

Ontology Integration Experienced on Medical Terminologies

Yugyung Lee and Kaustubh Supekar
School of Computing and Engineering
University of Missouri, Kansas City
Kansas City, MO 64110

James Geller*
CS Department
New Jersey Institute of Technology
Newark, NJ 07102

Abstract

To build a common controlled vocabulary is a formidable challenge in Medical Informatics. Due to vast scale and multiplicity in interpretation of medical data, it is natural to face overlapping terminologies in the process of tracing medical informatics [40]. A major concern lies in the integration of seemingly overlapping terminologies in the medical domain and this issue has not been well addressed. In this paper, we describe a novel approach for medical ontology integration that relies on the theory of Algorithmic Semantic Refinement we previously developed. Our approach simplifies the task of matching pairs of corresponding concepts derived from a pair of ontology, which is very vital to terminology mapping. A formal theory and algorithm for our approach is devised and the application of this method to two medical terminologies is developed. The result of our work is an integrated medical terminology and a methodology and implementation ready to use for other ontology integration tasks.

Keywords: Terminology, Ontology, Ontology Integration, Semantics, Semantic Enrichment, Semantic Refinement, Semantic Integration.

1 Introduction

In medical research, there is a need to exchange valuable information between different researchers or research groups, for the purpose of independent analysis or the verification of experimental results. Increasingly, we are also seeing the emergence of distributed scientific processing. The Internet provides an important platform for this activity of medical information exchange to take place. However, there are still some difficulties to resolve before seamless interoperability and interchange can occur. The main cause for these limitations arise from the fact that different research groups rely on heterogeneous research data sources.

Over the past several years, one promising line of approach that has been very active in building and using terminologies, ontologies and vocabularies has been Medical Informatics. There exist a large number of terminologies developed for different purposes (literature indexing and retrieval, electronic patient records, statistical reports on mortality, billing), in different subdomains (diseases, genomes, micro-organisms, diagnoses, medical devices, procedures, drugs). These terminologies have been built by different institutions (World Health Organization, National Library of Medicine, College of American Pathologists, etc.) on different continents, for different purposes. Yet, attempts to represent the whole medical domain are usually limited in scope (GALEN) [41] or lack a strong organizational structure, as in the Unified Medical Language System (UMLS). The UMLS [56] is a compendium of about 100 individual terminologies, but integrating a new terminology into it is a difficult task. Existing methods for integration of ontologies use structural and semantic methods (see Section 6), however, there is still room for improvement.

Most ontologies [13] are organized around a concept hierarchy (a tree or a directed acyclic graph). Many ontologies add rules, axioms, or other additional mechanisms to this backbone. In this paper, as a first step, we will only deal with integration of the concepts in the concept hierarchy. When integrating two ontologies, it

*This research was supported in part by the New Jersey Commission for Science and Technology.

is necessary to identify pairs of concepts that have the same meaning in both of them. Clearly, these concepts should occur only once in an integrated ontology. The existence of synonyms and homonyms causes problems for this kind of integration. However, a much bigger problem is the existence of subtle differences between *implemented* concepts that have the same name and stand, vaguely, for the same (concrete or abstract) real world objects. These differences occur because different ontology designers may bring different world views to the task, conceptualizing the world at different levels of granularity and abstraction. Such differences are commonly considered semantic problems.

Implementations add structural problems also. We note that for human communication the lack of exact matches does not normally make communication impossible. For the example of a calendar date, people might greatly disagree on the exact structure *beyond the obvious data type and name issues*. Thus, a person *A* may use the US calendar. However, *A* is a professor, and his academic calendar defines the Wednesday before Thanksgiving to be a Friday. Thus, *A* distinguishes between two kinds of dates. For *A*, any date needs to be annotated as a US date or an academic date. Person *B*, on the other hand, might not know about academic calendars and will not need any additional annotations. Clearly, *A* and *B* have concepts of calendar dates that are sufficiently different to make an exact match impossible. Nevertheless, *A* and *B* are able to communicate about dates in most situations. Thus, differences between concept representations should not automatically exclude matches.

In this paper, we propose to address some issues in ontology integration of medical terminologies by extending our semantic refinement methodology, called Algorithmic Semantic Refinement (ALSER) [11, 14]. To capture the essence of ALSER in two sentences, we "compute" small sets of concepts of similar semantics based on a superficial human specification that does *not* have to be free of contradictions. The "computation" is based on intersections of concept sets. Our methodology was previously applied to the problem of auditing the very large UMLS [14].

The ALSER methodology is based on a kind of terminology that we call a *Terminological Knowledge Base (TKB)*. However, there are many important medical ontologies or controlled medical terminologies, which are not structured as TKB. Thus, we first address the question of how to transform existing medical terminologies into a TKB. Due to the great difficulty of this problem, we are only addressing the case where a global reference ontology (i.e., UMLS) exists in the same domain as the given medical terminology. Our methodology, called *semantic enrichment*, builds a TKB by finding its concepts in a global ontology that has semantic types. Secondly, we describe the Algorithmic Semantic Refinement methodology (ALSER) generating precompute sets of semantically similar terms in both ontologies that need to be integrated. Finally, we introduce the *semantic integration* methodology, which compares and integrates terms from two ontologies, if they are already classified as semantically similar.

We believe that this *incremental* approach composed of *semantic enrichment*, *semantic refinement* and *semantic integration* will reduce the likelihood of false positives, since we avoid matching concepts of different semantics, which out of principle cannot be the same. The semantic approach has an additional advantage. It greatly reduces the computational effort of the matching operations. Fewer terms will have to be matched against each other.

In Section 2, we describe the difficulties in handling medical terminologies and also propose our semantic enrichment approach. In Section 3, we introduce our methodology of Algorithmic Semantic Refinement. In Section 4, we show how this methodology can be used to create pairs of candidate sets of concepts that have to be matched with each other. A prototype implementation and results are described in Section 5. In Section 6, we review related research work. Section 7 contains our conclusions.

2 Semantic Enrichment

The semantic enrichment process of finding semantic types for concepts is difficult, even in the medical domain, with the UMLS readily available. In this section, we briefly survey the two medical terminologies to describe some obstacles that we have encountered during the integration of medical ontologies and to highlight the necessity of semantic enrichment as a preprocess of ontology integration. The American College

Relationship	Concept	Category
IS-A	Gender	Demographics
	Weight	History and Risk Factors
IS-A (Prefix-of)	RF-Diabetes	History and Risk Factors
	Meds-Digitalis	Pre Operative Medications
IS-A (Postfix-of)	Thrombolysis-Intvl	Previous Interventions
	Ace-Inhibitors - Discharge	Discharge
Attribute-of	Participant ID	Administrative
	Hospital ZIP Code	Hospitalization
Attribute-of (Compound)	Patient SSN/Country Code	Hospitalization
	Clopidogrel/Ticlopidine	Medications
Instance-of	Left Main Dis > 50%	Diagnostic cath procedure-findings
	Comps-Neuro-Cont Coma $\geq$ 24Hrs	Complications

Table 1: Examples of relationships between Concepts and Categories in ACC/STS

of Cardiology (ACC) has provided a list of 142 terms with definitions [1]. These concepts are separated into 22 "categories." The Society of Thoracic Surgery (STS) has created a classification of 248 terms, subdivided into 21 categories [48].

2.1 Difficulties in Handling Medical Terminologies

Both these ontologies have concepts and categories for describing cardiovascular domain knowledge. Two major issues we faced were: (1) A great degree of inconsistency exists among the (perceived) relationships between concepts and categories; (2) Inconsistent patterns appeared in either the concept names or the category names. The inconsistent naming created major obstacles in matching and automated categorization. Above we wrote "perceived relationships," because the ontology itself does not name the relationship that is supposed to hold between one concept and its category. Thus, the user is left with the task of guessing each relationship.

Intuitively, categories should have been introduced for the purpose of categorizing concepts (similar to semantic types in the UMLS). This is what the name "categories" seems to imply. However, as mentioned before, there exist different kinds of relationship between concepts and categories of the ACC and STS. This forces us to evaluate each relationship and to incorporate its treatment in the semantic enrichment algorithm. If there exists an IS-A relationship between a concept and a category, then a semantic type of the category can be propagated to the concept. Otherwise we use the category information provided, according to our understanding of the relationship that is presumably holding between the concept and the category.

In Table 1, each *IS-A* relationship describes a super/subclass relationship between a concept and a category (e.g., *Gender* is a *Demographics [Item]*). In the ACC and STS terminologies, the category occasionally appears as prefix or postfix in the concept name. Those prefixes or postfixes provide additional context, which is useful for determining the semantic type of a concept (e.g., *Thrombolysis-Intvl* contains *Intervention* as a postfix and *RF-Diabetes* contains *Risk Factor* as a prefix). Thus, we define *IS-A (Prefix-of)* and *IS-A (Postfix-of)* relationships as IS-A relationships. Occasionally, like above, a prefix or postfix occurs as an abbreviation. However, this does not have to be the case. In order to handle acronyms, a list of domain specific acronyms can be stored in a database and converted into full names such as *Risk Factor* for *RF*, *Medications* for *Meds*, *Valve Surgery* for *VS* and *Vessel Disease* for *VD*.

