes, product types (reference, similar, or generic), pricing, and regulatory categories by colored stripes (black, red, or none). Additionally, the DCB table enriched the ontology with detailed chemical data, including the substance name (active ingredient), its Chemical Abstracts Service (CAS) Number, and classifications such as API (Active Pharmaceutical Ingredient), Biological, or Homeopathic, enhancing the ontology's utility for diverse pharmaceutical applications [44]. For details on the three lists, see Table 1 for a summary of records and associated comments. Figure 1 displays the high-level structure of the ontology, highlighting its primary classes. The central class, *GenericMedicine*, represents the specific dosage at which a drug is available, combining active ingredients (*activeIngredient* property) with its concentration and unit of measure (*strengthUnit* and *strengthValue* properties), based on data from the CMED list. This list offers a comprehensive catalog of medications along with their attributes [7]. The *Drug* class corresponds to the active ingredients and is derived from the DCB listing.

Medications listed in RENAME are categorized under the *GenericMedicine* class. They correspond to entries found in the CMED list and include additional details such as the pharmaceutical care group(*RenameGroup*) to which each medication belongs, whether basic, strategic, or specialized [8].

The *CommercialMedicine* class describes commercial drug products from private laboratories. These products are linked to the *GenericMedicine* class and, through it, to the *Drug* class, allowing for a connection based on shared attributes like active ingredient, dosage form, and concentration. This linkage enables multiple products to be associated with a single *GenericMedicine* if they have identical properties. Such a structure guarantees that medications, identified by their commercial name or active ingredient, are consistently represented in the ontology, enhancing its practicality.

All medications in the *GenericMedicine* class listed in RENAME are designated as belonging to the *Rename* class. To verify RENAME membership, we determine if a medication is a *GenericMedicine* of type *Rename* or a *CommercialMedicine* equivalent to such a *GenericMedicine*. If there is no direct match, we check whether the medication belongs to a *CommercialMedicine* or *GenericMedicine* with an active component in RENAME. We then gather all *GenericMedicine* entries with this active ingredient to see if alternative combinations
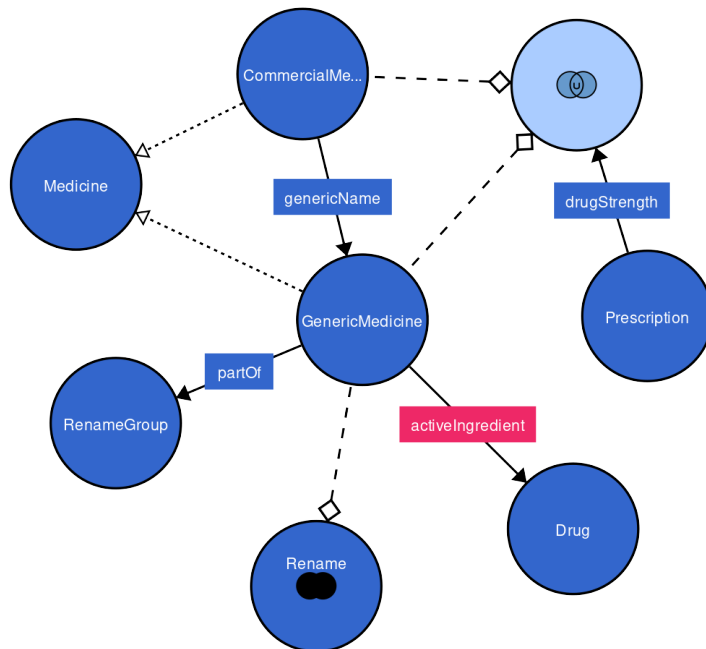
Figure 1: OntoDrug: Overview of the Ontology Structure.

match the required strength. For example, if *Ibu-tab 400 mg* is prescribed and unavailable, the system might suggest two doses of the generic medicine *Ibuprofen 200 mg* as an alternative.

The controlled vocabulary detailing pharmaceutical forms, routes of administration, and packaging, as published by Anvisa [3], was successfully and manually transformed into an RDF taxonomy for integration with the OntoDrug ontology. This transformation included:

- **Pharmaceutical Forms**: Solids, liquids, semi-solids, and gases were identified as categories, each containing specific entries under them, as in the case of tablets, capsules, creams, and inhalants.
- **Routes of Administration**: This consisted of classifications under mechanisms like oral, intravenous, topic, and inhalation.
- **Packaging**: It listed all kinds of packaging forms under this category as ampoules, bottles, blister packs, and tubes, which were grouped under categories of primary, secondary, and special packaging.

In RDF, each term from the controlled vocabulary was represented as an instance within its respective class, with properties to denote relationships and characteristics, enhancing interoperability and standardization across healthcare systems.

