Personalized Depression Treatment Ontology

Cole Feuer (feuerc@rpi.edu)
Gunnar Eastman (eastmg2@rpi.edu)
Nancy Zhang (zhangn7@rpi.edu)

Abstract

In this paper, we discuss the Personalized Depression Treatment Ontology, an ontology developed to describe the symptoms, genetic factors, and treatments for Major Depressive Disorder (MDD) and related psychiatric disorders. With ongoing advances in the medical field, new treatments are continually being developed. However, predicting how individual patients will respond to specific treatments remains a challenge due to interactions between different factors such as genetics, environment, and medical history. This ontology seeks to fill that gap by creating a detailed framework that combines medication treatment recommendations with genetic and phenotypic data, both of which are essential for customizing treatment plans. Using this ontology, healthcare professionals will be able to go through a broad range of current medications that will be recommended and make more informed decisions based on both the presenting symptoms and the patient's genetic makeup. Specifically, the system will offer treatment recommendations that take into account genetic markers and alleles associated with better responses to particular medications, allowing for more precise, individualized care. Through the use of this ontology, we hope to contribute to the growing field of personalized medicine, offering a patient-specific, more data-driven approach to the treatment of depression and related mental health conditions. This approach has the potential to significantly improve patient outcomes and reduce trial-and-error prescribing, which often leads to delays in achieving effective treatment.

Introduction

The mental healthcare field has been in rapid development, with a large number of medicinal products being developed, tested, and released to treat similar psychiatric disorders. While this progress offers hope to many, it also challenges clinicians with navigating through the vast array of available treatments to find the most suitable option. Recognizing this complexity, we have developed the Personalized Depression Treatment Ontology (PDTO) which categorizes these treatments and highlights key differences in their effectiveness across various demographic groups. This ontology aims to make it easier to identify treatments that are more likely to work for specific individuals based on their unique characteristics.

The motivation for creating this ontology comes from the realization that depression treatments are not universally effective. Every patient is different, and their responses to medications can vary significantly based on factors such as age, gender, ethnicity, and genetic predispositions. This variation often leads to a frustrating process of trial and error in finding the right treatment. To address this, our ontology integrates data from studies that investigate treatment effectiveness with genetic markers, when available. This ensures that recommendations are evidence-based and tailored to each patient as much as possible, increasing the likelihood of success and reducing unnecessary side effects or delays in finding an effective treatment.

In creating the Personalized Depression Treatment Ontology, our goal is and continues to be to improve the precision of depression treatment and reduce the reliance on trial-and-error approaches. By integrating demographic-specific and genetic data, the ontology provides a more individualized approach to mental healthcare, ultimately helping patients find effective treatments faster and with fewer setbacks. This project reflects the growing need for personalized medicine and offers a pathway toward more effective and patient-centered care in the mental health field.

Related Work

Our work is far from the first to categorize medicinal products, nor is it the first to collect published scientific data. We owe many of the concepts captured by our ontology to prior work in the field.

First, we make use of several terms from the ontology for the Identification of Medicinal Products (IDMP). The central term that appears in both ontologies is Medicinal Product, representing a single drug that the patient would take. The IDMP's representation of a Medicinal Product and its associated concepts go into much more detail than is required for our work. As such, our ontology has reduced the level of detail.

Another ontology that we reference is the Study Cohort Ontology (SCO), which is dedicated to populations for clinical studies. We use this ontology for its representations of clinical trials, which are vital to our ontology as a method by which to provide provenance for the treatments that the system recommends.

Technical Approach

The approach taken for this work involved extensive documentation and iterative maintenance of several key artifacts throughout the project's lifespan. These included the Use Case Document, the Conceptual Model, and a Terminology List. Each of these artifacts played a vital role in ensuring the ontology's alignment with its intended purpose while allowing for flexibility to adapt to evolving requirements. Over the course of the semester, these documents underwent meticulous refinement to keep pace with the expanding scope and technical details of the project. By maintaining and refining these artifacts, the project established a robust, scalable, and contextually relevant ontology. This technical foundation was essential for addressing the complex and dynamic nature of personalized depression treatment. The resulting ontology is not only practical for clinical use but also extensible to accommodate future advancements in personalized mental healthcare.