The *Attribute-of* relationship describes that a concept is a database field of a category (e.g., *Participant ID* is a field of the *Administrative* table). The *Instance-of* relationship defines a concept as a specific instance of a category (e.g., *Comps-Neuro-Cont Coma $\geq$ 24Hrs* is an instance of *Complications*). There are some *ambiguous* categorizations that exhibit a lack of evidence for determining a concept as belonging to a category (e.g., *Hypertension* is a category of *History and Risk Factors*, *Diabetes* is a category of *History and Risk Factors*).

Table 2 shows some patterns that appeared in ACC or STS concepts. The *Instance-of* relationship

Pattern	Name	Description
Instance-of	Valve Disease - Mitral	<i>Mitral</i> is an instance of <i>Valve Disease</i> .
Acronym-of	VS-Aortic Proc-Procedure	<i>VS</i> is a <i>Valve Surgery</i> .
	VD-Insuff-Mitral	<i>VD</i> is <i>Vessel Disease</i> .
Synonym-of	Patient DOB	<i>DOB</i> is <i>Date of Birth</i> .
Multi-words	Conversion to Std Incision	<i>Conversion</i> determines the semantic type.
	Skin Incision Start Time	<i>Time</i> determines the semantic type.
	Primary Cause of Death	<i>Cause</i> determines the semantic type.
Redundant word	Cross Clamp Time (min)	<i>(min)</i> is redundant.
Redundant word	Unique Patient ID	<i>Unique</i> is redundant.
Symbol	CAB During This Admission - Date	"-" is a symbol.
Abbreviation	Comps-Op-ReOp Other Card	<i>Comps</i> is an abbreviation of <i>Complications</i>
Compound words	Comps-Op-ReOp Bleed/Tamponade	<i>Bleed</i> and <i>Tamponade</i> are compound words.
Inconsistency	$\geq$ and "Greater than Equal"	Different notations for the same concept

Table 2: Complications that appeared in ACC/STS Concept Names

describes a relationship between words in the concept (e.g., *Mitral* is an instance of *Valve Disease*). In the multi-word case of the form *Skin Incision Start Time* the last word *Time* determines the semantic type while in the case of *Primary Cause of Death*, the word *Cause* before *of* determines the semantic type. In a noun-noun phrase, the determining word is typically the second noun, which is referred to by linguists as *head noun*. However, there are famous exceptions to this rule, such as *toy gun*, which is a toy, not a gun. In this case, the first noun would be used to determine the semantic type of the noun-noun phrase.

The string "(min)" is marked as *redundant*, as it is not really a part of the concept term, but provides additional information about this concept. In this specific case, it provides the unit of measurement of the quantity that is measured by the concept.

2.2 Details of Data Enrichment

By first classifying the concepts of a new terminology using the Semantic Network of UMLS, this task becomes more manageable. A new concept does not have to be compared with every concept in the UMLS, but only with those UMLS concepts that have the same assignments to semantic types as the new concept. This considerably reduces the difficulty of integration, as long as all assignments to semantic types have been made correctly and consistently.

In order to perform semantic enrichment we need to identify pairs of concepts from different ontologies (or concept–category pairs) that have the same meaning. This step, called concept matching, requires that we overcome many issues of inconsistent naming which are usually obvious to humans but difficult to handle for algorithms.

For instance, many terminologies freely mix the use of terms with the acronyms or abbreviations of those terms. Thus, these abbreviations need to be expanded for easier concept matching. We call this expansion step *data enrichment*. Data enrichment also performs other preprocessing and clean-up steps, such as dealing with non-alphabetic characters occurring in many medical terms.

For example, the acronym RF needs to be replaced by its expansion, Risk Factor. The abbreviation "Meds." is replaced by Medications. Other common medical acronyms in our terminologies are DOB (Date of Birth), MI (Myocardial Infarction), and many more. Whenever terms contain special characters such as "/" or "." they are replaced by blanks. When there exist the prefix or postfix cases (e.g., RF-Smoker), it is converted into a special form (e.g., Risk Factor:Smoker). In this way, those terms can be matched with other terms of the same meaning which are lacking the special characters. In some cases, precise mathematical symbols are expressed by imprecise English words. For example, the mathematical notation "greater than equal to" is transformed from its English representation into its well defined symbolic representation $\geq$. This symbolic representation is unique, while the English representation may equally appear as "greater equal" or "greater than or equal to," etc.

Third, the existence of synonyms and homonyms causes problems for concept matching. The use of synonyms is absolutely necessary, because medical terminologies are full of variant terminologies (e.g., Heart and Coeur, Heart Block and Lev's disease).

While acronyms can be dealt with by expansion into a canonical form, this is harder for synonyms. Rather, we have decided to include the use of synonyms during the concept matching step itself. If no match is found for a concept, then it is attempted to use its synonyms for matching. The synonyms of terms were derived both from STS and ACC documentation and from the UMLS (currently still manually, in a preprocessing step).

2.3 Details of Semantic Enrichment

We now explain the semantic enrichment process in more detail (Figure 1). Semantic types are drawn as little squares and named with capital letters, while concepts are drawn as circles and named with small letters. ACC and STS categories are drawn as filled in circles. The medical terminologies (ACC, STS) do not have semantic types. Thus, a global ontology (the UMLS) is used to add semantic type assignments for the concepts in the ACC and the STS.




Figure 1: Local and Global Ontologies

There are two possible cases how semantic enrichment can be done. In *Case 1*, a concept of the local terminologies is found in the global ontology. This concept has one or several assigned (global) semantic types which are returned. In Figure 2, the concept *a* matches the global concept *z*. This might happen if *a* is a synonym of *z* in the global ontology. Because *z* has assigned semantic types *U* and *S*, *a* acquires *U* and *S* as its semantic types.

In *Case 2*, a local concept without a match in the global terminology belongs to a category such that an IS-A relationship holds between the concept and the category. In this case it is reasonable to look in the UMLS for the semantic type of the category and assign this semantic type to the concept for the following reason. In some cases, concepts are ambiguous. However, the category eliminates this ambiguity. Thus, the semantic type of the category should help to better define the meaning of the concept.

In Figure 1, the concept *b* is connected to the category *m* by an IS-A link. The category *m* matches the concept *y*, which has a semantic type *W* assigned to it. Thus, we assign the semantic type *W* to the concept *b*. But *a* is also connected to the category *m* by an IS-A link. Thus, *a* now gets the additional semantic type *W* assigned to it, resulting in *a* having *S*, *U*, and *W* as semantic types (see Figure 3).

We note that whenever a concept is assigned several semantic types, we make sure that there are no *redundant* assignments. The assignment of a concept to a semantic type R_p is *redundant* if that concept is also assigned to another semantic type R_c , and R_p is a parent or ancestor of R_c in the semantic network




Figure 2: Mapping between Local and Global Ontologies

[14, 37]. For example, in Figure 2, assigning the semantic types W and Z to the concept w would be redundant, because w is already assigned to T , and W is a parent of T and an Z is an ancestor of T .

Figure 3 shows the results of semantic enrichment. The concepts c and d are crossed out, because there were no semantic types corresponding to them in the global ontology. We do not allow concepts without semantic types. In our experiments, these cases were rare. Only 4 out of 248 concepts were dropped for this reason.

3 Overview of Semantic Refinement

3.1 The Unified Medical Language System

Algorithmic Semantic Refinement (ALSER) was invented [14, 11] in the context of the Unified Medical Language System (UMLS). Thus we start with a brief summary of the UMLS. The UMLS [16, 17, 25, 56], designed by the National Library of Medicine (NLM), consists of three Knowledge Sources of which we are interested in two, the Metathesaurus [46, 54] and the Semantic Network [29, 30, 31]. The *Metathesaurus* is a unified collection of many different medical terminologies. It is a compilation of terms, concepts, relationships, and associated information. The January 2003AA edition includes 875,255 concepts and 2.14 million concept names in over 100 biomedical source vocabularies.¹

The *Semantic Network* of the UMLS contains 135 semantic types (e.g., **Disease or Syndrome**, **Virus**). One may think of semantic types as high-level concepts, i.e., broad categories. These semantic types are organized in a hierarchy of IS-A links. The hierarchy consists of two trees, rooted in the semantic types **Entity** and **Event**. In addition, there are 53 kinds of non-IS-A relationships among these semantic types, e.g., *causes*, used in: **Virus causes Disease or Syndrome**. Every concept in the Metathesaurus is assigned to at least one, but often several, semantic types in the Semantic Network. One can say that a concept (in the Metathesaurus) is assigned some semantics by being assigned to a semantic type in the Semantic Network.