In developing the OntoDrug ontology, we included the Units of Measurement Ontology (UO) to standardize measurements like mg, g, mcg, mg/mL, and ml, ensuring clarity in medication properties across health systems. The UO facilitates data and knowledge integration, enabling interoperability and semantic information processing across various biomedical resources and domains. [31] This integration supports precise drug administration and enhances

OntoDrug's interoperability with health informatics systems[4]. The ontology data were exported in RDF format using the RDF extension in OpenRefine and stored in GraphDB 10.0.1[5], a triplestore. A triplestore, or RDF store, is a database specifically designed for storing and retrieving data via semantic queries, which facilitates the handling and querying of ontology data.[36]

4 ONTODRUG ONTOLOGY

This section presents the ontology structure, such as the class hierarchy, primary axioms, and retrieval of individuals. OntoDrug contains about 33 classes, 18 object properties, 19 data properties, and over 78,000 individuals. It includes around 430 axioms that make questioning and reasoning capabilities possible.

Structured relationships between key classes, such as *GenericMedicine*, *Drug*, and *CommercialMedicine*, are shown in Figure 2. In this regard, having a diagram is important as a way to make sure that different elements relate in the proper way to one another.
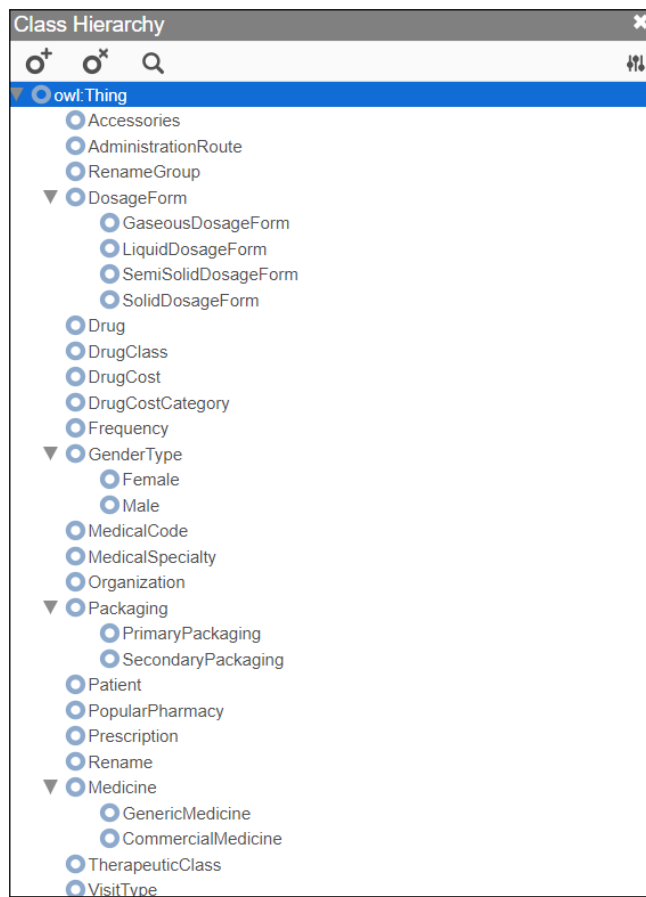

Figure 2: Class Hierarchy of the OntoDrug Ontology

In our work, a significant emphasis was placed on the integration of instances from diverse data sources (DCB, CMED, and RENAME) into the ontology. For instance, we created an axiom to determine

if a medicine is part of Rename. The formal definition specifies that within the entire set of medicines, only those classified under the GenericMedicine class are considered part of the Rename list. This classification is established through a relationship indicated by the property *:categorizedAs* linking to an instance of the RenameGroup class. Additionally, a commercial medicine can be indirectly included in the Rename list because it has an equivalent generic. This inclusion occurs through the association of its generic name (*:genericName* property) with an entity within the RenameGroup. The RenameGroup class is an enumerated type comprising 3 elements, rename-basic, renames-strategic, and rename-specialized, each representing a different categorization within the Rename list. This axiom is represented in Equation 1. It is key to the correct classification of medicines as belonging to the Rename list.

$$
\begin{aligned}
\text{Rename} \equiv\ & (\text{CommercialMedicine} \wedge \\
& (\exists \text{:genericName}\,(\text{GenericMedicine} \wedge \\
& (\exists \text{categorizedAs.RenameGroup})))) \\
& \vee (\text{GenericMedicine} \wedge \\
& (\exists \text{categorizedAs.RenameGroup})) \qquad (1)
\end{aligned}
$$

