

# Subjective Databases

Yuliang Li                      Aaron Feng                      Jinfeng Li  
Saran Mumick                  Alon Halevy                      Vivian Li                      Wang-Chiew Tan  
Megagon Labs  
{yuliang, aaron, jinfeng, saran, alon, vivian, wangchiew}@megagon.ai

## ABSTRACT

Online users are constantly seeking experiences, such as a hotel with clean rooms and a lively bar, or a restaurant for a romantic rendezvous. However, e-commerce search engines only support queries involving objective attributes such as location, price, and cuisine, and any experiential data is relegated to text reviews.

In order to support experiential queries, a database system needs to model *subjective* data. Users should be able to pose queries that specify subjective experiences using their own words, in addition to conditions on the usual objective attributes. This paper introduces OPINEDB, a subjective database system that addresses these challenges. We introduce a data model for subjective databases. We describe how OPINEDB translates subjective queries against the subjective database schema, which is done by matching the user query phrases to the underlying schema. We also show how the experiential conditions specified by the user can be combined and the results aggregated and ranked. We demonstrate that subjective databases satisfy user needs more effectively and accurately than alternative techniques through experiments with real data of hotel and restaurant reviews.

### PVLDB Reference Format:

Yuliang Li, Aaron Feng, Jinfeng Li, Saran Mumick, Alon Halevy, Vivian Li, Wang-Chiew Tan. Subjective Databases. *PVLDB*, 12(11): 1330-1343, 2019.

DOI: <https://doi.org/10.14778/3342263.3342271>

## 1. INTRODUCTION

Database systems model entities in a domain with a set of attributes. Typically, these attributes are objective in the sense that they have an unambiguous value for a given entity, even if the value is unknown to the database, known only probabilistically, or recorded erroneously. Typical examples of such attributes include product specifications, details of a purchase order, or values of sensor readings. The boolean nature of database query languages reinforces the primacy of objective data—a tuple is either in the answer to the query or is not, but cannot be anywhere in between.

However, the world also abounds with *subjective* attributes for which there is no unambiguous value and are of great interest to

This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>. For any use beyond those covered by this license, obtain permission by emailing [info@vlldb.org](mailto:info@vlldb.org). Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

*Proceedings of the VLDB Endowment*, Vol. 12, No. 11

ISSN 2150-8097.

DOI: <https://doi.org/10.14778/3342263.3342271>

users. Examples of such attributes occur in a variety of domains, including the cleanliness of hotel rooms, the difficulty level of an online course, or whether a restaurant is romantic. Currently, the data for these attributes, when it exists, is typically left in text reviews or social media, but not modeled in the database and therefore not queryable. As a result, making decisions that involve subjective preferences is labor intensive for the end user.

Query	Subjective	"Hotels in London of reasonable price"	"Restaurants that serve delicious food"
	Objective	"List all hotels in London <= £180 per night"	"Restaurants with avg. food_rating > 4.9"
		Objective	Subjective
		Data	

Figure 1: Objective/subjective data and queries.

Figure 1 illustrates that subjectivity can occur in both the data and queries. The lower left quadrant, where both the data and the queries are objective, represents the vast majority of databases today. The lower right quadrant represents how some subjective data has been shoehorned into database systems to date. For example, user ratings of a restaurant are stored in a numerical field, and queries are objective in that they refer to predicates or aggregates on that field (e.g., *return restaurants with an aggregate rating of more than 4.5*). As shown in the upper left hand quadrant, users can pose subjective queries even on objective data. However, in many cases a proper visualization of the answer (e.g., with a histogram or a map) suffices to help the user make their subjective judgment.

This paper introduces the OPINEDB System, a database system that explicitly models subjective data and answers subjective queries, thereby addressing the challenges in the upper right quadrant of Figure 1. With this capability, OPINEDB enables a new set of applications where users can search by their subjective preferences.

### 1.1 Example: experiential search

An important motivating application for subjective databases is *experiential search*. Today's e-commerce search engines support querying by objective attributes of a service or product, such as price, location or square footage. However, as we illustrate in Section 5.1, users overwhelmingly want to be able to search by specifying the experiences they desire, and these are often expressed as subjective predicates. Table 1 shows examples of queries that users should be able to pose.

We use the domain of hotel search to illustrate some of the challenges of experiential search. A subjective database in the hotel domain will have a schema that models hotel rooms with a mixture of objective and subjective attributes.

**Table 1: Queries that express experiential preferences.**

Domain	Example Query
Hotels	<i>a hotel with a lively bar scene and clean rooms</i>
Dining	<i>a restaurant with a sunset view of Tokyo Tower</i>
Employment	<i>a job with a dynamic team working on social good</i>
Housing	<i>a 2-bedroom apartment in a quiet neighborhood near good cafes</i>
Online education	<i>a 1-week course on python with short programming exercises</i>
Travel	<i>a relaxing trip to a beach on the Mediterranean</i>
Multi-domain	<i>a quiet Thai restaurant next to a cinema that shows Ocean's 8</i>

Consider a user who is searching for a hotel in London that costs less than 180 pounds per night, has really clean rooms, and is a romantic getaway. The first condition is objective and simple to satisfy, but the second and third conditions are subjective. Conceivably, techniques from sentiment analysis and opinion mining [31] can be used to extract relevant descriptions from hotel reviews. However, OPINEDB faces several novel challenges.

First, OPINEDB needs to aggregate the reviews in a meaningful way and so they can be queried effectively and efficiently. To do so, OPINEDB introduces the concept of *markers*, which are the distinctions about the domain that the designer thinks are important for the application. OPINEDB then aggregates phrases from the reviews along these markers to create *marker summaries*. An application based on OPINEDB may also decide to expose these markers in the user interface for easier querying. The choice of markers is based on a combination of mining the review data and knowing the requirements of the application and is crucial to the quality of the query results. For example, the designer might decide that room cleanliness can be modeled by {clean, dirty}, while bathroom style requires a scale such as {old, standard, modern, luxurious}.

Second, OPINEDB needs to answer complex queries in a principled way. In our example, OPINEDB has to combine two subjective query predicates and an objective predicate. The user may also decide to complicate the query and only consider opinions of people who reviewed at least 10 hotels. In that case, OPINEDB needs to refer back to the review data and consider only reviews by qualified reviewers, which will involve recalculating the room cleanliness scores for every hotel.

Finally, users may not always use terms that fit neatly into the database schema. For example, the user may ask for romantic hotels, but the schema does not have an attribute for romantic. However, OPINEDB may have background knowledge that suggests that hotels with exceptional service and luxurious bathrooms are often considered romantic. Since the quality of service and bathroom luxury are captured in its schema as subjective attributes, OPINEDB can reformulate the query for romantic rooms into a combination of attributes in the schema. Users may also query for properties that are not even close to the database schema, such as hotels that are good for motorcyclists. In this case, OPINEDB will verify this requirement by falling back on text search in the reviews to see if any reviews mention amenities for motorcyclists. □

The above example illustrates the broader dichotomy that exists between the search for structured objects and search for documents. As a typical example, online shoppers will typically go to a search engine like Google or Bing to find the top-rated espresso machines, and then go to Amazon to purchase one (a situation that none of these companies are happy with and are working hard to tilt in their favor). The fundamental reason for the dichotomy is that the experiential aspects of the object (or service) being purchased are not

queryable, and therefore users rely on Web search engines as the best option to discover them. This paper focuses on the core issues involved in making subjective data queryable, but ultimately these techniques can be used to extend database systems as well as document retrieval systems.

## 1.2 Contributions

Specifically, we make the following contributions:

- We introduce a data model for subjective databases and an associated variant of SQL that supports subjective predicates for querying such databases. The key feature of a subjective database is that every subjective attribute is associated with a new data type called the *linguistic domain*, which is a set of phrases for describing the attribute. The designer specifies a set of *markers* that are the phrases in the linguistic domain that represent the important concepts of the domain. The linguistic domain and markers are effective intermediaries between text and queries; They help produce meaningful aggregation from information extracted from text and help produce high-quality answers to queries.
- We present OPINEDB, our query processing system for subjective databases. OPINEDB effectively interprets subjective predicates against the subjective database schema through a combination of NLP and IR techniques; by matching against the linguistic domains and markers, or finding correlations between subjective predicates and attributes in reviews. It is also able to fall back to exploiting traditional text retrieval methods as needed. After interpretation, OPINEDB uses a variant of fuzzy logic to combine the results of multiple subjective query predicates.
- One of the major challenges that OPINEDB raises is the construction of the subjective database, that requires extracting relevant information from text and designing the subjective database schema. We developed a novel extraction pipeline which requires little NLP expertise from the schema designer and also facilitates the automatic discovery of potential markers for subjective attributes. We show that our extraction pipeline achieves state-of-the-art performance by leveraging the most recent advances in NLP techniques such as BERT [11].
- We demonstrate the effectiveness and efficiency of OPINEDB with real-world review data from two domains — hotels and restaurants. Our experimental results demonstrate the need for subjective databases, OPINEDB outperforms two search baselines by up to 15% and 10% even when evaluated conservatively, and OPINEDB achieves a speedup of up to 6.6x through the use of marker summaries for query processing.