3.2 Methodology for Semantic Refinement

We formally present the ALSER methodology that we have published in [11].

¹<http://www.nlm.nih.gov/pubs/factsheets/umlsmeta.html>




Figure 3: Enrichment of Local Terminology after the Mapping

Definition 1: Terminological Knowledge Base. We call any structure that consists of (1) a semantic network of semantic types; (2) a thesaurus of concepts; and (3) assignments of every concept to at least one semantic type a *Terminological Knowledge Base* (TKB). Thus, a TKB is a triple:

$$TKB = \langle \hat{\mathcal{C}}, \hat{\mathcal{S}}, \mu \rangle \quad (1)$$

in which $\hat{\mathcal{C}}$ is a set of concepts, $\hat{\mathcal{S}}$ is a set of semantic types, and μ is a set of assignments of concepts to semantic types. We will use capital letters for semantic types and small letters for concepts.² Finally, μ consists of pairs (c, S) such that the concept c is assigned to the semantic type S .

$$\hat{\mathcal{S}} = \{W, X, Y, \dots\}; \quad \hat{\mathcal{C}} = \{a, b, c, d, e, \dots\} \quad (2)$$

$$\mu \subset \{(c, S) \mid c \in \hat{\mathcal{C}} \text{ \& } S \in \hat{\mathcal{S}}\} \quad (3)$$

In [11], $\hat{\mathcal{S}}$ and $\hat{\mathcal{C}}$ form DAG structures. In this paper, they are sets. Furthermore, it holds:

$$\forall c \in \hat{\mathcal{C}} [\exists p \in \mu [p = (c, S) \text{ \& } S \in \hat{\mathcal{S}}]] \quad (4)$$

In words, every concept must be assigned to at least one semantic type. The opposite condition does not hold. ALSER takes as input one Terminological Knowledge Base and returns as result a second Terminological Knowledge Base with better semantics. What “better semantics” exactly means will become clear by the end of this Section. The refinement can be written as a function:

$$TKB' = \text{ALSER}(TKB); \quad TKB = \langle \hat{\mathcal{C}}, \hat{\mathcal{S}}, \mu \rangle \quad (5)$$

Next, we need a notion when two concepts are semantically similar. We will consider two concepts c, d as similar, $c \simeq d$, when they are assigned to exactly the same set of semantic types of a TKB.

$$c \simeq d : \forall S \in \hat{\mathcal{S}} [(c, S) \in \mu \Leftrightarrow (d, S) \in \mu] \quad (6)$$

Intuitively, if two concepts c, d are assigned to exactly the same semantic type X , then we can say that these two concepts have similar semantics. If two concepts c, d are assigned to X and are also assigned to Y , then these two concepts also have similar semantics. On the other hand, if a concept a is assigned to X and a concept b is assigned to both X and Y , then a and b will have semantics that are not similar in the formal sense defined above. (They are still similar, to a lower degree, in the real world.)

²Both roman and italic fonts




Figure 4: Algorithmic Semantic Refinement Example in Graph Notation

We will use a running example to explain our rather abstract definitions. Figure 4(a) shows a TKB before applying the ALSER algorithm. Figure 4(b) shows the corresponding TKB' that is the output of ALSER. Capital letters represent semantic types, small letters represent concepts. A line connecting a semantic type to a concept means that this concept is assigned to this semantic type. The following is an informal explanation of Figure 4, which will be formalized shortly. As the figure shows, the result of ALSER is that semantic types are refined, and every concept in the output is assigned to exactly one refined semantic type. Thus, refined semantic types never share concepts with one another. All concepts assigned to any one refined semantic type in the output are similar ($\simeq$).

Definition 2: Original Semantic Type. A semantic type of a given TKB (i.e., one that is intended as input to ALSER) is called an original semantic type. *Example:* X, Y, Z and W are original semantic types in Figure 4(a).

Definition 3: Simple Concept. A simple concept is assigned to one single semantic type. *Example:* a, c, f and h are simple concepts in the input.

Definition 4: Compound Concept. A compound concept is assigned to two or more semantic types. *Example:* b, d, e, and g are compound concepts in the input in Figure 4(a).

Definition 5: Refined Semantic Type. Every semantic type which is part of the output of ALSER is called a refined semantic type. *Example:* X', Y', Z', W', X∩Y, X∩Z, W∩X∩Y∩Z are refined semantic types in Figure 4(b).

Definition 6: Simple Semantic Type. A simple semantic type P' is a refined semantic type defined to be assigned all and only the simple concepts assigned to the original semantic type P . *Example:* The refined semantic types X', Y', Z' and W' are simple semantic types.

Definition 7: Intersection Type. An intersection type I is a refined semantic type defined to be assigned all and only the compound concepts assigned to exactly one specific set of original semantic types. An intersection type is denoted by the sequence of those original semantic types, where the intersection symbol appears between any two consecutive original semantic types in the sequence. *Example:* X∩Y, X∩Z and W∩X∩Y∩Z are intersection types.

The intersection type W∩X∩Y∩Z is the same as the intersection type Y∩X∩W∩Z, etc. For clarity, an intersection type may have a shorter direct name which may be chosen to reflect the semantics of the

concepts. Many times there are obvious names for intersection types. For example, the intersection type of the semantic type Body Part with the semantic type Manufactured Object is commonly known as prosthesis.

Definition 8: Semantically Uniform TKB. A TKB is called semantically uniform if and only if each concept is assigned to exactly one semantic type.

In summary, the input to ALSER consists of concepts which may be assigned to several semantic types, which is undesirable.³ In the output of ALSER, every concept is assigned to a single refined semantic type. Therefore, all concepts assigned to any one refined semantic type are now similar, according to the previous definition of similarity $\simeq$. The output of ALSER may contain empty semantic types.

As even our small example shows, the output will (normally) contain more (refined) semantic types than the input had (original) semantic types. Each simple semantic type has, on average, fewer concepts assigned than the corresponding original semantic type. We will show this effect with real data in Section 5. We can now advance to a **procedural definition** of ALSER. The easiest part of the ALSER algorithm deals with the set of concepts. Algorithmic Semantic Refinement does not change the number or names of concepts in a TKB. It only reassigns the concepts to refined semantic types.

$$\hat{\mathcal{C}}' = \hat{\mathcal{C}}. \quad (7)$$

In order to find the set $\hat{\mathcal{S}}'$, we first need to identify the simple semantic types. As noted above, every original semantic type with at least one simple concept is mapped into a simple semantic type. All the simple concepts assigned to an original semantic type are reassigned in the output to the corresponding simple semantic type. In the second step, we need to identify all intersection types. Every compound concept identifies a set of original semantic types. Every such set needs to be mapped into one intersection type. The compound concept will then be reassigned to this intersection type in the output. Formally, for every combination of k original semantic types $S_{i_1}, S_{i_2}, S_{i_j}, \dots, S_{i_k}$ from TKB such that there is at least one single concept c that is assigned to all S_{i_j} , for $1 \leq j \leq k$, we create an intersection type $S_{i_1} \cap S_{i_2} \cap \dots S_{i_k}$ in TKB'.

The derivation of μ' also consists of two parts. μ'_s contains the assignments of concepts to simple semantic types. μ'_i contains the assignments of concepts to intersection types. For a specific simple semantic type X' , μ'_s can be found as follows:

$$\mu'_s = \{(c, X') \mid \text{simple}(c) \ \& \ (c, X) \in \mu\} \quad (8)$$

4 Semantic Integration

In the real world, there is a spectrum of requirements one could impose to accept two concepts as matching. On one extreme, one might insist that there be only perfect matches between two concepts. In the example of calendar dates, one might require that both are called “date” and have exactly three fields, two numbers and one character string (for the month). Furthermore, there also must be a match between the attribute names. For example, one might insist that both of them have an attribute called “month.” Such requirements would lead to a few matched concepts and large numbers of unmatched concepts.

The other extreme is to insist that all (or almost all) concepts of the smaller ontology are matched against concepts in the larger ontology, as long as there is at least some structural similarity. This extreme could be based on the assumption that both ontology designers did a reasonable job to cover the domain, and thus a fundamental concept such as “date” simply has to appear in both ontologies, no matter what it is called, and no matter how exactly it is structured. Our solution is closer to the second extreme.

³Why it is undesirable was explained in great detail in [14, 11]. From the integration point of view, it is impossible to tell whether two compound concepts are similar without looking at all semantic types they are assigned to.