Figure 3 shows a part of the OntoDrug ontology related to the active ingredient *Ivermectin*. In the knowledge graph, products containing different strengths of ivermectin are related, such as the generic medicine *Ivermectin 6 mg* and the commercial *Revectina 6mg*. From this visualization, one can now realize that *Soolantra 10mg/g* is not part of the RENAME list as it is related to the generic medicine Ivermectin 10mg/g, which is not part of the Rename list.
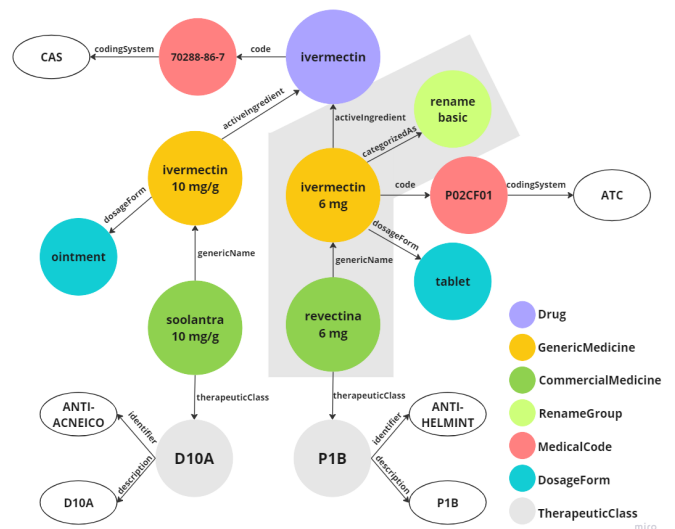

Figure 3: Knowledge Graph Generated from OntoDrug

5 USE CASE SCENARIOS

The application of the OntoDrug ontology can extend beyond its integration with electronic health records (EHRs). This section outlines some real-world scenarios where OntoDrug can enhance

---

[4]Units of Measurement Ontology details available at https://ontobee.org/ontology/UO
[5]https://graphdb.ontotext.com/documentation/10.0/index.html

pharmaceutical practices and healthcare services. The ontology can aid in the customization of drug therapies, streamline medicine stock management, and serve as an educational resource. Each scenario demonstrates the practical benefits and potential of OntoDrug to improve patient care, optimize operational efficiencies, and support the continuous education of healthcare professionals.

5.1 Personalized Medication Therapy

Integrating OntoDrug into a clinical decision support system allows medication therapies to be personalized based on the individual specifics of each patient. It could document drug interactions with allergies, other drugs, and pre-existing conditions. Such systems can use the detailed and structured information available in OntoDrug to help identify problems and recommend safer or more efficacious medications. This approach helps to optimize patient safety and therapeutic outcomes by ensuring that suggestions of optimal medication options are made available. Options more suitable to the unique health profile of each patient.

5.2 Medication Inventory Management

The OntoDrug ontology can also be leveraged to optimize inventory management in pharmacies and hospitals by offering possible options for medicines that are unavailable in the Rename list.

5.3 Pharmacy Education and Training

The OntoDrug ontology can serve as an educational tool for students and healthcare professionals, offering a comprehensive knowledge base on drug characteristics, classifications, and proper usage. This resource can be integrated into e-learning modules that simulate prescription writing and drug administration scenarios. Such training enhances learning through interaction with an ontology-based system, reinforcing the application of theoretical knowledge in real-world settings, leading to better-informed clinical decisions

The scenarios discussed illustrate the versatility and possible uses of the OntoDrug ontology across various facets of healthcare and education. The next section, *Proof of Concept*, shows a practical use of the OntoDrug ontology. It was integrated into the electronic prescribing module of a major hospital's EHR system.

6 PROOF OF CONCEPT

This section reports the results of integrating the **OntoDrug** ontology into the prescription module of the Marília Medical School's (FAMEMA) Hospital **EHR system**. This real-world deployment showcases how the ontology can promote adherence to a prescription standard, enhancing the efficiency of medication management in a clinical setting.

As shown in Figure 4, the prescription module GUI incorporates a user-friendly interface for entering medical prescriptions, which allows healthcare personnel to easily create free-text prescriptions.

This functionality is provided by the **Prescription Service**, implemented using .NET 4.5, which enhances the capabilities of the hospital's existing EHR system. It serves as an intermediary that orchestrates the data exchange between the existing EHR and a prescription database through the use of JSON objects (Listing 1). This

Figure 4: Proof of Concept: Medication Order Form

service ensures robust data integrity and secure communication [24][38].

```
1  {
2    "prescription": "INTERNAL USE new prescription:
3    1) Riluzole 70 mg Take 01 tablet every 12/12 hours.
4    2) Methyldopa 250 mg Take 02 tablets every 8/8 hours.
5    3) Carbamazepine 200 mg Take 02 tablets every 6/6 hours.
6    4) L-carnitine 1g continuous use Take 03 tablets every 12/12
       hours.
7    5) Omalizumab 150 mg continuous use Take 01 tablet once a day
8    6) Simvastatin 40 mg continuous use Take 10 tablets 10 times a
       day",
9    "cns": "001",
10   "name": "patient name"
11 }
```

Listing 1: Translated JSON for Medicine Prescription

Data interchange between the EHR and the Prescription Service involves object serialization, which is managed by the **Prescription Backend**. This backend handles business logic, data transformations, and interactions with the ontology [39][22].