**Outline:** We introduce our data model for subjective databases and subjective queries in Section 2. Section 3 describes query processing in OPINEDB. Section 4 describes how OPINEDB constructs a subjective database by extracting opinions from reviews and aggregating them into summaries. We present our experimental results in Section 5. We discuss related work in Section 6. We also provide more details and examples in the appendix of the full version [30].

## 2. DATA MODEL

A relation in OPINEDB includes objective and subjective attributes. Informally, a subjective attribute represents an aggregate view of textual phrases extracted from reviews. In this section we introduce *linguistic domains* that capture these phrases and *marker summaries* that represent the aggregates.

Consider the example of an attribute `room_cleanliness` in the domain of hotels. The raw data for this attribute exists in reviews and social media and consists of a wide variety of phrases such as

1. “The floor in my room was filthy dirty ...”,
2. “The room was clean, well-decorated and ...”, or
3. “Spotlessly clean and good location”.

The challenge OPINEDB faces is to *aggregate* these phrases into a meaningful signal and to rank hotels appropriately in response to queries that may themselves include different linguistic phrases. The ability to aggregate linguistic phrases is one of the key aspects that distinguishes OPINEDB from text retrieval systems where ranking is typically based on similarity of unstructured text.

To define an aggregation function, there needs to be a scale onto which the aggregation is performed. However, when dealing with text, we cannot always arrange the main landmarks in the domain into a linearly-ordered scale. Hence, OPINEDB lets the schema designer define a set of *markers* in the domain that represent the points onto which we map the reviews. In our example, the markers might be a linear scale [very\_clean, average, dirty, very\_dirty] for room cleanliness and a set of markers for bathroom styles may be [old, standard, modern, luxurious], which is a set of categories (not a linear scale) that describes the different styles of bathroom. OPINEDB maintains a *marker summary*, which is a view that aggregates the phrases from the reviews onto the markers. Specifically, OPINEDB computes a value representing the membership of each hotel for each marker. For example, the record [very\_clean: 20, average: 70, dirty: 30, very\_dirty: 10] for a hotel would represent that the hotel is closer to being a member of *average* than to the other markers. As we see later, there are different possible aggregation functions OPINEDB can use and the appropriate choice depends on the semantics of the attribute. At present, OPINEDB’s marker summaries are histograms that tabulate the number of phrases of reviews closest to each marker. Each marker summary also records features useful for query processing including the average sentiment score and the average phrase embedding vector.

In addition to aggregating the raw data, the marker summaries serve two other important goals. First, by aggregating the review data offline, query processing can be much more efficient. Accessing the raw data at query time would be prohibitively expensive. In our experiments, query processing was accelerated by a factor of up to 6.6x by only accessing the marker summaries. Second, the markers enable the system designer to shape which important aspects of the reviews to highlight in the application. Even though OPINEDB allows end users to specify arbitrary keywords in their query, markers can be useful for end users and applications in gauging the range of linguistic terms that are supported by marker summaries.

We now describe each of the above components in detail.

**Linguistic domains:** A *linguistic domain* is defined to be a set of short linguistic phrases (which we refer to as *linguistic variations*) that describe a particular aspect of an object. The linguistic domain allows OPINEDB to capture the different ways of describing the cleanliness of a room. For example, { “very clean”, “pretty clean”, “spotless”, “average”, “not dirty”, “dirty”, “stained carpet”, “very dirty”... } is a linguistic domain for the cleanliness aspect of the room object.

The linguistic domain is not enumerated in advance. OPINEDB bootstraps it by extracting phrases from the review data. Phrases in reviews can express opinions about an object directly (e.g., “very clean room”), or indirectly (e.g., “the room has stained carpet”). OPINEDB’s extraction module supports both types of opinions.

**Markers and Marker summaries:** A *marker summary* is defined over a linguistic domain and is a record type  $\text{Rcd} : [f_1, \dots, f_k]$  where  $\text{Rcd}$  is the name of the record,  $f_i$  are names of *markers*, and the type for each field is assumed to be a number. There are two types

<b>Hotels</b>	hotelname	capacity	address	price_pn
<b>HRoomCleanliness</b>	hotelname	* <i>room_cleanliness</i>		
<b>HBathroom</b>	hotelname	* <i>style</i>		
<b>HService</b>	hotelname	* <i>service</i>		
<b>HBed</b>	hotelname	* <i>comfort</i>		

#### Marker Summaries

\* room\_cleanliness: [very\_clean, average, dirty, very\_dirty]  
 \* style: [old, standard, modern, luxurious]  
 \* service: [exceptional, good, average, bad, very\_bad]  
 \* comfort: [very\_soft, soft, firm, very\_firm, ok, worn\_out]

**Figure 2: A subjective database schema for the hotel domain. Subjective attributes are prefixed with \*.**

of marker summaries: *linearly-ordered* or *categorical*. A linearly-ordered marker summary is one where  $f_1, \dots, f_k$  form a linear scale. An example of a such a marker summary for room cleanliness over the linguistic domain described earlier is shown below.

room\_cleanliness : [very\_clean, average, dirty, very\_dirty]

A phrase can contribute *in part* to multiple markers of a linearly-ordered marker summary. For example, the phrase “rooms are quite clean” can contribute in equal proportions (0.5 each) to the markers “very\_clean” and “average”. An example instance of the room\_cleanliness marker summary is [very\_clean: 90.5, average: 60.5, dirty: 30, very\_dirty: 20]. We use the term “marker summary” to refer to either the record type or the record instance when the context is clear.

Note that phrases in a linguistic domain do not always fall into a simple linearly-ordered scale of sentiment. For example, the fields we obtained by mining reviews from *booking.com* show that the quietness of a room may be described with words such as “annoying”, “peaceful”, “very noisy”, “traffic noise”, “constant noise”, which do not follow a natural linear order. To handle such cases, we also allow for categorical marker summaries in OPINEDB. A categorical marker summary is one where no two markers form a linear scale. An example of a categorical marker summary is

style : [old, standard, modern, luxurious]

A phrase can contribute *as a whole* to multiple markers in a categorical marker summary. For example, “extravagant old-fashioned bathrooms” contributes to both “old” and “luxurious” (1 count each).

At present, we assume that a marker summary is either linearly-ordered or categorical. Of course, for each categorical marker such as “luxurious”, there may be a linearly-ordered marker summary on the degree of “luxuriousness”. As with any database design, it is the task of the schema designer to decide the appropriate level of granularity to model the domain. In our context, OPINEDB assists her by clustering the linguistic domain (See Section 4.2.1 for details).

**Schema of a subjective database:** The schema of an OPINEDB application consists of three elements: (1) the main schema that is visible to the user and the application programmer, (2) the raw review data, and (3) the extractions of relevant phrases from the reviews from which we compute marker summaries. Parts (2) and (3) of the schema are intended to support queries that might qualify the reviews (e.g., consider only reviewers who have reviewed more than 10 hotels) and to support OPINEDB’s ability to fall back on raw text when a query cannot be answered using the database schema. We discuss some of the details of these components of the schema in Section 4. In what follows we focus on (1), which illustrates the main novel aspects of our data model.

The schema that is visible to the user or the application is a finite sequence of relation schemas each of the following form:  $R(K, A_1,$

$\dots, A_n$ ) where  $K$  is the key for  $R$  (for simplicity, we assume it's a single attribute), and  $A_1, \dots, A_m$  are a set of attributes.

We distinguish between two types of attributes. An *objective attribute* is an attribute whose value is based on facts and is largely indisputable. In contrast, there is no ground truth for the value of a *subjective attribute*. The value of a subjective attribute is “influenced by or based on personal beliefs or feelings, rather than based on facts”.<sup>1</sup> Figure 2 shows an example of a subjective database schema for the hotel domain.

The type of a subjective attribute is a marker summary over a linguistic domain. The linguistic domain of the subjective attribute `style`, for example, is a set of phrases that may include { “modern faucets”, “old shower”, “OK”, “adequate”, “luxurious bath towels”, ... }. As described earlier, the intuition is that the marker summary keeps a summary (e.g., histogram) of the subjective phrases for that attribute w.r.t. the markers.

The key of the `Hotels` relation is `hotelname` and it has three objective attributes, `capacity`, `address`, and `price_pn`. There are four additional relations with the same key attribute that contain subjective attributes. The attributes `room_cleanliness`, `style`, `service`, `comfort` are subjective attributes of the relations `HRoomCleanliness`, `HBathroom`, `HService`, and `HBed` respectively and their marker summaries are shown at the bottom of Figure 2.

A core issue that OPINEDB needs to address is that of aggregating a large collection of linguistic phrases onto marker summaries, which we will describe in Section 4. In what follows, we describe the query language of OPINEDB.