4.1 Which Pairs Need to be Matched?

We now assume that we have to integrate two Terminological Knowledge Bases TKB and TKB_2 . Before we can integrate, we need to identify for which concepts to attempt integration. This requires three steps: (1) We transform TKB into TKB' by applying ALSER. (2) We transform TKB_2 into TKB'_2 by applying ALSER. Step (2) is the transition that happens from the right top part of Figure 5 to the right middle part of the figure. Note that concepts of TKB are assigned by semantic types of UMLS during semantic enrichment. Thus, *Cardiogenic Shock* and *Arrhythmia Type* are assigned to *Functional Concept* and *Therapeutic or Preventive*, *Diabetes* is assigned to *Disease or Symptom*, etc. The set of refined semantic types of TKB'_2 is not necessarily the same as the set of semantic types of TKB' . In our Figure, TKB'_2 has an additional intersection type *Therapeutic or Preventive* $\cap$ *Disease or Symptom* which does not occur in TKB' . It is assigned the concept *Congestive Heart Failure*. (3) We identify pairs of sets in TKB' and TKB'_2 that are potential matches (Figure 5). At the bottom of Figure 5, no match will be attempted for the concept *Congestive Heart Failure*. An attempt is made to match the concepts in the set $\{Cardiogenic Shock, Arrhythmia Type\}$ against the concept *Cardiogenic Shock*. Similarly, the match of *Previous CAB* and *Previous CAB* is attempted. Formally:

1. Assume a refined semantic type S exists in TKB' that has assigned concepts $x, y, z, \dots$. Further assume that S does not exist in TKB'_2 or, there are no concepts assigned to S in TKB'_2 . Then, by the similarity assumptions made above, no concepts corresponding to $x, y, z, \dots$ exist anywhere in TKB'_2 . Thus, these concepts do not need to be matched at all. Alternatively, the concepts $x, y, z, \dots$ might have been misclassified along the way ([11]) by human experts. In that case, the integration effort would make those problematic concepts easier to detect [11].
2. The above observation applies in reverse also. If an intersection type S has been generated for TKB'_2 that does not exist in TKB' , then the concepts $x, y, z, \dots$ assigned to S will not have corresponding concepts anywhere in TKB' . Thus, these concepts do not need to be matched at all. The same observation as above applies to possible misclassifications.
3. All pairs of concepts (q, r) with $q \in \text{CONCEPTS}(S, TKB')$ and $r \in \text{CONCEPTS}(S, TKB'_2)$ are similar ($\simeq$) and need to be matched. The matching algorithm follows in Section 4.2.

A brute force approach would be to compare every concept from TKB' with every concept from TKB'_2 . In our approach, we only compare pairs of concepts for which we already know that they are similar.

4.2 Scoring Concept Similarities

Now we describe details of how scores for concept similarities are computed. We use three aspects to determine whether a match exists between similar concepts. Initially we rank pairs of concepts according to their terms (or synonyms) and then according to attribute similarity. After establishing some initial matches in this way, we use relationships that point from one concept to another concept to iteratively recompute the similarity value between two concepts.

4.2.1 Ranking Concepts by Terms

If two concepts have similar names (defined below, based on bigrams) then they are possibly matches. The existence of synonyms and homonyms causes problems for concept matching. The use of synonyms is absolutely necessary, because medical terminologies are full of variant terminologies (e.g., Heart and Coeur, Heart Block and Lev's disease). We include the use of synonyms during the concept matching step itself. If no match is found for a concept, then it is attempted to use its synonyms for matching.

The bigram approach is known to be a very effective, simply programmed means of determining a similarity measure between strings. The bigram approach consists of three steps. (1) Sequences of two consecutive letters within a string are extracted (e.g., the word "heart" contains 'he', 'ea', 'ar' and 'rt'); (2)




Figure 5: Overall Process of Semantic Integration

Two sequences of bigrams are compared, and a raw similarity score is computed; (3) A matched score is computed from the raw score, i.e., the number of the common bigrams is divided by the average number of bigrams in the two strings.

4.2.2 Ranking Candidates by Attributes

Assume that we are given a pair of concepts from two different ontologies. These concepts have different terms, therefore, a priori there is no reason for a computer to assume that they are in fact describing the same concepts. In order to establish whether they are indeed the same concept, we need to compare attributes.

We assign to every pair of concepts a score as follows.

1. Two concepts that have the same number of attributes, and for every attribute in one concept there is an attribute in the other concept of the same name and same data type, are considered perfectly matched, with a score of 1.
2. If two attributes (of two concepts from two ontologies) have the same name but are of different data types, we assign them a score of k ($k < 1$, $k \gg 0$).

3. Then we compute the ratio of matched attribute scores divided by the number of attributes of the concept that has more attributes.
4. The final decision about similarity is made, based on a minimum threshold for the computed combined score.

4.2.3 Ranking Candidates by Relationships using Propagation

In the previous step we have established matches between concepts from two different ontologies, based on pairs of terms and attributes. However, two concepts that point to exactly the same concepts with the same relationships are presumably very similar to each other. We view the relationship targets as data types, and two concepts that point to all the same data types are likely to be quite similar. However, we would have a chicken and egg problem here, if we start with considering relationships from the beginning. That is the case because the relationships targets cannot be used for matching if they themselves have not been matched up.

This is why we start by matching up a few concepts using terms and attributes alone. By this step, we create an initialization for matching up additional concepts by using relationships. Thus, two concepts with different names that point to several target concepts that all have been matched up between two ontologies are presumably themselves a match. We can use a similar ratio criterion as for attributes, however, now the targets carry more semantics than the undifferentiated data types of attributes. Thus, we are willing to assign a pair of relationships a high score if the targets are the same OR if the relationship names are the same. Let us assume now that a set of concept pairs has been established such that the concepts in each pair match and are from two different ontologies. Then any pair of concepts that point to these matched concepts would also be considered highly ranked for being matches. Thus, after establishing initial matches, we continue ranking concepts by similarity using a process similar to a Waltz filtering [57].

Thus, the process of finding matches needs to be recomputed until a score change of one concept pair does not result in a score change of any concepts pointing to that pair anymore. Note that this state of equilibrium can be easily achieved, as we are using a threshold. If there are only changes that do not cross the threshold, the update process would terminate.

4.2.4 Combining Matching Scores

Two concepts are considered matched if their terms, their attributes and their relationships are (on average) similar. A weight is assigned to each similarity aspect of a concept (term similarity, average attribute similarity, average relationship similarity). Considering these three criteria, we now compute the degree of the similarity of concepts from two distinct ontologies. For this purpose, we use a Multiple Attribute Decision Making (MADM) approach, a simple additive weight-based solution [18]. This approach determines a combined score of concept matches between ontologies. Let $C_i = \{C_{i1}, C_{i2}, \dots, C_{im}\}$ and $C_j = \{C_{j1}, C_{j2}, \dots, C_{jn}\}$ be sets of concepts for given ontologies, and let $F = \{F_1, F_2, \dots, F_p\}$ be a set of p features (in this paper $p = 3$) that characterize the degree of similarity. The weight vector W reflects the importance of each attribute $W = \{W_1, W_2, \dots, W_p\}$, where $\sum W_i = 1$. We compute the scores for each of the p features for each of l matching cases ($l \gg n$ or m) in a decision matrix $D = d_{ij}$.

The method is comprised of three steps: first, scale the scores into a range $[0, 1]$, with the best score represented by 1, using

$$r_{ij} = (d_{ij} - d_{jmin}) / (d_{jmax} - d_{jmin}) \quad (9)$$

Second, apply weights and third, sum up the values for each of the alternatives, using

$$S_i = \sum W_j r_{ij} \quad (10)$$

After a combined score has been computed, we compare the weighted sum with a given threshold α . Some matches may be lacking attributes or relationships. In this case, a weight of zero will be assigned to these aspects of a concept. All combined similarity values greater than α are stored in a matrix G_T . Subsequently, concept pairs with similarity values above the threshold are constructed, starting with the

maximal similarity value. If there are several equal maximal similarity values, they are processed in random order. Whenever the next largest similarity value has been identified between two concepts c and d , then the complete row of c and the complete column of d in the similarity matrix G_T are set to 0. This is because c and d are not available for matching anymore.

4.2.5 The SEMINT Algorithm

Thus, the algorithm for matching concepts from TKB' with concepts from TKB'_2 , assigned to the same semantic types, follows below. The three output lists L_1 , L_2 , and L_3 correspond to pairs of matched concepts (L_1) and unmatched concepts of TKB' and TKB'_2 (L_2 and L_3 , respectively). The algorithm makes use of a matrix G_T that is dynamically resized at every iteration of the main loop.

Algorithm: SEMINT

(Input: TKB' , TKB'_2 ; Output: L_1 , L_2 , L_3)

$L_1 = L_2 = L_3 = \{\}$;

FOR ALL semantic types $S \in \hat{S}'_2$ {

// Note that $\hat{S}'$ (SemanticTypeSet(TKB')) = $\hat{S}'_2$ (SemanticTypeSet(TKB'_2))

$G_T[,]$; // Two dimensional array of floats for similarity values;

// Cardinality of CONCEPT(S , TKB') defines number of rows.

// Cardinality of CONCEPT(S , TKB'_2) defines number of columns.

FOR ALL pairs of concepts (c , d),

$c \in \text{CONCEPT}(S, TKB')$,

$d \in \text{CONCEPT}(S, TKB'_2)$ {

B = Bigram similarity of the name of c and the name of d ;

A = Attribute similarity of c and of d ;

R = Relationship similarity of c and of d ;

// Combine these three similarities using weighted average for given weights W_i , W_j , W_k .