In the workflow of the Prescription Backend, an **Entity Extraction Service** utilized a traditional pre-trained model specifically tailored for Named Entity Recognition (NER) in medical prescriptions. We tested the quality of the NER extraction by asking physicians how many medications were successfully recognized in each prescription (Figure 5). We had feedback for 3,761 prescription reports. NER achieved good precision with 61% or 2,295 prescriptions where all medicines were successfully recognized and 33.5% or 1,261 prescriptions where at least one medicine was recognized [25]. This is the current version in production.

After recognizing medicines in the prescription, the system uses the ontology, employing the techniques explained earlier, to verify if these medicines are part of Rename and, if they are not, try to find a medicine or combination of medicines that can replace the original one. In this way, the system can not only warn physicians about non-compliant medicines but also suggest, if possible, substitutes.

The system uses SPARQL queries to a triple store that houses the ontology to interact with it. This arrangement permits sophisticated handling of queries related to medication, encompassing drug classifications and potential alternatives based on active ingredients.

The Prescription Backend constructs and dispatches SPARQL queries to the GraphDB triple store, which executes them and returns results in JSON format. Listing 2 shows a sample JSON response from the backend. The **valid** section lists medications successfully verified.

```
1  {
2      "valid": [
3          {
4              "name": "Carbamazepine",
5              "dosage": "200 mg",
6              "rename": true,
7              "atc": "N03AF01"
8          },
9          "...other valid items..."
10     ],
11     "invalid": [
12         {
13             "originalname": "Riluzole 70 mg",
14             "name": "riluzole",
15             "dosage": "50 mg",
16             "rename": "N",
17         }
18     ]
19 }
```

Listing 2: JSON Representation for Medication Validation

Figure 5 displays the user interface, which presents a list of medications identified by the ontology along with information on whether they belong to the Rename list. Additionally, the interface requests physicians' feedback to assess the information's accuracy and relevance.
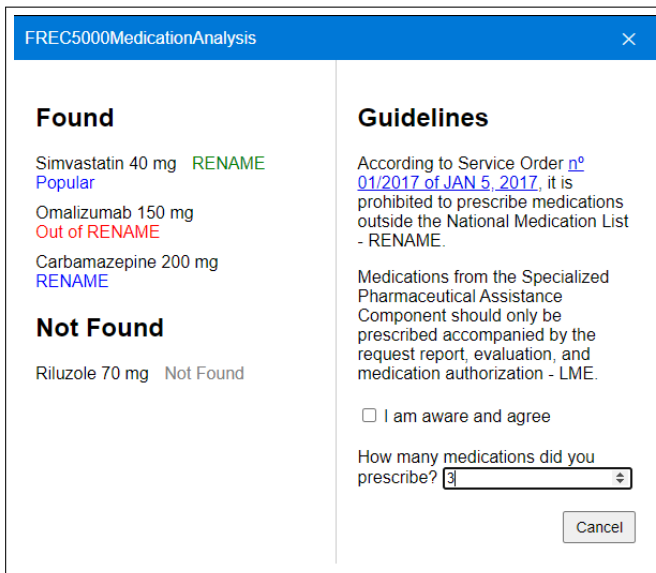

Figure 5: Feedback Dialog: Display of Serialized Data in the User Interface

6.1 Using LLMs for medicine recognition

The last section described the system currently in production at the FAMEMA school hospital. Subsequently, the recognition of medicines on the prescription based on traditional NER techniques was replaced by more advanced LLM-based implementations, such as LLM GPT-4 Turbo, which enhanced natural language processing capabilities. The rest of the system remained the same.

Figure 6 shows the data transmission from the Prescription Service to the EHR system, offering a succinct visualization of the new workflow, which now incorporates the use of LLMs for enhanced processing.
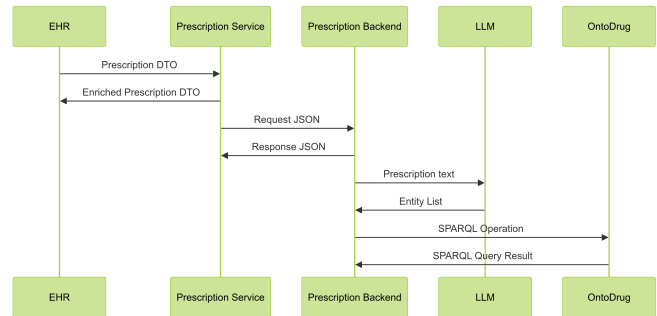
Figure 6: EHR to OntoDrug Data Flow Overview