**Query:** The OPINEDB query language is essentially SQL with the extra ability of specifying atomic conditions in natural language in the `where` clause. For the purposes of this paper, we assume that an OPINEDB query consists of a single `select-from-where` clause. For example, assume that our hotel database is for hotels in London, then the query for hotels in London that cost less than 150 Euros per night, has clean rooms, and is good as a romantic getaway can be expressed as shown below:

```
select * from Hotels
where price_pn < 150 and
      “has really clean rooms” and “is a romantic getaway”
```

The `where` clause is a logical expression over a set of conditions. The example above shows a conjunction of a condition on price and two *query predicates* (or *predicates* in short), which are conditions specified in natural language and they need not be phrases that occur in the reviews. The user can also directly query for clean rooms using the attribute `room_cleanliness` in the `HRoomCleanliness` relation. However, to do so she will need to understand the exact semantics of the schema and specify the precise predicate, such as whether “clean rooms” should be interpreted as “very clean”, “average”, “dirty” or “very dirty” rooms. By extending the query language to accept natural language predicates, we can support a broader range of user interfaces to subjective databases. We will describe how OPINEDB automatically compiles the query predicates against the underlying schema.

Of course, natural language queries can involve non-atomic conditions, but OPINEDB relies on techniques such as [57] to decompose a complex query into atomic conditions. As such, we assume atomic query predicates throughout this paper.

As we shall describe in Section 3, our subjective query interpreter compiles the query predicate “has really clean rooms” into a predicate over the `room_cleanliness` attribute and the query predi-

<sup>1</sup><https://dictionary.cambridge.org/us/dictionary/english/subjective>

cate “is a romantic getaway” into a predicate over the `service` and `bathroom` attributes. If the query predicate cannot be satisfactorily interpreted into a predicate over the existing schema, OPINEDB falls back to the source reviews to arrive at a ranked set of answers.

**Benefits of a query language:** One of the major advantages of OPINEDB is that subjective data can be queried declaratively, and therefore we can express complex queries. For example, the semantics of the following query is well defined: “find hotels with clean room based on reviews after 2010”. Another important example is queries involving joins, such as “find a hotel with a lively bar on the same street as a cafe with a relaxing atmosphere.” (Figure 3, we leave the discussion of the join semantics to future work). Being implemented on an RDBMS, OPINEDB is also able to leverage any query optimization capability of the underlying engine to boost the querying performance.

```
select H.name, C.name from
(select * from Hotels where “has a lively bar”) as H
JOIN ON H.street == C.street
(select * from Cafes where “relaxing atmosphere”) as C
```

Figure 3: An example of an OPINEDB query with joins.

Next, we describe how OPINEDB processes queries. Although we continue to exemplify our technical discussions with examples from the hotel domain, the techniques we develop are not dependent on a particular domain. In fact, we have conducted our experiments in Section 5 on two domains: hotels and restaurants.

### 3. PROCESSING SUBJECTIVE QUERIES

To highlight the technical challenges OPINEDB faces in query processing, we begin with a simple class of queries and then move on to more complex ones. Figure 4 illustrates the entire process.

#### 3.1 Predicates with markers

We begin with queries that contain predicates where each predicate maps to a specific subjective attribute and one of its markers. In this discussion we ignore objective attributes because they do not introduce new challenges.

Consider a subjective query that contains a conjunction of the following query predicates:

```
select * from Hotels h
where “has firm beds” and “has luxurious bathrooms”
```

Here, we will assume that the query predicates can be interpreted into the following subjective attributes and their respective markers: `comfort`.“firm” and `style`.“luxurious”. In general, however, a query predicate might not match exactly with one specific marker of a subjective attribute. Mapping the predicate approximately and into multiple subjective attributes are necessary in many cases. Section 3.2 describes in more detail how OPINEDB interprets arbitrary query predicates.

Once we have the interpretation of each query predicate, we need to compute the *degree of truth* for each interpreted predicate for each entity in the database. We describe that process in detail in Section 3.3. In this simple case, since the query references markers of the subjective attribute, we assume that the degrees of truth have been computed in advance. Hence, we can focus on the last part of query processing which is to *combine* the degrees of truth of multiple predicates in a principled fashion.

**Combining degrees of truth** The degrees of truths are combined using fuzzy logic [28, 13]. Fuzzy logic generalizes propositional logic by allowing truth values to be real numbers in the range  $[0, 1]$ . The real truth value of a logical formula  $\varphi$  represents the degree of

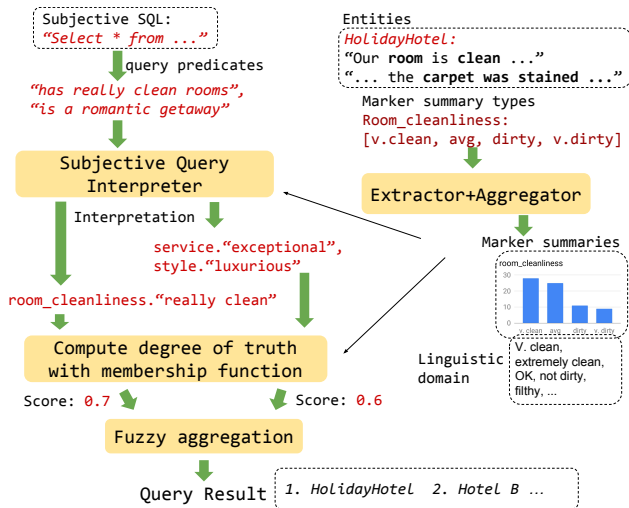


Figure 4: Processing subjective queries with OPINEDB.

$\varphi$  being satisfied, where a higher value means a higher degree of satisfying  $\varphi$ .

With our previous example, we now have the following query. The logical AND is replaced with  $\otimes$  and the query predicates have been replaced with their respective interpretations.

```
select * from Hotels h
where h.comfort  $\doteq$  "firm"  $\otimes$  h.style  $\doteq$  "luxurious"
```

In the query above,  $h.comfort$  denotes the marker summary of the bed comfort of hotel  $h$ . The condition  $h.comfort \doteq$  "firm" computes the degree of truth of how well the summary of bed comfort of hotel  $h$  represents the word "firm". Similarly, a degree of truth is computed for  $h.style \doteq$  "luxurious".

The fuzzy logical operator AND is denoted with  $\otimes$ . Later we will see queries with OR, which will be denoted by  $\oplus$ . In the most classic variant of fuzzy logic [13],  $x \otimes y$  is interpreted as  $\text{MIN}(x, y)$ ,  $\text{NOT}(x)$  is interpreted as  $1 - x$  and  $x \oplus y$  is interpreted as  $\text{MAX}(x, y)$ . Other variants include the multiplication variant [27] which we use in OPINEDB. In this variant, following De Morgan's law,  $x \otimes y$  is simply  $xy$ , negation is  $1 - x$  and  $x \oplus y$  is  $1 - (1 - x) * (1 - y)$ . Note that an objective predicate will simply be interpreted as 0 or 1.

Going back to our example, the conditions in the where clause are combined using the multiplication variant, which takes the product of the two degrees of truth. The result is a ranked list of hotels based on the final degree of truth that the entire expression in the where clause evaluates to.

**Why fuzzy logic?** An alternative to fuzzy logic would be to translate a subjective SQL query into a classic SQL query with hard selection constraints. For example, the previous two conditions can be written as:

$(h.comfort \doteq \text{"firm"}) > 0.8$  and  $(h.style \doteq \text{"luxurious"}) > 0.6$

where the thresholds 0.8 and 0.6 are specified by the application or the end-user. Aside from the inherent difficulty of specifying such thresholds in a meaningful fashion, this method may miss entities that may fall slightly out of the specified constraints. For example, a hotel with  $comfort$  score slightly less than 0.8 will be discarded from the result set for the above query. In contrast, the fuzzy interpretation is more forgiving and may therefore yield entities with good overall relevance to the query even if it may not satisfy the threshold of 0.8 on comfort. (See Appendix A of [30] for a visual illustration of this point). Furthermore, as the number of conditions

increases, the number of relevant entities that are potentially missed by the hard constraints only increases.

### 3.2 Predicates with arbitrary phrases

In the previous section, the query predicates were simple in the sense that it was clear which subjective attribute they refer to and which value (the marker) the user is specifying. The designer of an OPINEDB application may constrain the user to such queries, but one of the important benefits of subjectivity is that users can specify queries using their own terms. This, in turn, raises two challenges:

- Interpreting the phrase specified by the user: The user may specify a phrase for a subjective attribute that is not a marker. For example, she may ask for a hotel with rooms that are "really clean" or "meticulously clean". In some cases, these phrases may be in the linguistic domain of the subjective attribute and in other cases it may be a phrases that the application has never seen before.
- Determining the subjective attribute(s): in the simplest case, the challenge is to map the predicate to a single subjective attribute (e.g., mapping the predicate "has really clean rooms" to the subjective attribute  $room\_cleanliness$  with the phrase "really clean"). However, the user may specify predicates that do not correspond directly to a subjective attribute, such as "is a romantic getaway". In this case, a combination of subjective attributes may be equivalent to the predicate, or may at least provide strong evidence for it. For example, OPINEDB may know that hotels with  $service$ . "exceptional" and  $bathroom$ . "luxurious" are usually considered romantic. In other cases, the user may specify a predicate that does not correspond to any subjective attribute such as "has great towel art".