Use Case

The Use Case Document served as an important artifact throughout the development process, providing a structured framework to demonstrate the system's intended functionality. It detailed the competency questions, which are high-level queries the system must be able to address, ensuring that our ontology met the needs of end users. Additionally, it outlined the flow of events, mapping how various processes would unfold within the system and ensuring logical consistency. The document also incorporated usage scenarios that depicted real-world applications, enabling us to consider the diverse ways the ontology might be utilized by different stakeholders. Resources, such as datasets and research articles were also recorded in this document, providing a foundation for system design and development. Figure 1 below is an overview diagram from this use case document, showcasing the different ways in which different classes of users might make use of the system.

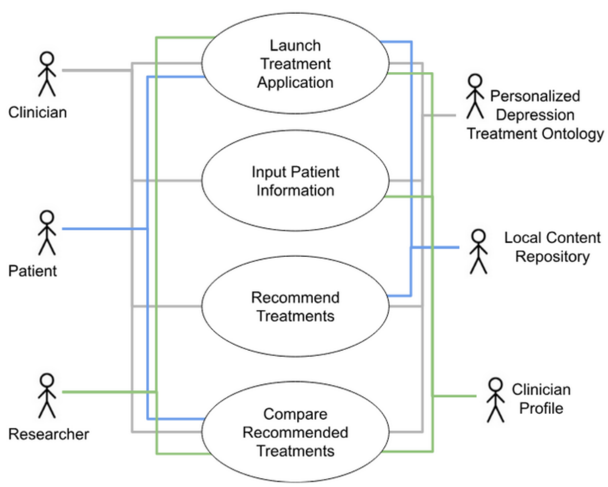

Figure 1: Activity Diagram showcasing use of PDTO

The full use case can be found on our website. [1]

Conceptual Model

The concept model for this project is in the form of a Google Slides presentation, which enabled us to break down the model and keep it readable. The first slide serves as an overview diagram which showcases the ontology's most important concepts: Condition, which was borrowed from the IDMP; Diagnosis; Gene; Patient; Treatment, and Treatment Response. The other 4 slides each go into further detail on a subset of the concepts from the first slide, with the second slide having more detail on the auxiliary concepts and properties of Patient and Condition; the third and fourth elaborating on Treatments and Responses; and the final slide containing more information on the Gene class.

The overview slide is shown below, in Figure 2. This slide is meant to capture the simplest version of the relations between these concepts: Each Patient experiences some number of symptoms, has some number of diagnoses for one condition each, a Genome comprised of some number of relevant genes, and has one response for each treatment they are undergoing.

The full conceptual model can be found on our website. [2]

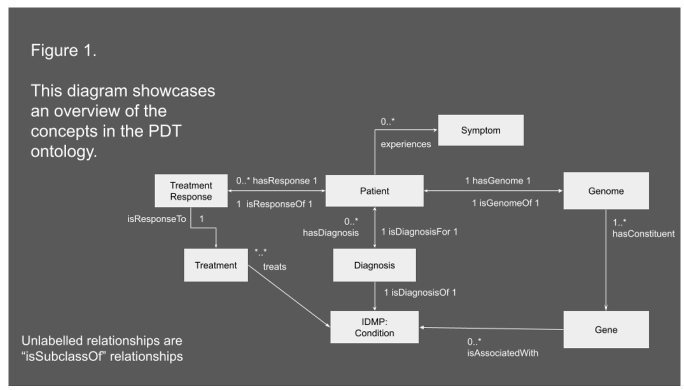

Figure 2: Concept Model Overview Slide

Evaluation

The ontology was evaluated by measuring its ability to answer each of the following competency questions. These four competency questions were converted into SPARQL queries, the results of which were used to judge the success of the project.

1. What are effective treatments for patients with the SLC6A4 long (L) genetic marker?

   To evaluate this, we wrote a SPARQL query that retrieved treatments and their effectiveness based on specific genes (e.g., SLC6A4-L). The query filtered treatments like citalopram and sertraline for patients with a high likelihood of positive response due to the presence of specific genetic variants. The evaluation

[1] https://personalized-depression-treatment–rpi-ontology-engineering.netlify.app/oe2024/personalized-depression-treatment/usecase

[2] https://personalized-depression-treatment–rpi-ontology-engineering.netlify.app/oe2024/personalized-depression-treatment/ontology

demonstrated that the ontology could successfully represent genetic influences on treatment responses, confirming the system's ability to answer such questions accurately.