We could have asked the LLM to directly determine if the recognized medicines were or were not listed in the Rename. However, that could lead to hallucinations [42]: the LLM falsely stating that a medicine is part of Rename or, worst still, suggesting nonexistent substitutes. We decided to use the ontology to "ground" the LLM [14]. We limited the LLM to medicine recognition and left the rest to the ontology. In this way, all possible answers are limited by the ontology contents, reducing hallucinations. Notice that the LLM can still make a mistake and not recognize a medicine, but the system will not misclassify a correctly recognized medicine or suggest nonexistent medicine substitutions.

For this new implementation, we tested different LLMs, remembering that we still use the OntoDrug to find a medication Rename membership and to suggest alternative medications whenever possible. The GPT-4 Turbo model achieved perfect recall (1.0) and high precision (0.97), resulting in an overall F1-score of 0.99, showing highly effective performance metrics.

7 EXPERIMENTAL STUDY AND RESULTS

We evaluated the new implementation across four different LLMs: GPT-4 Turbo, LLaMA3-8b, LLaMA3-70b, and GEMMA-7b-IT to assess their effectiveness [4, 30, 40]. One goal was to compare these models to the original NER approach and determine their relative performance.

In the original approach (NER + ontology), we used 3,761 prescriptions obtained with physician feedback provided in response to the dialog box shown in Figure 5. The results for this approach [25] were:

(1) **Complete Recognition** (61% or 2,295 Prescriptions): The system accurately identified all medications as reported by the physicians in these instances.

(2) **Partial Recognition** (33.5% or 1,261 Prescriptions): The system recognized only some of the medications reported by the doctors.

(3) **No Recognition** (5.5% or 205 Prescriptions): The system did not identify any medications in these cases, indicating a need for further investigation.

These results highlighted the possibility that better techniques could enhance the system's accuracy.

In our current study, we replaced the traditional NER-based recognition with LLMs to recognize medications and dosages directly from the prescription texts.

We conducted tests with four models. A massive, more than 300 billion parameters, proprietary top model, gpt-4 Turbo, and three open source smaller models, llama3-70b (70 billion parameters), llama3-8b (8 billion), and gemma-7b-it (7 billion). We evaluated their effectiveness and accuracy in comparison to the original NER-based approach.

For each prescription, we used the list of the medicines actually in the prescription, supplied by a physician, and the list of machine-detected medicines. All medicines in both lists were considered true positives (TP), medicines only on the physician's list were false negatives (FN), and only in the machine-detected list were false positives (FP). There are no true negatives (TN).

We could not use all the original 3,761 prescriptions (from the first study). For simplicity of the GUI, the physicians only informed the total number of medicines prescribed and how many were correctly recognized in each prescription. We do not know which medicines were correctly identified. For that, physicians needed to go on each prescription and manually annotate which medicines were present.

Another problem is price restrictions. Currently, high-end, very expensive hardware is needed to run LLMs locally. So, it is more cost-effective to run them remotely using paid systems. That option is cheaper but still quite expensive to run 3,761 examples. So, we opted to use a random sample of 203 (5.4%) prescriptions from the original set.

For the experiments, gpt-4 Turbo models run using OpenAI's cloud [6] and llama3-70b, llama3-8b, and gemma-7b models used the Groq Cloud[7].

After each of the four models recognized the medicines in the prescriptions, we calculated the precision, recall, and F1-score shown in Table 2. We did not include measurements, such as accuracy, that depend on true negative (TN) values, as these values are not available.

Model	Precision	Recall	F1-score
gpt-4 turbo	0.974	1.000	0.987
llama3-70b-8192	0.990	0.985	0.988
gemma-7b-it	0.995	0.971	0.983
llama3-8b-8192	0.985	0.973	0.979
ner-model*	*0.976*	*0.501*	*0.662*

*(*Included for comparison purposes.)*

Table 2: Metrics of different language models.

As expected, gpt-4 turbo had the best numbers, even achieving 100% recall. But the three open source models were not far behind

[6]https://openai.com/
[7]https://groq.com/

it. For instance, all four models achieved a precision above 98%. In general, the bigger the model, the better (Figure 7). However, Gemma-7b-it was slightly better than llama3-8b.