$G = \text{MADM}(F, W)$ where $F = B, A, R$ and $W = W_i, W_j, W_k$

IF ($G > \alpha$) { // α is a threshold value

$G_T[c, d] = G$;

} ELSE {

$G_T[c, d] = 0$;

}

}

// Loop as long as there is a value greater than 0 ANYWHERE in the 2-D array G_T .

// The WHILE loop hides a 2-D search.

WHILE there are values greater than 0 in $G_T[,]$ {

// The line below also involves a 2-D search.

Find the maximum value in $G_T[,]$;

Say, this maximum value is at (c , d)

Store (c , d) as a match in the list L_1 ;

Set all $G_T[c,] = 0$ // Remove c from consideration

Set all $G_T[, d] = 0$ // Remove d from consideration

}

Append all remaining elements of CONCEPT(S , TKB') to the list L_2 ;

Append all remaining elements of CONCEPT(S , TKB'_2) to the list L_3 ;

}

The above algorithm performs a matching operation on pairs of concepts. The structure of the algorithm is such that two concepts are considered matched if their names, their attribute names and their relationship names are (on average) similar. Inside of the algorithm, attribute similarity and relationship similarity are computed by subalgorithms (not shown) that perform matching operations on pairs of attribute names and

relationship names, respectively. The attribute similarity is the average of all the bigram similarities of all the matched attribute names. Attributes that don't match are ignored. We can use exactly the same way to compute the relationship name similarity for given pairs of concepts. Thus, the two subalgorithms are structurally similar to the algorithm SEMINT itself. A weight is assigned to each similarity aspect of a concept (name, attribute, and relationship similarity). Then we compare the weighted sum G with a given threshold α . For ontologies that lack attributes or relationships, a weight of zero is assigned to those aspects of similarity.

All combined similarity values greater than α are stored in a matrix G_T . Subsequently, concept pairs with similarity values above the threshold are constructed, starting with the maximal similarity value. If there are several equal maximal similarity values, they are processed in random order. Whenever the next largest similarity value has been identified between two concepts c and d , then the complete row of c and the complete column of d in the similarity matrix G_T are set to 0. This is because c and d are not available for matching anymore. To demonstrate the functioning of the above algorithm, assume that $\text{CONCEPT}(S, \text{TKB}') = \{a, b, c\}$ and $\text{CONCEPT}(S, \text{TKB}'_2) = \{d, e\}$. Assume that G_T would look as follows (numbers are "random" but above the given threshold):

G_T	d	e
a	0.8	0.9
b	0.7	0.75
c	0.95	0.85

Stage 1

G_T	d	e
a	0	0.9
b	0	0.75
c	0	0

Stage 2

G_T	d	e
a	0	0
b	0	0
c	0	0

Stage 3

In Stage 1 G_T contains all values that are above α . First, we append (c, d) to L_1 , giving $L_1 = \{(c, d)\}$. Then, we set the row of c and the column of d to 0 (Stage 2). Next, we create the pair (a, e) and append it to L_1 , giving $L_1 = \{(c, d), (a, e)\}$. Now, the row of a and the column of e have to be set to zero (Stage 3). At this point, all values in G_T are zero. The concept c was not matched with any other concept and is appended to L_2 . $L_2 = \{c\}$. $L_3 = \{\}$ stays empty for this semantic type. The algorithm would now advance to the next semantic type.

At this point, we have reached the limit of what can be done algorithmically. A human expert will need to review all matches in L_1 , or only those with a relatively small G value (greater than α). Similarly, he will need to review L_2 and L_3 to look for any missed matches or misclassifications. By our matching philosophy, we are trying to keep L_2 and L_3 small, limiting the manual effort that goes into this step.

To repeat one more time the effect of ALSER, consider a term like "date," which can refer to (at least) a calendar date, a sweet fruit or a romantic meeting between two people. If the assignment to semantic types was done properly, then SEMINT will not attempt to match date (the fruit) with calendar date.

Our approach to ontology integration simplifies the matching task by identifying sets of semantically similar concepts before starting with the actual matching steps. Terms from two ontologies only need to be compared for integration if they are already classified as semantically similar. Our methodology has an additional advantage. It reduces the computational cost of the matching operations. Fewer pairs of terms have to be matched against each other. For more details on the Terminological Knowledge Base Framework see [12, 14].

5 Implementation and Experimental Results

5.1 Implementation Architecture

We have implemented an ALSER/SEMINT prototype system [50] following the paradigm of component-oriented development [9]. The component-based development approach allows a complex system to be considered as a composition of an arbitrary number of smaller components with well-defined interfaces. Our system architecture is shown in Figure 6. The UI manager handles a user's requests by invoking an appropriate subsequent component. The two main components of the system are the ALSER component




Figure 6: The Architecture

and the SEMINT component which have been extensively described above.




Figure 7: The Data Flow of the ALSER/SEMINT Framework

ALSER and SEMINT expect as input TKBs coded in XML. We are using XML, because it allows us to quickly extract data and exchange information between components. Unfortunately, most existing terminologies are not in that format. The *SEMENCH* component, composed of three subcomponents (XML Reader, XML Converter, and Ontology Mapper), performs the required translation of the input. If the input format is not already XML, then the ontological input data has to be transformed into XML, using the *XML Converter*. Then the *XML Reader* component is invoked. The XML Reader component extracts concepts and their corresponding semantic types from the XML input. The XML Reader is implemented using JAVA SAX [45].

Secondly, the given terminology or ontology has to be transformed into a TKB. The only preexisting knowledge base that strictly follows the TKB format is the UMLS. We performed experiments with two terminologies from the medical domain, namely the American College of Cardiology terminology (ACC) and the Society of Thoracic Surgery terminology (STS). Both those medical terminologies are not structured like TKBs. Thus, they first had to be transformed into TKBs. Before discussing the *Ontology Mapper* we briefly describe those two ontologies.

Case	STS Concept	ACC Concept	After Data Enrichment
Synonym	<i>Date of Birth</i>	Patient <i>DOB</i>	Date of Birth
	<i>Readmission Reason</i>	<i>Readmit Reason</i>	Readmission Reason
Acronym	<i>RF-Diabetes</i>	Diabetes	Risk Factor:Diabetes
Redundant word	MI	<i>Previous MI</i>	Myocardial Infarction
	Patient ID	<i>Unique Patient ID</i>	Patient ID
	Payor	Insurance Payor	Payor
Compound word	Comps-Op-ReOp Bleed/Tamponade	Tamponade	Tamponade

Table 3: Some Examples of Actual Mapping Before/After Data Enrichment

The American College of Cardiology has provided a list of 142 terms with definitions. These concepts are separated into 44 “categories.” The Society of Thoracic Surgery has created a classification of 248 terms, subdivided into 32 categories. Regretfully, the categories are not always assigned consistently. Thus, the optimistic assumption that “term IS-A category” holds cannot be made. For example, in many cases a term describes an attribute of a category. More details about the ACC and the STS follow in the next subsection.

The *Ontology Mapper* transforms terminologies into TKBs using wrappers. In our case, three wrappers are needed, the *ACC Wrapper*, the *STS Wrapper*, and the *UMLS Wrapper*. As we will explain in the next subsection, the Ontology Mapper performs the “semantic enrichment” of the two local terminologies (ACC and STS) through a global ontology, the UMLS. The ACC Wrapper and the STS Wrapper directly access their respective terminologies. The UMLS Wrapper component communicates with the Unified Medical Language System Knowledge Source (UMLSKS) server [56]. It takes concepts as input and returns corresponding UMLS semantic types. We have found that implementing the UMLS Wrapper is difficult, as it requires Natural Language Processing. Thus, the results of the UMLS Wrapper needed to be hand-checked by a human. If the UMLS Wrapper did not find any semantic type, then the human expert found one by looking up the concept on the UMLSKS server. Secondly, if the human expert judged that the semantic type was incorrect, she looked it up on the UMLSKS server. If the semantic type appeared correct, it was left alone. Thus, in this paper, the task of the UMLS Wrapper is augmented by a human. We have collected observations how the human performs this task and intend to improve the UMLS Wrapper in future research.

The UMLSKS server offers several options. We are using “advanced search” with approximate matching. Terms are derived from all source vocabularies in the UMLS 2003AA. Using semantic enrichment, two Terminological Knowledge Bases are generated, encapsulating the ACC and STS terminologies, respectively. The TKBs generated by SEMENCH, refined by ALSER and/or integrated by SEMINT are stored in the *TKB Repository* for future use. The ALSER component needs to be invoked for both TKBs. SEMINT is then applied to the two refined TKBs, resulting from the application of ALSER (Figure 7).