The subjective query interpreter is the component of OPINEDB that translates query predicates onto subjective attributes and their markers or to combinations thereof. The interpreter computes an *interpretation* for every predicate in the query. The interpretation consists of expressions of the form  $A.m$  where  $A$  is a subjective attribute and  $m$  is a marker of  $A$ . The goal of the interpreter is to find the expression over the  $A.m$ 's that best matches  $q$ . Each  $A.m$  replaces the original query predicate as a condition  $A \doteq m$  (e.g.,  $comfort \doteq$  "firm"). If a predicate interprets into multiple  $A.m$ 's, OPINEDB replaces the original query predicate as a disjunction of the results. For example, there are two subjective query predicates in our running example: "has really clean rooms" and "is a romantic getaway". The predicate "has really clean rooms" is interpreted into  $room\_cleanliness$ . "very clean" due to the high similarity between the predicate with the marker "very clean". However, the second predicate does not bear sufficiently high similarity to any of the existing markers. In this case, OPINEDB uses an alternative approach to map the predicate to markers by finding all markers that frequently co-occurs with the predicate. With this approach, the second predicate is interpreted into a disjunction of  $service$ . "exceptional" and  $bathroom$ . "luxurious" because the phrase "romantic getaway" frequently co-occurs with exceptional service or luxurious bathroom in the review corpus. Note that "exceptional service" or "luxurious bathrooms" may not reflect the true meaning of "romantic getaway". However, they are proxies of "romantic getaway" derived in an entirely data-driven way based on the reviews. We obtain the following fuzzy SQL snippet after this step.

```
select * from Hotels h
where price_pn < 150  $\otimes$ 
h.room_cleanliness  $\doteq$  "really clean"  $\otimes$ 
(h.service  $\doteq$  "exceptional"  $\oplus$  h.style  $\doteq$  "luxurious")
```




Figure 5: Fallback mechanisms in OPINEDB.

As we mention later, the co-occurrence method sometimes outputs a conjunction of  $A.m$ 's instead. For example, if “exceptional service” and “luxurious bathrooms” are frequently mentioned together along with romantic getaway instead of individually, then the above  $\oplus$  will be  $\otimes$  instead.

As noted earlier, it may not be possible to completely interpret a query predicate in terms of the database schema. Hence, in parallel with trying to interpret the query, OPINEDB also relies on a text retrieval system (described later in this section) to produce matching scores between database entities and query predicates. In principle, OPINEDB should combine the scores of the interpreter and the text retrieval system to produce the final ranking. In our current implementation, OPINEDB uses a threshold to determine whether it has enough confidence in the interpretation, and only if it does not, OPINEDB falls back on the text-retrieval method.

**Predicate interpretation algorithm** This algorithm takes as input a query predicate and returns as output an expression over the set of  $A.m$ 's as introduced above. OPINEDB currently uses a three-stage approach to interpret query predicates in a best-effort manner (see Figure 5): it first applies the word2vec method to find a direct interpretation of the query predicate. If this method fails to produce a satisfactory interpretation, it uses the co-occurrence method to find an approximate interpretation of the query predicate. If the second method fails to produce a satisfactory interpretation, it falls back to the text retrieval method to produce a ranked list of answers. When an interpretation is successfully obtained from the first or second method, the SQL query is rewritten based on the interpretation and executed to obtain a ranked list of result.

**Word2Vec method:** Given a query predicate, this method finds the linguistic variations of all subjective attributes having the highest similarity with the query predicate and returns the attributes and markers that correspond to the most similar variations as the interpretation. This method is based on the observation that most query predicates are simple so they likely to match closely with some linguistic variations already captured in the subjective database. Such common queries include “clean room”, “good breakfast”, and “nice location” for the hotel domain and “tasty food”, “friendly staff”, and “ambience” for the restaurant domain. These phrases and their synonyms also frequently appeared in reviews.

Word2vec [33] allows one to compute a vector representation of a word or a short phrase (e.g. bi-gram). The vector representation is typically a dense vector with hundreds of dimensions. Two phrases have similar vector representations if they share similar contexts in the text corpus, and so the two phrases are *semantically* similar to each other. The query predicates and the linguistic variations can contain multiple words or short phrases. To compute their vector representations  $\text{rep}(\cdot)$ , we use the IDF-weighted sum method commonly used in the NLP community:

$$\text{rep}(p) = \sum_{w \in p} w2v(w) \cdot \text{idf}(w). \quad (1)$$

Here,  $p$  is the query predicate or the linguistic variation,  $w$  is a word or short phrase of  $p$ ,  $w2v(w)$  is the word vector of  $w$ , and  $\text{idf}(w)$  is the Inverse Document Frequency (IDF) [9] of  $w$  in the review corpus. Intuitively,  $\text{idf}(w)$  measures the importance of the word  $w$  so that less frequent words are weighted higher in  $\text{rep}(p)$ . For example, the short phrase “very-clean” has a higher weight than “clean”

since “very-clean” is less frequent than “clean”. Then, to measure the closeness of a query predicate  $q$  to a linguistic variation  $p$ , we simply compute the cosine similarity of their representations:

$$\text{similarity}(q, p) = \cos(\text{rep}(q), \text{rep}(p)). \quad (2)$$

There are more sophisticated methods for computing short text representation (i.e., sentence embedding) like Skip-Thought Vectors [26] and InferSent [10]. These methods have shown good performance in tasks like sentence classification, entailment, and similarity search. One can build an interpreter method that first converts the query predicate into the sentence embedding then performs a similarity search over the linguistic domains. However, these methods usually involve computation with a neural network and similarity search which can be expensive. Such operations can lead to less efficient query processing. On the other hand, due to its simplicity, the IDF weighted method enables OPINEDB to reduce the cost of these expensive operators with efficient indexing schemes. We introduce one such method in Appendix B of [30].

The word2vec method can fail if there is no similar linguistic variation found in the database. Specifically, when the highest similarity returned is below a certain threshold (e.g., 0.5), OPINEDB will turn to the co-occurrence method, which we describe next.

**Co-occurrence method:** We can decide whether or not a predicate  $q$  should map to an expression  $A.m$  according to whether  $q$  frequently co-occurs with linguistic variations of  $m$  in the source text of the subjective database. We use this method only when the query predicate cannot be satisfactorily mapped to markers of the existing set of subjective attributes as described above. For example, the predicate “*is a romantic getaway*” is not sufficiently similar to any linguistic variation (i.e., the highest similarity is below the threshold 0.5). Hence, OPINEDB uses instead the co-occurrence method for this predicate. It discovers that this predicate frequently occurs in positive reviews where “*excellent service*” and “*five-star bathrooms*” are also mentioned. Hence, it is likely that the query predicate is correlated to the `service` attribute and “*exceptional service*” is the closest marker of that attribute. In addition, the query predicate is also close to the `style` attribute and “*luxurious bathrooms*” is the closest marker of that attribute to “*five-star bathrooms*”.

Specifically, given a query predicate  $q$ , OPINEDB first searches the source reviews of the subjective database to find all positive reviews where  $q$  occurs. We measure the positiveness of a review by applying sentiment analysis [5]. Among the set of related reviews, we find the top- $k$  reviews ranked by the following scoring function

$$\text{rank\_score}(d) = \text{BM25}(d, q) \cdot \text{senti}(d), \quad (3)$$

where  $d$  is a review text,  $\text{BM25}(d, q)$  is the classic Okapi BM25 [9] ranking function measuring relevance of  $d$  and  $q$  based on tf-idf and  $\text{senti}(d)$  is the sentiment score computed on the review  $d$ . Efficient computation of the top- $k$  documents ranked by BM25 is well-studied with mature implementations such as Elasticsearch [19].

Afterwards, we collect the set of linguistic variations extracted from the top- $k$  reviews. The most correlated attributes are the ones with the highest tf-idf score. Formally, for each subjective attribute  $A$ , we let  $\text{freq}_k(A)$  be the number of times that linguistic variations of attribute  $A$  are extracted among the top- $k$  search result. Let  $\text{idf}(A)$  be the inverse document frequency of attribute  $A$ . The final interpretation of a query predicate consists of a disjunction of  $n$  expressions  $A.m$  where (1)  $A$  is an attribute with the top- $n$  highest  $\text{freq}_k(A) \cdot \text{idf}(A)$  and (2)  $m$  is the marker of  $A$  with the highest frequency in the top- $k$  reviews. When the top- $n$  highest score is below a certain threshold, OPINEDB considers the result to be less confident and will turn to the results of text retrieval.

Table 2 illustrates the strength of the co-occurrence method with outputs from real examples.



**Table 2: Example output of the co-occurrence method.**

	Query Predicates	Top-1 Interpretations
Hotels	for our anniversary multiple eating options kid friendly hotel	staff. “ <i>helpful concierge</i> ” food. “ <i>good options</i> ” staff. “ <i>very kind staff</i> ”
Restaurants	dinner with kids close to public transportation private dinner	table. “ <i>high chair</i> ” general. “ <i>great place</i> ” vibe. “ <i>quiet place</i> ”

After interpretation, OPINEDB computes how well each resulting  $A.m$  matches with a database entity by computing the degree of truth (Section 3.3).

**Text-retrieval method:** In the event that both word2vec and co-occurrence failed to interpret a query predicate with high confidence, we fall back to traditional information retrieval techniques to compute the degree of truth based on ranking scores of each entity w.r.t. the query phrase.