It is important to highlight that the smaller LLMs had values very competitive in relation to GPT 4 performance. That indicates that medicine identification is not such a hard problem for LLMs, opening the possibility that sooner such models can run locally, reducing costs and improving privacy.
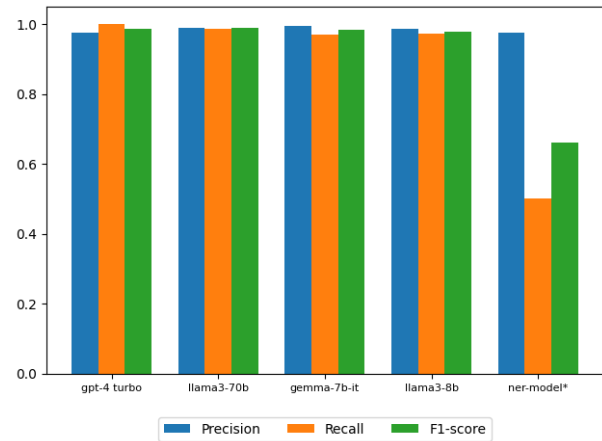
Figure 7: Performance metrics of different language models.

All models outperformed the pre-trained NER model, which achieved a Recall of 0.501, a Precision of 0.976, and an F1-score of 0.662 on the same sample.

Despite the significant improvement in accuracy with the LLM use, they have disadvantages. One primary concern is the cost; running such advanced models can be expensive, especially when processing large datasets, making it less feasible for continuous or large-scale applications (a hospital may produce thousands of prescriptions daily).

Additionally, service availability can be an issue; reliance on an external API means that any downtime or service disruption from the provider or the internet could impact the accessibility and functionality of the medication identification system. Privacy is another issue because the sensitive data on patients is to be transferred and processed by the API provider.

8 CONCLUSION

In this work, we presented the development and implementation of the OntoDrug ontology, designed to enhance medication management within the Brazilian healthcare system. By integrating regulatory frameworks and standardizing terminologies, OntoDrug supports accurate identification and classification of medications, reducing ambiguities and improving patient safety. Our proof-of-concept application demonstrated the practical utility of OntoDrug in a clinical setting, showing significant improvements in the precision and efficiency of medication management.

The experimental study revealed that the incorporation of advanced language models like GPT-4 turbo significantly outperforms traditional NER methods in terms of recall, precision, and F1-score.

However, these models provide substantial accuracy gains but come with cost, privacy, and service availability challenges.

Our findings highlight the potential of OntoDrug to be a valuable tool for various stakeholders, including healthcare providers, pharmacists, and educators. OntoDrug can enhance clinical decision-making by seamlessly integrating medication information into electronic health records.

Future research should focus on expanding the prescription validation process to include warnings of possible harmful interactions between prescribed drugs. Additionally, exploring the use of local open-source LLMs for entity recognition, as they can be more cost-effective and privacy-preserving implementations of language models.

9   ACKNOWLEDGMENTS

REFERENCES

[1] Burhan Ud Din Abbasi, Iram Fatima, Hamid Mukhtar, Sharifullah Khan, Abdulaziz Alhumam, and Hafiz Farooq Ahmad. 2022. Autonomous schema markups based on intelligent computing for search engine optimization. *PeerJ Computer Science* 8 (2022), e1163.

[2] Agência Nacional de Vigilância Sanitária. 2004. *Política Vigente para a Regulamentação de Medicamentos no Brasil.* Gabinete do Diretor-Presidente, Núcleo de Assessoramento em Comunicação Social e Institucional - Comin/Anvisa. https://bvsms. saude.gov.br/bvs/publicacoes/anvisa/manual_politica_medicamentos.pdf Impresso no Brasil. Permitida a reprodução parcial ou total desta obra, desde que citada a fonte..

[3] Agência Nacional de Vigilância Sanitária (Anvisa). 2022. *Vocabulário Controlado de Formas Farmacêuticas, Vias de Administração e Embalagens.* https://www.gov. br/anvisa/pt-br/centraisdeconteudo/publicacoes/medicamentos/publicacoes-sobre-medicamentos/vocabulario-controlado.pdf/view Accessed: 2022-05-19.

[4] AI@Meta. 2024. Llama 3 Model Card. (2024). https://github.com/meta-llama/ llama3/blob/main/MODEL_CARD.md

[5] Dean Allemang and James Hendler. 2011. *Semantic web for the working ontologist: effective modeling in RDFS and OWL.* Elsevier.

[6] Caio Viktor S Avila, Wellington Franco, Amanda DP Venceslau, Tulio Vidal Rolim, Vania MP Vidal, and Valéria M Pequeno. 2021. MediBot: an ontology-based chatbot to retrieve drug information and compare its prices. *Journal of Information and Data Management* 12, 2 (2021).

[7] Letícia de Andrade Barbosa. 2013. Assistência farmacêutica no sistema único de saúde. (2013).

[8] Ricardo José Magalhães Barros, Marco Antônio de Araújo Fireman, and Mateus Rodrigues Westin. 2017. Relação nacional de medicamentos essenciais. (2017).