SPARQL Query and Results:

```
PREFIX pdt: <https://tw.rpi.edu/ontology-engineering/oe2024/personalized-depression-treatment/PersonalizedDepressionTreatment/>

SELECT ?treatment ?effectiveness
WHERE {
  ?gene rdf:type pdt:SLC6A4 .
  FILTER (?gene = pdt:5-HTTLPR_L)
  ?gene pdt:hasResponse ?treatmentResponse .
  ?treatmentResponse pdt:isResponseTo ?treatment .
  ?treatmentResponse pdt:hasTreatmentEffectiveness ?effectiveness .
  BIND(
    IF(?effectiveness = pdt:high, 1,
      IF(?effectiveness = pdt:medium, 2,
        IF(?effectiveness = pdt:low, 3, 4)
      )
    ) AS ?priority
  )
}
ORDER BY ?priority
```

Treatment	Effectiveness
Citalopram	High
Escitalopram	High
Fluoxetine	High
Fluvoxamine	High
Paroxetine	High
Sertraline	High
Vilazodone	High

Table 1: Effectiveness of Treatments

2. What are effective treatments for a 35 year old patient with depression?

This question involved identifying treatments applicable to specific demographic groups, such as age ranges. The query filtered treatments by population target (e.g., treatments suitable for patients aged 35 years) using quantity value ranges. By evaluating the query results, we confirmed that the ontology adequately captured demographic-specific treatment applicability, enhancing its utility for personalized healthcare.

SPARQL Query and Results:

```
PREFIX pdt: <https://tw.rpi.edu/ontology-engineering/oe2024/personalized-depression-treatment/PersonalizedDepressionTreatment/>

SELECT ?treatment ?effectiveness
WHERE {
  ?gene rdf:type pdt:SLC6A4 .
  FILTER (?gene = pdt:5-HTTLPR_L)
  ?gene pdt:hasResponse ?treatmentResponse .
  ?treatmentResponse pdt:isResponseTo ?treatment .
  ?treatmentResponse pdt:hasTreatmentEffectiveness ?effectiveness .
  BIND(
    IF(?effectiveness = pdt:high, 1,
      IF(?effectiveness = pdt:medium, 2,
        IF(?effectiveness = pdt:low, 3, 4)
      )
    ) AS ?priority
  )
}
ORDER BY ?priority
```

Treatment Name
Citalopram
Escitalopram
Fluoxetine
Paroxetine
Sertraline

Table 2: List of Treatments

3. What alternative treatments are recommended for patients that have experienced poor response to SSRIs?

To address questions regarding ranking treatments by effectiveness, we developed a query that assigned priority values based on treatment effectiveness (e.g., high, medium, or low). The evaluation involved sorting treatments for specific responses (e.g., SSRIs for low responders) and verifying that the prioritized results matched expectations from real-world clinical scenarios. This confirmed that the ontology could handle treatment prioritization tasks, a critical functionality for personalized treatment recommendations.

SPARQL Query and Results:

```
PREFIX pdt: <https://tw.rpi.edu/ontology-engineering/oe2024/personalized-depression-treatment/PersonalizedDepressionTreatment/>

SELECT ?treatment ?effectiveness
WHERE {
  ?patient rdf:type pdt:Patient .
  ?patient pdt:hasResponse ?response .
  FILTER(?response = pdt:SSRI-Low)
  ?treatmentResponse rdf:type pdt:TreatmentResponse .
```

```
  ?treatmentResponse pdt:isResponseOf ?patient .
  ?treatmentResponse pdt:hasTreatmentEffectiveness ?effectiveness .
  ?treatmentResponse pdt:isResponseTo ?treatment .
  BIND(
    IF(?effectiveness = pdt:high, 1,
      IF(?effectiveness = pdt:medium, 2,
        IF(?effectiveness = pdt:low, 3, 4)
      )
    ) AS ?priority
  )
}
ORDER BY ?priority
```

Treatment	Effectiveness
Bupropion	High

Table 3: Effectiveness of Treatment

4. Which gene alleles impact the effectiveness of SSRIs and what are their impact?

   We tested whether the ontology could match patients with responses such as ”SSRI-Low” to effective treatments. A SPARQL query retrieved treatments, their effectiveness, and priorities for such patients. The results validated that the ontology accurately linked treatment responses to appropriate recommendations, demonstrating its reliability in complex, patient-centered scenarios.