5.2 Experimental Results

In order to test our algorithms, we have performed extensive experimental work. As mentioned above, we used the UMLS for enriching the terminologies of the American College of Cardiology (ACC) [1] and the Society of Thoracic Surgeons (STS) [48]. While the goals and sources of the STS and ACC are different, there is some degree of overlap. Table 3 demonstrates the necessity of data enrichment. It shows matches of concepts in the ACC and the STS which became evident only after applying data enrichment to the terms in the table. Table 4 shows part of the STS terminology before enrichment and the enriched STS terminologies. The symbol $\cap$ in Table 5 indicates that a concept belongs to all the semantic types connected by $\cap$, i.e., the intersection type. Table 5 shows how concepts have been enriched through Semantic Enrichment by showing their categories and semantic types of STS and ACC. The STS terminologies contains 248 concepts and 32 categories. In the enriched terminology there are only 244 concepts, as we did not find semantic types for 4 concepts. The ACC terminology contains 142 concepts and 44 categories. The enriched ACC contains all 142 terms. The ACC and STS have neither attributes nor relationships. Thus, our integration is based on the names of concepts only.

Category Name	Concept Name	Semantic Type	Concept Name
Post Operative	Initial ICU hours	Temporal Concept	Initial ICU hours
Hospitalization	Additional ICU Hours Total Hours ICU		Additional ICU Hours Total Hours ICU
History and Risk Factors	Weight, Height	Organism Attribute $\cap$ Quantitative Concept	Weight, Height
	Risk Factor:Smoker, Risk Factor:Smoker Current	Population Group $\cap$ Finding $\cap$ Quantitative Concept	Smoker Smoker Current
Operative	Urgent Reason Emergent Reason	Idea or Concept	Urgent Reason Emergent Reason
Diagnostic cath procedure-findings	Valve Disease:Insufficient Aortic, Valve Disease:Insufficient Mitral	Sign or Symptom	Insufficient Aortic Insufficient Mitral
Operative	Valve Surgery:Aortic Procedure, Valve Surgery:Mitral Procedure	Therapeutic or Preventive Procedure	Aortic Procedure Mitral Procedure

Table 4: Partial STS Terminologies before/after Semantic Enrichment

Concept Name	Category (Before SE)	Semantic Type (After SE)
Initial ICU hours	Post Operative	Temporal Concept
Additional ICU Hours	Hospitalization	Temporal Concept
Readmission to ICU	Hospitalization	Health Care Activity
Date of Birth	Demographics	Finding
Weight	History and Risk Factors	Organism Attribute $\cap$ Quantitative Concept
Smoker	History and Risk Factors	Population Group $\cap$ Finding $\cap$ Quantitative Concept
Patient ID	Administrative	Idea or Concept
Urgent Reason	Operative	Idea or Concept
Patient Last Name	Demographics	Intellectual Product
Insufficient Aortic	Diagnostic cath procedure-findings	Sign or Symptom
Aortic Procedure	Operative	Therapeutic or Preventive Procedure

Table 5: Semantic Assignment in Partial STS Terminologies before/after Semantic Enrichment

In Table 6, we characterize the assignment of concepts to semantic types before and after running the ALSER algorithm. The left part of the table shows the expected increase in the number of semantic types due to ALSER, from 35 to 54 (for ACC) and from 38 to 68 for STS. Of these 54, 20 are intersection types, 26 simple semantic types and 8 empty semantic types. The right part shows the distribution of concepts over semantic types before and after running ALSER. As expected, the number of concepts assigned to each semantic type is reduced for minimum, maximum and average.

Table 7 (left part) shows results for running SEMINT with similarity threshold, α (90%). In an integrated ontology, every matched concept will occur only once. This is shown in the last two rows of the table. The right part of Table 7 shows more properties of the integrated TKB. The number of simple semantic types is 27. Looking back at the numbers of simple semantic types in Table 6, they are both 26. This indicates a high degree of overlap between the semantic types of ACC and of STS. There are 25 common semantic types, with only one semantic type that is unique to ACC (and to STS, respectively).

5.3 Recall and Precision

We have performed a series of experiments to test the effectiveness of ALSER with SEMINT. We first computed a desired matching result between ACC and STS by hand. Then we used *recall* and *precision* to compare the results of ALSER with SEMINT with our manual result. We also compared with a brute force approach that matches every concept from ACC with every concept from STS. For evaluation purposes, we

STS Concept	STS Category	ACC Concept	ACC Category	Semantic Type
Cardiogenic Shock	Pre Operative	Cardiogenic Shock	Diagnostic Cath	Disease or Syndrome
Arrhythmia	Cardiac Status	Arrhythmia	Procedure - Indications	$\cap$ Finding
MI	Pre Operative	Previous MI	History and Risk Factors	Disease or Syndrome
Meds-Asprin	Cardiac Status			
	Pre Operative Medications	Asprin	Cath Lab Visit - Medications	Organic Chemical $\cap$ Pharmacologic Substance
VD-Stenosis-Aortic	Pre & Cath	Valve Disease - Aortic	Diagnostic Cath	Disease or Syndrome $\cap$
VD-Stenosis-Mitral	Operative Hemodynamics	Valve Disease - Mitral	Procedure - Findings	Finding
RF-Renal Fail	Pre-Operative Risk Factors	Renal Failure	Adverse Outcomes	Disease or Syndrome $\cap$ Finding

Table 6: STS and ACC Concepts and Categories (Before Semantic Enrichment) and Their Semantic Types (After Semantic Enrichment)

Feature	ACC	STS
Concepts#	142	244
Simple Concepts	93	173
Compound Concepts	49	71
Original Semantic Types	35	38
Refined Semantic Types	54	68
Intersection Types	20	32
Simple Semantic Types	26	26
Empty Semantic Types	8	10

	ACC		STS	
Assigned to a Semantic Type	Bef.	Aft.	Bef.	Aft.
Min# of Concepts	1	0	1	0
Max# of Concepts	53	15	58	33
Avg.# of Concepts	3	2	5	3
Total# of Assignments	192	142	325	244

Table 7: Statistics of the Effect of ALSER on ACC and STS

use the following common terms:

True Positives (TP): The matching algorithm reports a match of two terms, and those two terms are considered a match according to a human expert.

True Negatives (TN): The matching algorithm does not report a match of two terms, and it should not report it, according to a human expert.

False Positives (FP): The matching algorithm reports a match of two terms, but a human expert would not consider those two terms a match.

False Negatives (FN): The matching algorithm does not report two terms as matched, but they should be matched, according to a human expert.

We can now use Equation 11 to compute the precision (P) and Equation 12 to compute the recall (R) of our experiment. Low recall means that our algorithm is missing many pairs that it should report, according to our human expert. Low precision means that our algorithm is reporting many pairs that it should not report. An additional measure that is commonly used in Information Retrieval is Van Rijsbergen’s F measure:

$$P = \frac{TP}{TP + FP} \quad (11)$$

$$R = \frac{TP}{TP + FN} \quad (12)$$

$$F = \frac{2 * R * P}{R + P} \quad (13)$$

Table 8 shows the threshold α , P , R , F and runtime for ALSER with SEMINT and for brute force, given three different threshold values. Using ALSER/SEMINT, increasing the threshold to 90% leads to better precision, without deterioration of recall. Brute Force is also better for $\alpha = 90\%$, as shown by the F value, although recall is slightly better for lower α . Thus, we concentrate on comparing the 90% entries in the table, which have the same P and R , but the run time is about 30 times faster for ALSER with SEMINT.

Feature	Thresh. $\alpha > 90$
Matched Concepts	37
ACC Unmatched Concepts #	105
STS Unmatched Concepts #	207
Simple Concepts #	349
Total # of Ass. to Semantic Types	349

Features	Number
Simple Semantic Types #	27
Intersection Types #	41
Empty Semantic Types #	15
Total Sem. Types	83
Minimum # of Concepts assigned to a Semantic Type	0
Maximum # of Concepts assigned to a Semantic Type	33
Average # of Concepts assigned to a Semantic Type	3

Table 8: (a) Impact of threshold matches (left); (b) Integrated terminology (right)

Method	Threshold	Precision	Recall	F Measure	Time (in ms)
ALSER with SEMINT	> 70 %	0.925	0.8604	0.8915	551
ALSER with SEMINT	> 80 %	0.9736	0.8604	0.9135	551
ALSER with SEMINT	> 90 %	1	0.8604	0.925	551
Brute Force	> 70 %	0.7916	0.8837	0.8351	16543
Brute Force	> 80 %	0.9268	0.8837	0.9047	16543
Brute Force	> 90 %	1	0.8604	0.925	16543

Table 9: Comparison of ALSER/SEMINT with brute force matching

In future work, we are expecting better precision and F results for the ALSER with SEMINT method, hopefully without degradation of recall, because the current implementation does not make use of attributes and relationships, as these are not available for the ACC and the STS. Thus, matching relies on a very weak method. We will therefore investigate the use of better test bed ontologies. With larger ontologies, the runtime differences become more important.