[9] Erik Bülow. 2020. coder: An R package for code-based item classification and categorization. *J. Open Source Softw.* 5 (2020), 2916. https://api.semanticscholar. org/CorpusID:233302997

[10] Wander de Almeida Limeira and Dilvan de Abreu Moreira. 2020. KGWE-A customizable editor for OWL ontologies. In *Anais Estendidos do XXVI Simpósio Brasileiro de Sistemas Multimídia e Web.* SBC, 91–95.

[11] Daniela Moulin Maciel de Vasconcelos, Gabriela Costa Chaves, Thiago Botelho Azeredo, and Rondineli Mendes da Silva. 2017. National Medicines Policy in retrospective : a review of (almost) 20 years of implementation. https://api. semanticscholar.org/CorpusID:22550643

[12] Agência Nacional de Vigilância Sanitária (Anvisa). 2020. *Novo Marco Regulatório de Farmacovigilância.* https://www.gov.br/anvisa/pt-br/assuntos/noticias-anvisa/ 2020/novo-marco-regulatorio-de-farmacovigilancia-confira Accessed: 2024-06-04.

[13] Frederico Freitas and Stefan Schulz. 2009. Ontologias, Web semântica e saúde. *Revista Eletrônica de Comunicação, Informação & Inovação em Saúde* 3, 1 (2009).

[14] Maurice Funk, Simon Hosemann, Jean Christoph Jung, and Carsten Lutz. 2023. Towards Ontology Construction with Language Models. *arXiv preprint arXiv:2309.09898* (2023).

[15] Thomas R Gruber. 1995. Toward principles for the design of ontologies used for knowledge sharing? *International journal of human-computer studies* 43, 5-6 (1995), 907–928.

[16] Ramanathan V Guha, Dan Brickley, and Steve Macbeth. 2016. Schema. org: evolution of structured data on the web. *Commun. ACM* 59, 2 (2016), 44–51.

[17] Curtis E. Haas, Mary Ann Kliethermes, Lori T. Armistead, Craig J. Beavers, Christie A. Schumacher, Lisa Smith, John A. Armitstead, Roshni P. Emmons, Lucy I. Darakjian, Krystal L. Edwards, and Michael Barr. 2023. Comprehensive medication management: Review and recommendations for quality measures. *Journal of the American College of Clinical Pharmacy* 6 (2023), 404 – 415. https: //api.semanticscholar.org/CorpusID:256718733

[18] Josh Hanna, Eric Joseph, Mathias Brochhausen, and William R. Hogan. 2013. Building a drug ontology based on RxNorm and other sources. *Journal of Biomedical Semantics* 4 (2013), 44 – 44. https://api.semanticscholar.org/CorpusID:2304775

[19] María Herrero-Zazo, Janna Hastings, Isabel Segura-Bedmar, Samuel Croset, Paloma Martínez, and Christoph Steinbeck. 2013. An Ontology for Drug-drug Interactions. In *Workshop on Semantic Web Applications and Tools for Life Sciences.* https://api.semanticscholar.org/CorpusID:187394

[20] María Herrero-Zazo, Isabel Segura-Bedmar, Janna Hastings, and Paloma Martínez. 2015. DINTO: Using OWL Ontologies and SWRL Rules to Infer Drug-Drug Interactions and Their Mechanisms. *Journal of chemical information and modeling* 55 8 (2015), 1698–707. https://api.semanticscholar.org/CorpusID:20599292

[21] Adriana M. Ivama-Brummell, Daniella PINGRET-KIPMAN, Priscila G. Louly, and Rosiene R. Andrade. 2022. Medicines regulation, pricing and reimbursement in Brazil. *Revista Brasileira de Farmácia Hospitalar e Serviços de Saúde* (2022). https://api.semanticscholar.org/CorpusID:247722354

[22] S Jagannatha, TV Suresh Kumar, and R RajaniKanth. [n. d.]. Comparative Study of Back-End Vs Front-End System by Performance Analysis during Preliminary Design stages. ([n. d.]).

[23] Christopher Manning and Hinrich Schutze. 1999. *Foundations of statistical natural language processing.* MIT press.

[24] Microsoft. 2023. .NET Framework 4.5. https://docs.microsoft.com/en-us/dotnet/ framework/. Accessed: January 2023.

[25] Nelson Miranda and Dilvan de A. Moreira. 2024. Improving Medication Identification Accuracy and Regulatory Compliance through NLP and Ontologies: An Analysis of Otorhinolaryngology Prescriptions. In *2024 IEEE International Conference on Enabling Technologies: Infrastructure for Collaborative Enterprises (WETICE).* IEEE.