   SPARQL Query and Results:

```
PREFIX pdt: <https://tw.rpi.edu/ontology-engineering/oe2024/personalized-depression-treatment/PersonalizedDepressionTreatment/>

SELECT ?gene ?treatment ?effectiveness
WHERE {
  ?gene rdf:type/rdfs:subClassOf* pdt:Gene .
  ?gene pdt:hasResponse/pdt:isResponseTo ?treatment .
  ?gene pdt:hasResponse/pdt:hasTreatmentEffectiveness ?effectiveness .
  ?treatment rdf:type/rdfs:subClassOf* pdt:SelectiveSerotoninReuptakeInhibitor .
}
```

Gene	Treatment	Effectiveness
SLC6A4-L	Citalopram	High
SLC6A4-L	Escitalopram	High
SLC6A4-L	Fluoxetine	High
SLC6A4-L	Fluvoxamine	High
SLC6A4-L	Paroxetine	High
SLC6A4-L	Sertraline	High
SLC6A4-L	Vilazodone	High
Val66Met_Met	Citalopram	Low
Val66Met_Met	Escitalopram	Low
Val66Met_Met	Fluoxetine	Low
Val66Met_Met	Fluvoxamine	Low
Val66Met_Met	Paroxetine	Low
Val66Met_Met	Sertraline	Low
Val66Met_Met	Vilazodone	Low
rs165599_A	Citalopram	Low
rs165599_A	Escitalopram	Low
rs165599_A	Fluoxetine	Low
rs165599_A	Fluvoxamine	Low
rs165599_A	Paroxetine	Low
rs165599_A	Sertraline	Low
rs165599_A	Vilazodone	Low

Table 4: Gene-Specific Treatment Effectiveness

The competency questions and SPARQL queries are found on our website [3].

For the ontology and all the queries written, a series of general evaluation methods were used. These methods were designed to test the ontology’s overall design, its ability to answer competency questions, and its readiness for practical application. Through these approaches, the ontology was refined iteratively to meet the requirements of personalized depression treatment recommendations.

One of the key methods used was coverage analysis, which focused on verifying whether all relevant concepts, relationships, and data elements were adequately represented within the ontology. This involved checking for critical entities, such as genetic markers, treatment options, and demographic factors, to ensure that no elements were missing. Instances, where ambiguities were identified, led to revisions in the ontology, improving its comprehensiveness and alignment with real-world treatment scenarios.

Another critical step was query testing, where SPARQL queries were tested against edge cases and unusual scenarios. This included examining how the ontology handled missing data, conflicting information, or boundary conditions in filtering and retrieval tasks. For example, when filtering treatments by target populations, queries were evaluated to ensure no results were returned for invalid or out-of-range conditions. This process was essential in identifying and correcting inconsistencies in the ontology’s structure and behavior.

[3]https://personalized-depression-treatment–rpi-ontology-engineering.netlify.app/oe2024/personalized-depression-treatment/demo

Competency question mapping was also a crucial part of the evaluation process. Each competency question was tied to specific SPARQL queries to verify that the ontology could fully answer these questions. By mapping queries to questions, we ensured that the ontology's design and content were sufficient to address its intended use cases. Iterative testing and adjustments to both the ontology and the queries ensured that the system evolved in alignment with these objectives.

Discussion

The development of the Personalized Depression Treatment Ontology (PDTO) represents an initial step toward integrating semantic technologies with the increasingly data-rich domain of personalized mental healthcare. By structuring key concepts—such as treatments, genetic markers, patient demographics, and clinical study cohorts—into an ontological framework, we enable meaningful reasoning about complex interactions that influence treatment efficacy. While this work, completed over a single semester, necessarily remains limited in its breadth and depth of coverage, it offers a tangible demonstration of how ontologies can support personalized medicinal treatment approaches.