Following a previous work [16], the text-retrieval method represents each entity by a single document  $D$  obtained by combining all source reviews of the entity. Then for a subjective query predicate  $q$ , OPINEDB computes the ranking score simply as  $\text{BM25}(D, q)$ . To convert the value into a degree of truth, we set a constant threshold  $c$  and apply the sigmoid function. The returned degree of truth is computed as  $\text{sigmoid}(\text{BM25}(D, q) - c)$ .

### 3.3 Computing the degrees of truth

After the predicates have been interpreted into an expression over a set of  $A.m$ ’s, OPINEDB now needs to compute how well the reviews of each entity represent each query predicate  $q$ . In other words, OPINEDB computes a degree of truth, which is a value between 0 (false) and 1 (true) for each interpreted predicate such as `room_cleanliness`. “very clean”.

As mentioned in Section 3.1, the degrees of truth for variations in the linguistic domain (i.e.,  $m = q$ ) of each subjective attribute can be pre-computed so that they can simply be looked up at query time. For phrases that are outside the linguistic domain ( $m \neq q$ ), the degrees of truth are computed during query time. These degrees of truth, once computed, can also be indexed and so they can be simply retrieved in future.

OPINEDB has access to the relevant marker summaries through the interpretations obtained. Next, we describe how OPINEDB translates marker summaries into degrees of truth w.r.t. a predicate.

**Membership functions:** OPINEDB constructs a *membership function* [28] to compute the degree of truth of an interpretation  $A.m$  based on the marker summaries of  $A$  and the interpreted marker  $m$  with original query predicate  $q$ . In effect, the membership function further aggregates the marker summary to compute the degree of truth. For example, the marker summary [“v. clean”: 20, “avg.”: 10, “dirty”: 1, “v. dirty”: 0] should have a value close to 1 (e.g., 0.95) for the query predicate “really clean room” since most reviews mentioned that the rooms are clean. In contrast, the marker summary [“avg.”: 10, “dirty”: 10] should have a much lower value (e.g., 0.2) for “really clean room” since half of the extraction results stated that the rooms are dirty.

OPINEDB uses machine learning to construct the membership functions. Specifically, OPINEDB trains classification models from labeled tuples  $\{(S_i, p_i, y_i)\}_{i \geq 0}$  where each  $S_i$  is a marker summary,  $p_i$  is a phrase and  $y_i \in \{0, 1\}$  is a binary label that indicates whether or not  $S_i$  satisfies  $p_i$ . Binary classification is suitable for this task because binary labels are less expensive to obtain compared to numeric labels. Furthermore, many popular models such as *Logistic Regression* compute intermediate values that can be interpreted as a degree of truth in  $[0, 1]$ . More specifically, logistic

regression learns the binary classifier by first learning a *logistic loss function* which can compute the probability (so  $[0, 1]$  is its range) of a label being positive given the input tuple. As a result, we can directly use the probability output as the membership function by interpreting the probability as the degree of truth.

The model makes use of features constructed from precomputed information in each marker summary. By doing so, OPINEDB can speed up query processing by avoiding scanning the full extraction tables. Such features include the sizes of the markers, the average sentiment scores, and the centers of the phrase vectors of the phrases mapped to each marker. OPINEDB trains high-quality models using these features as the markers are expected to be good representations of the underlying linguistic domain.

In our experiment, we found that with a set of 1,000 labeled tuples, we obtained Logistic Regression classifier of 71% to 75% accuracy (Section 5.4.2) on the hotel and the restaurant domains. This means that the features constructed from the marker summaries are high-quality and hence, the logistic loss is suitable as the membership function. In addition, the use of the marker summaries in query processing results in a speedup up to 6.6x.

## 4. DESIGNING SUBJECTIVE DATABASES

The creation of the schema and of the data in a subjective database are closely intertwined. Next, we describe how OPINEDB (1) extracts opinions from text and (2) based on the extractions, it constructs subjective attributes and marker summaries. Both processes are interactive in that the schema designer of OPINEDB provides input on what the important attributes are and what information needs to be extracted.

The problem of extracting opinions from reviews is a well studied problem in the NLP literature (e.g., [56, 31, 51, 20, 50, 49, 45, 39, 23]). Our focus is not on developing new techniques for opinion mining, but rather devising techniques that enable the schema designer to quickly develop a good schema for the database.

### 4.1 Extracting opinions from reviews

OPINEDB extracts all the pairs of *aspect term* and associated *opinion term*. For example, given the sentence:

*The room was very clean, but the staff was not so friendly.*

OPINEDB would extract pairs of the form

$\{(\text{“room”}, \text{“very clean”}), (\text{“staff”}, \text{“not so friendly”})\}$ .

Within each pair, the first element is the *aspect term*, which represents the target of the opinion. The second element is the *opinion term* containing an opinion on that aspect. This task is closely related to the *Aspect-Based Sentiment Analysis* (ABSA) problem [34, 36, 35], which aims at finding the opinionated aspect terms from text and predicting their sentiment scores (i.e., positive or negative). The solution was later extended to also extract the opinion terms [51, 50]. Hence, these proposed techniques are suitable for OPINEDB’s extractor.

Following previous work, we design the extractor as a two-stage procedure: tagging and pairing. This is illustrated in Figure 6 with a real example of our extractor applied to a hotel review. During the tagging stage, the tokens of the input sentence are classified as (part of) an aspect term (AS), an opinion term (OP), or irrelevant (O). In the pairing stage, the tagged aspect/opinion terms are paired to form the extracted opinions. Here, we focus on optimizing the quality of the tagging stage since the pairing stage can be implemented with a rule-based model and achieves comparable good performance to that of a learned model (More details in Appendix C of [30]).

The reported results for electronics and restaurants reviews [51, 50] were promising. However, the lack of labeled training data




Figure 6: Tagging and pairing.

makes their trained deep learning models hard to generalize to other domains [35], (e.g., hotels). Hence, instead of using these techniques, we built our extractor based on BERT [11], the recently developed pre-trained NLP model that achieved start-of-the-art performance in major NLP tasks including sentiment analysis and tagging. The transfer learning capability of BERT allows the extractor model to first be trained on a large set of unlabeled text data, and then fine-tuned on a labeled training set of a much smaller size. In our experiments we show that in the hotel domain, OPINEDB’s extractor achieved a good 74.71% F1 score when we use a pre-trained BERT model from [11] with 912 labeled review sentences for fine-tuning. In contrast, the method based on [51, 50] achieves only 68.04% F1 score. Labeling and training using these 912 review sentences data were done within a few hours. Another advantage of using a pre-trained model like BERT is that the model is fixed so that it does not require the schema designer to have NLP/ML expertise to program the neural networks or to tune the hyper-parameters. As such, the whole process of developing the extractor is extremely efficient.

## 4.2 Designing the subjective attributes

In the next step, the schema designer provides a set of subjective attributes and OPINEDB maps each extracted pair to one of the attributes. The design of the set of subjective attributes is analogous to the design of a relational database schema where we rely on the schema designer to decide what should or not be part of the schema. In future, we plan to provide the schema designer suggestions of possible subjective attributes by analyzing the extracted pairs. In our experience, the number of subjective attributes is quite small (11 attributes for the restaurant domain and 15 attributes for the hotel domain).

We formulate the problem of assigning extracted pairs to attributes as a text classification problem. For example, OPINEDB would classify the above pairs to attributes as follows:

$(\text{"room", "very clean"}) \rightarrow \text{room\_cleanliness}$   
 $(\text{"staff", "not so friendly"}) \rightarrow \text{staff}.$

We need labeled data to train such a classifier. To reduce the labeling cost, we construct the training set automatically via seed expansion. For each attribute  $A$ , the designer provides a pair  $(E, P)$  of seeds where  $E$  is a set of aspect terms that  $A$  describes and  $P$  is a set of opinion terms that refer to those aspects. For example, for the attribute `room_cleanliness`, the designer can provide:

$E = \{\text{"room", "bedroom", "carpet", "furniture"}\}$   
 $P = \{\text{"clean", "dirty", "very clean", "very dirty", "stained", "dusty"}\}$

OPINEDB then expands the seeds with synonyms using a word2vec model. The word2vec model is trained on the review corpus so it can capture similar phrases more accurately. For example, phrases such as “room” may be expanded into “suite”, “executive suite”, or “apartment”. Next, for each pair  $(e, p)$  in the cross product of  $(E \times P)$  and attribute  $A$ , OPINEDB constructs a labeled tuple  $(\text{concat}(e, p), A)$  where  $\text{concat}(e, p)$  is the example and the attribute  $A$  is the label.

This approach allows OPINEDB to train a high-quality attribute classifier with only little effort in creating the labeled dataset. For example, with only 235 seeds of 15 restaurant attributes (expanded into a training set of 5,000 tuples), OPINEDB is able to obtain a classifier with 88% accuracy on the test set.