[26] Ol'ga Alekseevna Mitina and Ivan Alexandrovich Yurchenkov. 2021. Data Classification in Medicine and Healthcare Service. *Artificial Intelligence in Intelligent Systems* (2021). https://api.semanticscholar.org/CorpusID:237974497

[27] Mark A Musen. 2015. The protégé project: a look back and a look forward. *AI matters* 1, 4 (2015), 4–12.

[28] Natalya Fridman Noy, Ray W Fergerson, and Mark A Musen. 2000. The knowledge model of Protege-2000: Combining interoperability and flexibility. In *International Conference on Knowledge Engineering and Knowledge Management.* Springer, 17–32.

[29] Natalya F Noy, Deborah L McGuinness, et al. 2001. Ontology development 101: A guide to creating your first ontology.

[30] OpenAI. 2024. GPT-4 Turbo: Advanced language model for natural language understanding and generation. https://www.openai.com/models/gpt-4-turbo. Accessed: 2024-03-22.

[31] Christian Pachl, Nils Frank, Jan Breitbart, and Stefan Bräse. 2020. Overview of chemical ontologies. *arXiv preprint arXiv:2002.03842* (2020).

[32] Cecilia Reyes Peña, Mireya Tovar, Maricela Claudia Bravo, and Regina Motz. 2020. Drug Ontology for the Public Mexican Health System. In *SWH@ISWC.* https://api.semanticscholar.org/CorpusID:229366837

[33] Deborah L. Pestka, Caitlin K. Frail, Lindsay A. Sorge, Kylee A. Funk, Kristin K. Janke, Mary T. Roth McClurg, and Todd D. Sorensen. 2020. Development of the comprehensive medication management practice management assessment tool: A resource to assess and prioritize areas for practice improvement. *Journal of the American College of Clinical Pharmacy* 3 (2020), 448 – 454. https://api. semanticscholar.org/CorpusID:208485233

[34] Joselio Emar Araujo Queiroz, Robson Willian Melo Matos, Elivan Silva Souza, Paula Xavier Santos, Laís Bié Pinto Bandeira, Daniel Catão Moreira Licio, Beatriz Faria Leao, and Rafael Santos Santana. 2023. Ontologia brasileira de medicamentos: rumo à padronização terminológica de medicamentos no Brasil. *JORNAL DE ASSISTÊNCIA FARMACÊUTICA E FARMACOECONOMIA* 8, s. 2 (2023).

[35] Bartira Dantas Rocha, Larysse Silva, Thais Batista, Everton Cavalcante, and Porfírio Gomes. 2020. An ontology-based information model for multi-domain semantic modeling and analysis of smart city data. In *Proceedings of the Brazilian Symposium on Multimedia and the Web.* 73–80.

[36] Tomer Sagi, Matteo Lissandrini, T. Pedersen, and Katja Hose. 2022. A design space for RDF data representations. *The VLDB Journal* 31 (2022), 347 – 373.

https://api.semanticscholar.org/CorpusID:246186959

[37] Mark E. Sharp. 2017. Toward a comprehensive drug ontology: extraction of drug-indication relations from diverse information sources. *Journal of Biomedical Semantics* 8 (2017). https://api.semanticscholar.org/CorpusID:26441866

[38] S. Stelting and B. Maassen. 2020. Understanding Data Transfer Object (DTO) in C#. https://code-maze.com/different-ways-handle-errors-net-core-webapi/. Accessed: January 2023.

[39] Clarence Tauro, Nagesswary Ganesan, Saumya Mishra, and Anupama Bhagwat. 2012. Article: Object Serialization: A Study of Techniques of Implementing Binary Serialization in C++, Java and .NET. *International Journal of Computer Applications* 45 (05 2012), 25–29.

[40] Gemma Team, Thomas Mesnard, Cassidy Hardin, and Robert Dadashi et al. 2024. Gemma: Open Models Based on Gemini Research and Technology.

arXiv:2403.08295

[41] Ruben Verborgh and Max De Wilde. 2013. *Using openrefine*. Packt Publishing Ltd.

[42] Zizhong Wei, Dongsheng Guo, Dengrong Huang, Qilai Zhang, Sijia Zhang, Kai Jiang, and Rui Li. 2023. Detecting and Mitigating the Ungrounded Hallucinations in Text Generation by LLMs. In *Proceedings of the 2023 International Conference on Artificial Intelligence, Systems and Network Security*. 77–81.

[43] Marieke Wermuth. 2021. Terminological cooperation in the biomedical field. *Terminology* (2021). https://api.semanticscholar.org/CorpusID:237833789

[44] Yijia Zhang, Qingyu Chen, Zhihao Yang, Hongfei Lin, and Zhiyong Lu. 2019. BioWordVec, improving biomedical word embeddings with subword information and MeSH. *Scientific data* 6, 1 (2019), 52.