The primary achievement of this ontology is the creation of a semantic foundation capable of answering key competency questions through SPARQL queries. These queries demonstrate that the ontology can identify treatments effective for patients harboring specific genetic markers (e.g., SLC6A4), highlighting relevant pharmacotherapies for particular demographic groups (e.g., 35-year-old patients), and suggesting alternatives for individuals who have shown poor responses to particular classes of medications (e.g., SSRIs). In doing so, PDTO showcases how semantic structures can reduce the guesswork in clinical decision-making and help keep clinicians better informed on the most up to date research. Clinicians, researchers, and other stakeholders can potentially use this framework to rapidly integrate new data, streamline treatment selection processes, and improve patient outcomes.

Value of Semantics

The ontology leverages semantics to support interoperability, data integration, and reasoning. Semantics play a pivotal role in integrating multiple data sources. For example, by linking genetic alleles, patient demographics, and treatment outcomes to standardized concepts and relationships, the system can integrate new datasets without needing to rebuild its foundational logic. This provides groundwork for large-scale integration of electronic health records, pharmacogenomic databases, and other clinical resources. Semantics also facilitate automated reasoning: using ontology-based inference, a system can reason about relationships among treatments, patient subgroups, and genetic factors that might not be evident from a manual review of studies and guidelines.

To read more and view our ontology, this is the link to the website:
https://personalized-depression-treatment--rpi-ontology-engineering.netlify.app/oe2024/personalized-depression-treatment/

Support of Claims

Each of our competency questions is able to be translated into a SPARQL query that will produce correct and intended results. This was achieved by designing each query to map directly to the concepts and relationships defined in the Personalized Depression Treatment Ontology. By aligning the structure of our queries with the ontology's classes and properties, we ensured that the queries retrieved relevant data related to treatment effectiveness, patient responses, and demographic factors. Additionally, we validated the accuracy of the results through testing, comparing the query outputs against expected outcomes based on real-world data and scientific studies. This process not only confirmed the reliability of the queries, but also demonstrated that the ontology supports meaningful insights and actionable information in the context of personalized depression treatment.

Limitations

The primary limitation for this work is the allotted time. This work taking place over the course of one semester means that our scope was severely limited. Specifically, the system is lacking in individuals, which leaves it unable to reason about neither all relevant genes nor all relevant treatment options. The system also does not contain the concepts that would allow it to reason about treatments outside of medicinal products, such as therapeutic interventions or lifestyle adjustments that might complement or substitute pharmacological treatments. Moreover, the ontology does not yet account for the full spectrum of demographic variables that could influence treatment efficacy, such as socioeconomic factors, mental health comorbidities, or environmental influences. These gaps limit the comprehensiveness of the ontology and its ability to fully support personalized treatment recommendations. Furthermore, the ontology's reliance on existing literature means that the data incorporated may be incomplete or biased by the availability of research in specific areas. These limitations highlight the need for future work to expand the ontology, incorporate a broader range of treatment modalities, and refine its ability to represent individual variability in response to treatments.

Future Work

Much to our chagrin, there are many directions that future work could take to improve the coverage and depth of the ontology. First is the very limited number of individuals, which could be remedied by adding more treatments, genes, and symptoms.

Future work could also focus on increasing the amount of personal data used to determine a recommendation. The system does not currently support the use of blood test results. Alternatively, non-medicinal treatments could be recommended by the system. Concepts currently exist within the ontology to capture therapeutic treatments, but the recommendation system does not utilize them, while things such as dietary changes or lifestyle interventions are not captured as concepts at all, despite their potential to treat some symptoms.

Conclusion

In conclusion, the Personalized Depression Treatment Ontology represents a significant step forward in the field of mental healthcare, particularly in the context of depression treatment. By providing a structured, comprehensive framework for categorizing medicinal products and understanding their efficacy across different demographics, this ontology addresses the complexities involved in selecting appropriate treatments for individuals. The inclusion of genetic markers and their potential influence on treatment outcomes further enhances the personalization aspect, making treatments more tailored to the needs of each patient. Throughout its development, the ontology has been evaluated against competency questions, ensuring that it meets the key requirements for both practical application and future research. As mental health treatments continue to evolve, this ontology will remain a valuable resource for healthcare providers and patients, enabling more informed, individualized decisions and ultimately improving patient outcomes in the fight against depression.

Acknowledgements

- Elisa Kendall
- Deborah McGuinness
- Jade Franklin
- Kelsey Rook

References

- Course Materials
- IDMP