### 4.2.1 Defining markers

Given the classification result, we can define the linguistic domain of each attribute to be the set of all possible phrases (concatenations of the aspect and opinion terms) assigned to it. Next, the schema designer needs to specify a marker summary for each attribute and whether or not the attribute is linearly ordered. OPINEDB alleviates the effort required for this step by providing two automated methods. The design of these two methods is based on the observation that most linguistic domains can be modeled as one of the two types:

- **Linearly-ordered domains.** The phrases of linguistic domains for attributes such as `room_cleanliness` can be ordered linearly. For example, “dirty” < “average” < “clean” < “very clean”. In such cases, we can generate the markers by leveraging sentiment analysis [5]. More specifically, we sort the phrases by their sentiment scores  $\text{sent}(\cdot)$  and divide the linguistic domain equally into  $k$  buckets. The markers are designated as the linguistic variation in the center of each bucket.
- **Categorical domains.** Linguistic domains can also be categorical, which means that the linguistic variations can be categorized into a few topics. The `bathroom` attribute is categorical – the phrases can be {“luxurious”, “modern”, “old-styled”} where there is no clear linear order but can be summarized into categories. In such cases, OPINEDB performs  $k$ -means clustering on the linguistic domain. Afterwards, OPINEDB suggests a set of markers by selecting the linguistic variations that correspond to the centroid of each cluster.

### 4.2.2 Calculating marker summaries

Once the marker summary is defined, the next step is to aggregate the data from the reviews according to the markers. In general, the appropriate aggregation function depends on the semantics of the attribute. Attributes like `friendlyStaff` will change over time more frequently than the attribute `quietLocation`. Hence, in the former case we may want to weigh recent reviews more heavily. As another example, some aspects of a hotel (such as `towelArt`) are mentioned much less frequently than others. In these cases, even a few mentions should be considered a strong signal.

The aggregation function can also depend on the specific needs of the OPINEDB application. For example, an application might decide to assign uniform weights to all reviews but another application might want to assign higher weights to reviews marked as “helpful” by other users. The full exploration of different possible aggregation functions is beyond the scope of this paper but we believe is an important aspect of building subjective databases.

In the current implementation, OPINEDB aggregates phrases of reviews based on the number of occurrences in the reviews. For example, a `room_cleanliness` marker summary is constructed for each hotel by counting the number of phrases of reviews from the extraction relation that contain linguistic variations closest to “very clean”, “average”, “dirty”, “very dirty”, respectively for that hotel. Even though our model for linearly-ordered marker summaries allows for a phrase to contribute in part to different markers, our initial implementation matches a phrase to only the best matching marker. We plan to explore techniques to weigh the proportion of contributions of each phrase to different markers in the future. The resulting histogram is the marker summary for that hotel and is stored in the `room_cleanliness` attribute of relation `HRoomCleanliness`.

The marker summaries can be incrementally computed. Furthermore, any result returned can be supported with evidence from the reviews on why that result is returned because OPINEDB keeps track of provenance of extracted phrases.



## 5. EXPERIMENTS

We present our implementation of OPINEDB and our experimental results.

**Overview** Our first set of experiments investigates the need for experiential search. We show that a significant proportion of user requirements are experiential in several domains.

In our second set of experiments, we compare the quality of OPINEDB’s query results with two baselines. To do this, we constructed subjective databases for two domains: hotels and restaurants and designed a method for evaluating subjective query results. We show that even when OPINEDB is evaluated conservatively, OPINEDB outperforms the other two baselines over a variety of subjective queries.

We also present experimental results on the quality of critical OPINEDB components. We show that the extractor which we use to produce the linguistic domains of our subjective database schema achieve F1 scores of close to 75% for hotels and 85% for restaurants with only a small amount of labeled data provided. We also show that the marker summaries significantly accelerate subjective query processing (from 3.3x to 6.6x) while maintaining the quality of the query results. Finally, we show that the predicate interpretation algorithm achieves a precision of up to 85% on the combined use of word2vec and co-occurrence interpretation methods.

### 5.1 The need for experiential search

Our very first experiment was to verify whether or not users search experientially. To do this, we conducted a user study on Amazon Mechanical Turk [7] to determine what are the important criteria users consider when they search for certain types of entities. More specifically, each MTurk worker is provided with a question for a particular domain. For example, we posed the following task for hotels: “*Suppose you have planned a vacation and are looking for a hotel. Other than cost, list 7 separate criteria you’d likely value the most when deciding on a hotel*”.

We asked 30 workers such questions for each domain. We collected the answers and manually (and conservatively) evaluate whether each criterion is subjective or objective. For example, wifi is a criterion that frequently shows up in hotel search but we interpret it as objective (*is there wifi*) rather than subjective (*fast and reliable wifi*). We asked similar survey questions for other domains such as Restaurant, Education, Career, Real Estate, Car, and Vacation. Table 3 summarizes our results. The table shows that a significant number of the desired properties are subjective for several domains. However, to the best of our knowledge, online services for these domains today provide keyword search over the reviews at best and do not directly support subjective querying.

**Table 3: Subjective attributes in different domains.**

Domain	%Subj. Attr	Some examples
Hotel	69.0%	cleanliness, food, comfortable
Restaurant	64.3%	food, ambiance, variety, service
Vacation	82.6%	weather, safety, culture, nightlife
College	77.4%	dorm quality, faculty, diversity
Home	68.8%	space, good schools, quiet, safe
Career	65.8%	work-life balance, colleagues, culture
Car	56.0%	comfortable, safety, reliability

### 5.2 Experimental settings

Before we present the evaluation of OPINEDB, we describe how we measure the quality of our query results and how we generated subjective queries for our experiments.

#### 5.2.1 Implementation and experiment settings

We implemented the extraction pipeline of OPINEDB in Python and used standard packages including Tensorflow [1] for neural networks, NLTK [5] for sentiment analysis, and Gensim [40] for word2vec. The core part of the pipeline is the adaptation of an existing neural network [17] based on BERT [11]<sup>2</sup>, BiLSTM, and Conditional Random Field (CRF), adapted to opinion extractions.

We implemented the querying engine of OPINEDB on top of PostgreSQL [43]. We store the results of the extraction pipeline in a postgres instance. To execute a subjective SQL query, OPINEDB simply parses it with sqlparse and applies the query interpretation algorithms described in Section 3 to translate the input query into an executable SQL query. The resulting SQL query computes the subjective predicates (translated into membership functions) as user-defined aggregates in postgres. For simplifying our experiments, we also implemented a version of OPINEDB without using PostgreSQL. Both implementations and all the experimental scripts are open-source and available in [18].

In our experiments, we designed two subjective database schemas (for hotels and restaurants) and their respective linguistic domains and marker summaries with OPINEDB. The reviews for hotels and restaurants are from two real-world datasets: Booking.com dataset [46] with 515,739 reviews for 1,493 hotels and a subset of the Yelp [47] dataset with 176,302 reviews for 860 restaurants in Toronto. We trained neural networks using an AWS p2.xlarge server with one NVIDIA Tesla K80 GPU. All other experiments were conducted on a server machine with Intel(R) Xeon(R) CPU E5 2.10GHz CPUs.

#### 5.2.2 Generating subjective queries

Since there is no benchmark for subjective queries, we had to create one. Along with the survey result in Table 3, we also collected a set of subjective queries for the hotel and restaurant domains. We collected 190 subjective query predicates for hotels and 185 query predicates for restaurants. We construct conjunctions of query predicates by uniform sampling of the predicates. These conjunctions will form the where clauses of our subjective SQL query. We consider 3 sets of queries (easy, medium, and hard) for each domain. The number of conjuncts in easy, medium, and hard queries are 2, 4, and 7 respectively. Each set consists of 100 subjective queries.

We further increase the complexity of the queries by adding two variations to each query. Specifically, each query is extended with each of the following options which are conditions over objective attributes. For hotel queries, the two options are: (1) find all hotels in London less than \$300 per night and (2) find all hotels in Amsterdam. For restaurant queries, the two options are: (1) find all low-priced restaurants (i.e., price range is one ‘\$’ sign according to yelp) and (2) find all Japanese restaurants.

Table 4 shows some statistics of the hotels/restaurants under each selection predicate. For example, there are 189 London hotels that are less than \$300 per night. The restaurant reviews tend to be longer (higher average number of words) and more positive (as indicated by the average polarity returned by the NLTK sentiment analyzer).

**Table 4: Review statistics in booking.com and yelp.**

	#Entities	#Reviews	avg #words	avg polarity
London, <\$300	189	139,293	34.27	0.19
Amsterdam	91	45,875	37.02	0.21
Low Price	112	22,302	104.01	0.71
JP Cuisine	108	24,701	126.02	0.72

#### 5.2.3 Evaluation metrics

In the experiments, we use a metric based on the well-known Normalized Discounted Cumulative Gain (NDCG) [9] to measure how

<sup>2</sup>We used the 12-layer uncased pre-trained model in all the experiments.