6 Related Work

Many lines of research have addressed ontology matching in the context of ontology construction and integration [6, 32, 35, 34]. The major goal of these approaches is to develop effective methodologies for automated mappings [27]. A linchpin of the ALSER methodology is that it defines a notion of similarity. There is an extensive literature on similarity measures in ontologies, which can be effectively used in ontology integration.

As a major line of similarity measurement research, a number of linguistics researchers have attempted to make use of conceptual distance in information retrieval. These approaches mainly focus on a measure of semantic relatedness using WordNet and medical terminologies such as MeSH, UMLS, etc. Although these terminologies are quite thorough in their coverage of concepts, the number of semantic relation types connecting these concepts is considerably limited for advanced feature-based similarity models. Also relations in terminologies tend to be lexical relations between words, which do not always reflect ontological relations between classes of entities of the world. Tversky’s feature-based similarity measure [55] is one of the most powerful similarity models, but it requires a rich knowledge base. A number of directions based on Tversky’s work have been explored. Most approaches [39, 21, 23] have limited their attention to IS-A links and *broader term* and *narrower term* links.

A similarity measure using edge-counting [39] relies on decomposing concepts by explicit features and inherited features and ranks concepts by counting feature differences. Lee *et al.* [23] extended this method by including the links between synsets or paths made up of meronym type links.

When using the edge distance between two concepts, the relative depth is important, because research [26] has shown that the distance between concepts shrinks as one descends down the ontology. Two siblings near the top of the ontology are conceptually far apart, compared to two siblings deep down in the ontology. Relative depth was taken into consideration by [51]. Another important similarity issue is related to ontology

density. Concept distances in denser regions should be smaller than distances in less dense regions. Hirst and St-Onge [15] also consider the direction and stability of links between concepts. In the HCG weighted conceptual distance methodology [2] the strength of connotation between parent and child nodes was considered together with ontology density and depth. Some research has combined thesauri with a corpus statistics for semantic similarity measurement [43, 24, 20].

Conceptual similarity is considered in terms of class similarity by Resnik [42]. In this approach, the extent of shared information between concepts was considered to determine their similarity. Lin [24] proposed an information-theoretic notion of similarity based on the joint distribution of properties. They accounted not only for commonalities but also for differences between the items being compared. In [4] a subclass must be of the same type as its parent, but must have some difference that distinguishes it from its parent. The subclasses of a concept are incompatible with each other. Jiang and Conrath’s similarity measurement [20] is based on the conditional probability of encountering a child synonym set given a parent synonym set.

More recently, a number of new similarity approaches has been introduced. One approach is the use of matching rules [34, 35, 49]. Another method compares the set of all possible correspondences [5, 33, 22, 8]. The names of the concepts or nesting relationships between concepts and the inter-relationships between concepts (slots of frames in [35]) are also criteria for comparison. The types of the concepts, or the labeled graph structure of the models [5, 33] may be used to estimate the likelihood of data instance correspondences [7, 22, 49, 3]. Rodriguez and Egenhofer [44] proposed computing semantic similarity for different ontologies from three perspectives (1) synonym set matching, (2) semantic neighborhood, measured by the shortest path between connected concepts, and (3) distinguishing features. Like in our approach, these three aspects are combined, using a weighted sum function. However, our approach is substantially different because of our use of semantic refinement prior to any matching. Furthermore, we do not use synonyms and distances; our term matching approach is based on bigrams.

Some similarity approaches [32, 35] allow for efficient user interaction or expressive rule languages [6] for specifying mappings. Several recent publications have attempted to further automate the ontology matching process. A general heuristic was used in [36] to show that paths between matching elements tend to contain other matching elements. In [22], similarity between two nodes was computed based on their signature vectors, which were computed from data instances. The above approaches argue for a single best universal similarity measure, whereas GLUE [8] allows for application-dependent similarity measures.

A methodology for learning to construct ontologies from other ontologies is described in [27]. The similarity measure in [19] is based on statistics, and can be thought of as being defined over the joint probability distribution of the concepts involved. The role of machine learning in the Semantic Web effort was described in [28]. [8] is another example of ontology integration using machine learning.

We now return to the problem of combining multiple ontologies into a single coherent ontology [35, 49]. There are a few different methodologies but there is no consensus on the methodology to follow to integrate ontologies either in the same or in different domains. It is not even clear whether there could be a domain independent methodology. When the semantics of an ontology integration system (whose sources are described by different ontologies) are defined, a global mediated ontology [5] might be required, whereas in some frameworks [52] a mediated ontology will exist only if it makes sense for the task at hand.

Pinto and Martins [38] addressed important issues of ontology engineering and ontology composition. Gangemi *et al.* [10] developed a domain ontology by the integration of existing repositories of knowledge, analyzing and selecting the relevant sets of terms from various terminological sources. While [10] starts with all ontologies at the beginning of the merge process, Skuce [47] starts with a selected initial group of ontologies that is incrementally enlarged. The work of [5] describes a specialization of a framework for ontology integration.

7 Conclusion and Future Work

Matching concepts from two different medical ontologies is a difficult task, which will become more and more important with the expected advancement of the Medical Informatics. Our approach to concept matching

is based on an ontology representation with a Terminological Knowledge Base which makes it possible to identify concepts with similar semantics before attempting the difficult matching task. The semantic enrichment method builds Terminological Knowledge Bases and the Algorithmic Semantic Refinement (ALSER) algorithm generates semantically uniform Terminological Knowledge Bases, for which it is possible to define a formal sense of semantic similarity between concepts.

Our approach is characterized by a philosophical assumption about future ontologies which will need to be integrated. Any ontology that will be important enough to be used at all will need to be “reasonably” complete in at least one subarea of a domain. Thus, if two ontologies cover a domain, one may expect that all the “important” concepts of the domain occur in both ontologies. Thus, a matching algorithm should attempt to maximize the number of matching concepts, even if there are structural differences between them. The Semantic Integration (SEMINT) algorithm performs this kind of matching on semantically uniform Terminological Knowledge Bases.

We presented an implementation of ALSER with SEMINT, applying it to the ACC and STS terminologies. Results of the implementation application were expressed in terms of precision, recall and F measure. ALSER with SEMINT greatly improves the run time of matching. It has a positive effect on precision for low thresholds and a negative effect on recall. Overall, as the F value shows, the improvement of precision outweighs the effect on recall.

This paper presents the beginning of a journey, not the end. At this point we discuss some limitations of our approach. For research on integration, two TKBs in the same domain are necessary, which have to be built by independent teams, to avoid any biases that would make integration unfairly easier. With increasing growth of ontologies, we hope that more and more medical ontologies available are available, which can be used in our research. For the case when a global TKB ontology (i.e., UMLS) is available, we have presented a method for semantic enrichment to create such two-level structures for two medical terminologies of STS and ACC. In general, there could exist some newly generated or empty semantic types as the result of ALSER. For the case of the larger scope of the problem or the involvement of many ontologies, we need to adopt a more systematic approach toward handling some new changes in TKBs. Our current testbed ontologies have no attributes and no relationships. We expect better results by adding these features. Also, our method relies on a threshold value and several weights. Tuning these values and analyzing their effect on precision and recall of SEMINT is a topic of future research. Lastly, too little is known about TKBs themselves. We will investigate structural parameters of TKBs, foremost the ratio of the size of the Semantic Network relative to the size of the thesaurus, to see what influence they have on running the ALSER and SEMINT algorithms.

References

- [1] American college of cardiology website. <http://www.acc.org/>.
- [2] E. Agirre and G. Rigau. Word sense disambiguation using conceptual density. In *Proceedings of the 16th International Conference on Computational Linguistics*, pages 16–22, 1996.
- [3] J. Berlin and A. Motro. Database schema matching using machine learning with feature selection. In *Proc. of the 14th Int. Conf. on Advanced Information Systems Engineering (CAiSE02)*, 2002.
- [4] J. Bouaud, B. Bachimont, J. Charlet, and P. Zweigenbaum. Methodological principles for structuring an ontology. In *IJCAI-95 Workshop on Basic Ontological Issues in Knowledge Sharing*, 1995.
- [5] D. Calvanese, D. G. Giuseppe, and M. Lenzerini. Ontology of integration and integration of ontologies. In *Proceedings of the 2001 Description Logic Workshop (DL2001)*, 2001.
- [6] H. Chalupsky. Ontomorph: A translation system for symbolic knowledge. In *Principles of Knowledge Representation and Reasoning*. Morgan Kaufmann, 2000.

- [7] A. Doan, P. Domingos, and A. Halevy. Reconciling schemas of disparate data sources: A machine-learning approach. In *Proceedings of SIGMOD*, pages 509–520, 2001.
- [8] A. Doan, J. Madhavan, P. Domingos, and A. Halevy. Learning to map between ontologies on the Semantic Web. In *Proceedings of WWW-2002*, 2002.
- [9] D.F. D’Souza and A.C. Wills. *Objects, Components, and Frameworks with UML: the Catalysis Approach*. Addison Wesley, 1999.
- [10] A. Gangemi, G. Steve, and F. Giacomelli. ONIONS: An ontological methodology for taxonomic knowledge integration. In *ECAI Workshop on Ontological Engineering*, 1996.
- [11] J. Geller, H. Gu, Y. Perl, and M. Halper. Semantic refinement and error correction in large terminological knowledge bases. *Data & Knowledge Engineering*, 45(1):1–32, 2003.
- [12] J. Geller, H. Gu, Y. Perl, and M. Halper. Semantic refinement and error correction in large terminological knowledge bases. *Data & Knowledge Engineering*, 45(1):1–32, 2003.
- [13] T. R. Gruber. Toward principles for the design of ontologies used for knowledge sharing. In N. Guarino and R. Poli, editors, *Proceedings of International Workshop on Formal Ontology*, Padova, Italy, 1993.
- [14] H. Gu, Y. Perl, J. Geller, M. Halper, L. Liu, and J. J. Cimino. Representing the UMLS as an OODB: Modeling issues and advantages. *Journal of the American Medical Informatics Association*, 7(1):66–80, 2000.
- [15] G. Hirst and D. St-Onge. Lexical chains as representations of context for the detection and correction of malapropisms. In M. Hoey, editor, *Patterns of Lexis in Text*, pages 305–332. Oxford University Press, 1998.
- [16] B. L. Humphreys and D. A. B. Lindberg. Building the Unified Medical Language System. In *Proc. Thirteenth Annual Symposium on Computer Applications in Medical Care*, pages 475–480, Washington, DC, November 1989.
- [17] B. L. Humphreys, D. A. B. Lindberg, H. M. Schoolman, and G. O. Barnett. The Unified Medical Language System: An informatics research collaboration. *JAMIA*, 5(1):1–11, 1998.
- [18] C. L. Hwang and K. Yoon. *Multiple Attribute Decision Making*. Springer-Verlag, Berlin, 1981.
- [19] R. Ichise, H. Takeda, and S. Honiden. Rule induction for concept hierarchy alignment. In *Proceedings of the Workshop on Ontology Learning at the 17th International Joint Conference on Artificial Intelligence (IJCAI)*, 2001.
- [20] J. Jiang and D. W. Conrath. Semantic similarity based on corpus statistics and lexical taxonomy. In *Proceedings of International Conference on Research in Computational Linguistics*, 1997.
- [21] Y. W. Kim and H. Kim. A model of knowledge based information retrieval with hierarchical concept graph. *Journal of Documentation*, 46(2):113–137, 1990.
- [22] M. Lacher and G. Groh. Facilitating the exchange of explicit knowledge through ontology mappings. In *Proceedings of the 14th Int. FLAIRS conference*, 2001.
- [23] J. H. Lee, M. H. Kim, and Y. J. Lee. Information retrieval based on conceptual distance in IS-A hierarchies. *Journal of Documentation*, 49(2):188–207, 1993.
- [24] D. Lin. An information-theoretic definition of similarity. In *Proceedings of the 15th International Conference on Machine Learning*, 1998.

- [25] D. A. B. Lindberg, B. L. Humphreys, and A. T. McCray. The Unified Medical Language System. *Methods of Information in Medicine*, 32:281–291, 1993.
- [26] W. Lord, R.D. Stevens, A. Brass, and C.A. Goble. Investigating semantic similarity measures across the gene ontology: the relationship between sequence and annotation. *Bioinformatics*, 2002. In press.
- [27] A. Maedche. A machine learning perspective for the Semantic Web. In *Proceedings of Semantic Web Working Symposium (SWWS)*, 2001.
- [28] A. Maedche and S. Saab. Ontology learning for the Semantic Web. *IEEE Intelligent Systems*, 16(2), 2001.
- [29] A. T. McCray. UMLS semantic network. In *Proceedings of the Thirteenth Annual SCAMC*, pages 503–507, 1989.
- [30] A. T. McCray. Representing biomedical knowledge in the UMLS semantic network. In N. C. Broering, editor, *High-Performance Medical Libraries: Advances in Information Management for the Virtual Era*, pages 45–55. Meckler, Westport, CT, 1993.
- [31] A. T. McCray and W. T. Hole. The scope and structure of the first version of the UMLS semantic network. In *Proceedings of the Fourteenth Annual SCAMC*, pages 126–130, 1990.
- [32] D. McGuinness, R. Fikes, J. Rice, and S. Wilder. The CHIMAERA ontology environment. In *Proceedings of the 17th National Conference on Artificial Intelligence*, 2000.
- [33] S. Melnik, H. Molina-Garcia, and E. Rahm. Similarity flooding: A versatile graph matching algorithm. In *Proceedings of the International Conference on Data Engineering (ICDE)*, 2002.
- [34] P. Mitra, G. Wiederhold, and J. Jannink. Semi-automatic integration of knowledge sources. In *Proceedings of Fusion’99*, 1999.
- [35] N. Noy and M. Musen. Prompt: Algorithm and tool for automated ontology merging and alignment. In *Proceedings of the National Conference on Artificial Intelligence*, 2000.
- [36] N. Noy and M. Musen. Anchor-PROMPT: Using non-local context for semantic matching. In *Proceedings of the Workshop on Ontologies and Information Sharing at the International Joint Conference on Artificial Intelligence (IJCAI)*, 2001.
- [37] Y. Peng, M. Halper, Y. Perl, and J. Geller. Auditing the umls for redundant classifications. In *Proceedings of AMIA 2002*, pages 612–616, San Antonio, TX,, 2002.
- [38] H. Sofia Pinto and J.P. Martins. A methodology for ontology integration. In *Proceedings of the First International Conference on Knowledge Capture (K-CAP2001)*, 2001.
- [39] R. Rada, H. Mili, Ellen Bicknell, and Maria Blettner. Development and application of a metric on semantic nets. *IEEE Transactions on Systems, Man, and Cybernetics*, 19(1):17–30, 1989.
- [40] A. Rector. Clinical terminology: Why is it so hard? *Methods of Information in Medicine*, 38:239 – 252, 1999.
- [41] A. Rector, S. Bechhofer, C. Goble, I. Horrocks, W. Nowlan, and W. Solomon. The GRAIL concept modelling language for medical terminology. *Artificial Intelligence in Medicine*, 9:139–171, 1997.
- [42] P. Resnik. *Selection and Information : A Class based Approach to Lexical Relationships*. PhD thesis, University of Pennsylvania, 1993.
- [43] P. Resnik. Using information content to evaluate semantic similarity. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, pages 448–453, 1995.

- [44] M. A. Rodriguez and M. J. Egenhofer. Determining semantic similarity among entity classes from different ontologies. *IEEE Transactions on Knowledge and Data Engineering*, 15(2):442 – 456, 2003.
- [45] SAX website. <http://www.saxproject.org/>.
- [46] P. L. Schuyler, W. T. Hole, M. S. Tuttle, and D. D. Sherertz. The UMLS Metathesaurus: Representing different views of biomedical concepts. *Bull Med Libr Assoc*, 81(2):217–222, 1993.
- [47] D. Skuce. How we might reach agreement on shared ontologies: A fundamental approach. In *AAAI Spring Symposium Series, Workshop on Ontological Engineering*, 1997.
- [48] Society of thoracic surgeon website. <http://www.sts.org/>.
- [49] G. Stumme and A. Maedche. FCA-MERGE: Bottom-up merging of ontologies. In *Proc. of the 17th Int. Joint Conf. on Artificial Intelligence (IJCAI)*, 2001.
- [50] K. Supekar. ALSER/SEMINT implementation report, 2002.
- [51] M. Sussna. Word sense disambiguation for free-text indexing using a massive semantic network. In *Proceedings of the Second International Conference on Information and Knowledge Management (CIKM-93)*, pages 67–74, 1993.
- [52] W. R. Swartout, R. Patil, K. Knight, and T. Russ. Towards distributed use of large-scale ontologies. In *AAAI Spring Symposium on Ontological Engineering*, 1997.
- [53] M. Tringali, W. Hole, and S. Srinivasan. Integration of a standard gastrinintestinal endoscopy terminology in the umls metathesaurus. In *The annual symposium of the American Medical Informatics Association (AMIA-2002)*, 2002.
- [54] M. S. Tuttle, D. D. Sherertz, N. E. Olson, M. S. Erlbaum, W. D. Sperzel, L. F. Fuller, and S. J. Nelson. Using meta-1 the first version of the UMLS Metathesaurus. In *Proceedings of the Fourteenth Annual SCAMC*, pages 131–135, 1990.
- [55] A. Tversky. Features of similarity. *Psychological Review*, 84:327–352, 1977.
- [56] Unified Medical Language System Website. <http://umlsks4.nlm.nih.gov>.
- [57] D. Waltz. Understanding line drawings with shadows. In P. Winston, editor, *The Psychology of Computer Vision*, pages 19–91. McGraw Hill, New York, 1975.