**Table 5: Query result quality by the top-10 results in the booking.com dataset (left) and the yelp dataset (right) where quality (NDCG@10) = #Satisfied-predicates / max. #predicates can possibly be satisfied. Each entry has a confidence interval within  $\pm 0.0168$  (on 0.95 confidence level).**

Method	London $\wedge$ <300			Amsterdam		
	easy	medium	hard	easy	medium	hard
GZ12 (IR-based)	0.75	0.75	0.76	0.66	0.68	0.71
ByPrice	0.65	0.67	0.68	0.37	0.38	0.39
ByRating	0.62	0.64	0.65	0.51	0.53	0.54
1-Attribute	0.71	0.72	0.72	0.70	0.71	0.72
2-Attribute	0.76	0.77	0.78	0.71	0.72	0.73
OpineDB	0.80	0.82	0.84	0.81	0.83	0.84

Method	Low Price			JP Cuisine		
	easy	medium	hard	easy	medium	hard
GZ12 (IR-based)	0.68	0.71	0.73	0.72	0.75	0.78
ByPrice	0.56	0.60	0.62	0.64	0.67	0.69
ByRating	0.65	0.70	0.72	0.70	0.73	0.75
1-Attribute	0.74	0.77	0.78	0.76	0.78	0.80
2-Attribute	0.77	0.78	0.79	0.80	0.81	0.83
OpineDB	0.76	0.80	0.82	0.78	0.81	0.84

well the entities in the result satisfy the predicates in the subjective query. More precisely, assume a subjective query  $Q$  with  $n$  query predicates  $\{q_1, \dots, q_n\}$  returns top- $k$  entities  $E = \{e_1, \dots, e_k\}$ . Let  $\text{sat}(q_i, e_j) \in \{0, 1\}$  denote whether  $e_i$  satisfies  $q_j$ . The quality of the result is measured by counting the total number of query predicates that are satisfied by all the  $k$  entities in the result:

$$\text{sat}(Q, E) = \sum_{j=1}^k \left( \sum_{i=1}^n \text{sat}(q_i, e_j) \right) / \log_2(j + 1).$$

Intuitively, a higher  $\text{sat}(Q, E)$  indicates that the top- $k$  entities are more relevant to the searched query predicates in  $Q$ . The term  $1/\log_2(j+1)$  logarithmically penalizes the irrelevant entities closer to the top of the result.

**Ground truth:** The ground truth of  $\text{sat}(q, e)$  is expensive to obtain as it requires one to go through all the reviews. So, we adopt a lighter-weight approach to generate the labels of  $\text{sat}(q, e)$ . First, we manually identify the subjective attribute  $A$  in the schema closest to a query predicate  $q$  (e.g., the closest attribute to “with clean rooms” is `room_cleanliness`). Afterwards, we ask a human labeler to label  $\text{sat}(q, e)$  by inspecting the marker summary of attribute  $A$ . We also notice that many summaries have (1) only a few number of reviews, (2) a large fraction of unmatched phrases, or (3) a large fraction of negative phrases. In these cases, we can further reduce the labeling cost by avoiding human labelers altogether as labels can be automatically generated with high accuracy using a set of rules. We verified 20  $\text{sat}(q, e)$  labels for restaurants and 20 for hotels by inspecting their source reviews. Both sets have 19/20 labels well-supported by the underlying reviews.

To better illustrate the quality of the query results in our experiments, we also compute  $\text{sat-max}(Q)$ , the maximum score of the quality of a query result can be when we know the ground truth. We have  $\text{sat-max}(Q_i) = \max_E \{\text{sat}(Q_i, E)\}$ . For a set of queries  $\{Q_1, \dots, Q_N\}$  with query results  $\{E_1, \dots, E_N\}$ , the *quality of the workload* is then computed as

$$\text{quality}(\{Q_1, \dots, Q_N\}) = \frac{1}{N} \sum_{i=1}^N \frac{\text{sat}(Q_i, E_i)}{\text{sat-max}(Q_i)}.$$

### 5.3 Comparing OPINEDB with baselines

We compare OPINEDB with two baselines: (1) IR-based search engine (IR) and (2) attribute-based query engine (AB).

The IR baseline is an implementation of [16] (GZ12), which applied the IR method Okapi BM25 [9] retrieval model to rank entities based on the opinions they received. Following [16], we also added the capability to perform query expansion and different methods for combining multiple query predicates to make the baseline more competitive.

The AB baseline represents what a user can obtain through online services such as `booking.com` or `yelp.com` by freely trying combinations of queryable attributes to obtain the best results. Hence, it is a strong baseline for comparing OPINEDB. For example, to search

for a hotel, the user can rank the hotels by price or rating or even the predefined filters for some subjective attributes. We scraped all 8 subjective attributes (Location, Cleanliness, Staff, Comfort, Facilities, Value for Money, Breakfast, Free Wifi) from `booking.com`. We assume that the user can choose two of the above attributes and rank the hotels by their sums. Among all the combinations of attributes, we pick the one that maximizes the score  $\text{sat}(Q, E)$ .

Similarly, for restaurant queries, the user can rank the restaurants by the number of stars or by the total number of reviews received. Additionally, the user can choose to filter the restaurants using one or two of the 33 available categorical attributes in the dataset. Some examples are `Attire`, `GoodForGroups`, `NoiseLevel`, and `Ambience`. The combination with the maximal  $\text{sat}(Q, E)$  is picked.

Table 5 shows the quality results on both datasets. Each column in the tables represents the type of query used. That is, whether it is easy, medium, or hard and extended with an objective predicate (e.g., in London). The first row tabulates the results for the IR method, followed by 4 variations of the AB baseline and finally, OPINEDB. As noted earlier, the numbers represent the proportion of the number of query predicates that are satisfied out of the maximal number of query predicates that can possibly be satisfied.

To ensure the results’ statistical significance, we repeated the experiment on 10 different samples of query sets (i.e., in total 1,000 queries per setting) and computed the averages with confidence intervals. The results are indeed statistically significant as the maximal size of the confidence interval is no larger than  $\pm 0.0168$ .

In both datasets, OPINEDB outperforms the IR baseline by a sizable margin (by  $\sim 0.05$  to  $\sim 0.15$  for hotels queries and by  $\sim 0.06$  to  $\sim 0.10$  for restaurant queries). This is not surprising as the IR baseline retrieves hotels with reviews that contain keywords in the query predicates (e.g., “clean”) even if the same reviews contain the opposite negative words (e.g., “dirty”) or may have used the phrase “not clean”. On the other hand, OPINEDB’s membership functions can carefully discern between entities based on the frequencies of positive versus negative phrases. We show one such example in the Appendix D of the full version [30].

The AB baseline has similar performance with the IR baseline. The tables clearly show that the result quality increases when more subjective attributes are used. The AB baseline also performs much better in the restaurant queries. This is because the yelp datasets contain more queryable attributes than the hotel dataset. These findings reaffirm our belief that utilizing subjective attributes is important for experience search engines. Still, OPINEDB outperforms the AB baseline especially when there are more subjective query predicates. We believe this is due to OPINEDB’s ability to accurately map those query predicates to subjective attributes.

Observe that the result quality of OPINEDB is higher in the hotel domain than in the restaurant domain resulting in larger margins of improvement compared to baselines. This is because the hotel dataset contains many more reviews per hotel and thus the generated marker summaries are more representative and statistically signifi-

cant. This result matches our intuition and suggests that OPINEDB brings more value to applications as the number of reviews grows.

## 5.4 Quality of OpinedB components

Next, we evaluate the quality of important parts of OPINEDB : the extractor, the marker summaries, and the predicate interpreter.

### 5.4.1 Extractor and subjective DB construction

We start by showing that OPINEDB’s extraction module achieves the state-of-the-art or better quality. Moreover, we show that with the recent advances in NLP, we are able to achieve the good performance with only a small amount of training data.

We evaluate the extractor on 4 datasets summarized in Table 6. The first 3 datasets are from ABSA competitions: SemEval 2014 Task 4 (Laptops and Restaurants) [36] and SemEval 2015 Task 12 (Restaurants) [35]. Each dataset contains a set of sentences labeled with *aspect terms* and *opinion terms* corresponding to the opinion targets and detailed opinions mentioned in Section 4.1 respectively. The aspect term labels are from the original datasets and the opinion term labels were added by [51] and [50]. Since there is no existing labeled opinion extraction datasets for hotels, we created our own (Booking.com Hotel) to train our extractor. The sizes of the datasets are listed in Table 6. Note that none of the datasets are big. The extractors trained on the hotel dataset and the SemEval-14 restaurant dataset are the ones used in the experiment reported in Table 5.

Similar to other extraction tasks like Named Entity Recognition (NER) [42], the extraction quality is measured by the *F1 scores* of the aspect terms and the opinion terms. An aspect/opinion term is considered correctly extracted only when the extracted term matches exactly with the ground truth term. As shown in Table 6, the model of OPINEDB’s extractor (BERT+BiLSTM+CRF) outperforms the previous state-of-the-art models in all the 4 datasets<sup>3</sup>. The improvement ranged from  $\sim 0.01\%$  to  $\sim 6.67\%$ . We noticed that the improvement is the highest for the hotel dataset which has the least number of training sentences. We believe that this is because of the transfer learning ability of the BERT model as similar observations were also reported in [11] for cases with a small amount of training data. We also found that the quality of extraction is robust to small training sets so that a high-quality extractor can be obtained at very low cost. Specifically, for the hotel domain, we found in our experiment that even with a training set of 20% of the original size (<200 sentences), the F1 score remains close to 70% which is still higher than the SOTA.

**Table 6: Datasets for training and evaluating the extractor of OPINEDB. The last two columns list the combined F1 scores (averaging the aspect/opinion terms’ F1 scores) of the previous state-of-the-art (SOTA) models [50, 51] and the one OPINEDB used. Each score of our model is the average score of 10 models trained separately. The  $\pm$  indicates the standard confidence interval computed on a 0.95 confidence level.**

Datasets	Train	Test	Total	SOTA	Our Model
SemEval-14 Restaurant	3,041	800	3,841	85.52	85.53 $\pm$ 0.40
SemEval-14 Laptop	3,045	800	3,845	78.99	79.82 $\pm$ 0.35
SemEval-15 Restaurant	1,315	685	2,000	72.21	75.40 $\pm$ 0.58
Booking.com Hotel	800	112	912	68.04	74.71 $\pm$ 0.72

OPINEDB also performs classification to map each pair of extracted aspect/opinion terms into the set of subjective attributes. To train such classifiers for the hotel and the restaurant domain, OPINEDB applies weak supervision with the seed expansion techniques described in Section 4.1. The schema designer provided 15

<sup>3</sup>We collected the F1 scores of the SemEval datasets from [50, 51] and re-trained their model on the hotel dataset (10 times to get the average).

subjective attributes with 277 seed phrases for the hotel domain and 11 attributes with 235 phrases for the restaurant domain. For each domain, the seed expansion generates a training set of 5,000 records and we manually labeled 1,000 addition records for testing purpose. Both classifiers performed well: the hotel attribute classifier achieves an accuracy of 86.63% and the accuracy of the restaurant attribute classifier is 88.29%.

Overall, the process of creating the subjective DB is efficient. As mentioned above, the effort of creating the extractor for Hotels was < 5 hours of human labeling and OPINEDB’s extractor performs better than SOTA techniques. Writing each seed set took us no more than 2 hours with 1 developer. These costs are small compared to the entire process of developing a travel application.

### 5.4.2 Marker summaries and membership functions

In addition to being a key component of OPINEDB’s data model, the marker summaries benefit a subjective DB in two ways: (1) accelerating query processing, and (2) creating high-quality features for entity ranking. In this section we experimentally evaluate these benefits. For both the hotel and the restaurant domain, we created 10 markers for each subjective attribute by applying the automatic approach described in Section 4.2.1. For each set of queries (the London, Amsterdam, Low-Price, and JP Cuisine queries listed above), we compared OPINEDB with a small variant of it which does not leverage the markers. Specifically, when the markers are used, the logistic regression (LR) model for membership scoring uses features precomputed for each marker (see Section 3 for details). Without the markers, the model uses another set of engineered features similar to the set when the markers are used with the addition of new features (e.g., the number/fraction of phrases that are similar to the query predicate) directly computed from the extracted phrases. Each LR model is trained on 1,000 labeled pairs of entity and query.

**Table 7: OPINEDB w. 10 markers (10-mkrs) vs. no marker (no-mkrs). The running time is per 100 queries. For each query set, LR-accuracy is the test accuracy of the Logistic Regression model, NDCG@10 measures the query result quality, and Runtime is the total running time of 100 queries in seconds. Each value is the average of 10 repeated runs. The max.CI column contains the maximal confidence interval of each row (on a 0.95 confidence level).**

	London	Amsterdam	Low-Price	JP Cuisine	max.CI	
10-mkrs	LR-accuracy	0.71	0.75	0.73	0.73	$\pm 0.016$
	NDCG@10	0.82	0.83	0.79	0.81	$\pm 0.012$
	Runtime	18.84s	9.89s	12.55s	13.95s	$\pm 0.726$
No-mkrs	LR-accuracy	0.71	0.76	0.71	0.71	$\pm 0.02$
	NDCG@10	0.76	0.83	0.81	0.83	$\pm 0.016$
	Runtime	68.66s	33.00s	70.05s	92.68s	$\pm 4.689$
Speedup	3.65x	3.34x	5.59x	6.65x	$\pm 0.237$	

Table 7 summarizes the results. There is a significant performance improvement when markers are used, ranging from 3.34x (Amsterdam) to 6.65x (JP Cuisine). The overall average time per query is 0.14 sec when markers are used and 0.93 sec without the use of markers. Note that this gap will be much larger in real-world review datasets and queries over a larger number of entities. We also observed that the quality of the membership functions (LR-accuracy) and query results (NDCG@10) remain mostly unchanged even with the speedup in performance. This is because on small training sets (1,000 in our case), a smaller number of good features can help improve accuracy without overfitting. By aggregating the extracted phrases onto the marker summaries, OPINEDB reduces the number of features while keeping the most relevant information.

### 5.4.3 Query predicate interpretation

We executed our predicate interpretation algorithms on the hotel and the restaurant sets of query predicates from Section 5.2.2. For each subjective query predicate, we manually labeled it with the closest subjective attribute that the predicate should be mapped. An interpretation result is counted as correct if the attribute matches exactly with the ground truth.

**Table 8: Accuracy of different methods for query interpretation with confidence intervals over 10 runs.**

Query sets	size	w2v	co-occur	w2v+co-occur	max.CI
Hotel queries	190	84.05	72.63	84.89	0.60
Restaurant queries	185	81.62	68.65	82.16	0.52

Table 8 shows the accuracy of the two methods (word2vec and co-occurrence) when used independently in the predicate interpretation algorithm and when used in combination (with the fallback similarity threshold set to 0.8). The word2vec method produces reasonably high-quality (>80% accuracy) interpretations. The co-occurrence method has relatively lower accuracy (68% to 72%), but it still improves the accuracy of the base w2v method when combined (by 0.84% for hotel queries and 0.54% for restaurants). This is because although the co-occurrence method is relatively less accurate, it captures nicely the hard cases (long and uncommon text) that the word2vec method fails to capture.

## 6. RELATED WORK

The fields of sentiment analysis and opinion mining [56, 31, 51, 20, 50, 49, 45, 39, 52, 23] have developed techniques for extracting subjective data from text. While sentiment analysis tries to decide whether a particular text is positive or negative about an object or an aspect of an object, opinion mining aims to summarize a large collection of sentiments in a way that is informative to the user. In contrast, OPINEDB incorporates subjective opinion data into a general data management system, and addresses the challenges involved in doing so.

As described, a primary challenge in OPINEDB is to answer subjective queries over opinion data. A subjective query can be complex involving multiple subjective attributes and objective attributes in addition to filters that restrict the reviews of interest (e.g., prolific reviewers, or reviewers that agree with the user’s taste). To the best of our knowledge, OPINEDB is the first system to answer complex subjective queries over review data in a principled way. Opinion-based entity ranking [16, 32] are the only works that considered subjective queries and utilized reviews for ranking entities. However, that work did not aggregate the reviews or support complex queries like OPINEDB. Trummer et al. [48] note that many queries to web search engines are of subjective nature and consider the problem of aggregating subjective opinions about (entity, attribute) pairs (e.g., cute animals). Aroyo and Welty [3] note that in the process of annotating training data for machine learning, there are several fallacies in assuming that there is an objective truth for the annotations. They develop a measure that supports differing subjective opinions from annotators. Finally, subjective databases are different from probabilistic databases [44] in that the latter still assume that there is an objective ground truth but it is not known to the database.

The second challenge OPINEDB faces is to enable application designers to apply domain semantics to subjective data management. We introduce the concept of marker summaries to provide the designer such flexibility. The designer can tailor the linguistic domains and also which distinctions in the data to highlight in marker summaries. For example, one can have a coarse-grained notion of

bathroom cleanliness (clean vs. dirty) or finer distinctions (shower cleanliness, faucet etc.)

The extractor of OPINEDB for extracting opinion expressions and forming linguistic domains is closely related to opinion mining [31, 34]. The extraction task is known as aspect term extraction [23, 22, 6, 54, 21] and opinion lexicon construction [23, 22, 31, 45, 49, 39, 31, 41, 15, 20, 8] that are well-studied in opinion mining. Following the recent trend of applying deep learning to opinion mining, OPINEDB leverages the BERT pre-trained model [11] and achieved quality surpassing the state of the art while requiring a small amount of human labeling effort.

OPINEDB explores a variant of fuzzy logic to combine the scores of multiple query predicates. Fuzzy logic has been used in a myriad of applications in AI, control theory, and even databases with the capability of reasoning with vague and/or partial predicates like “warm” or “fast” [53, 55, 28, 13]. The efficient evaluation of fuzzy selection queries has been broadly studied in databases, with the Threshold Algorithm [14] and its descendants [24] as the most widely used techniques. In contrast to previous work where fuzzy logic is used to reason about “partial truth” or subjective perception of objective attributes like temperatures or speed, our work considers processing queries on data that is itself subjective.

The problem of building natural language interfaces to databases is a long-standing one [2] and more recent work (e.g., [25, 29, 38, 37]) has focused on learning how to parse natural language into a corresponding semantic form (e.g., SQL) based on examples of pairs of such. OPINEDB does not translate natural language into SQL. Instead, it supports query predicates, which are short phrases, that are already embedded in an SQL-like query. Furthermore, the main focus of these works is on parsing objective queries but OPINEDB interprets and evaluates subjective queries.

## 7. CONCLUSION

As user-generated data becomes more prevalent, it plays a critical role when users make decisions about products and services. However, by nature, user-generated data touches upon subjective aspects of these services for which there is no ground truth. We introduced subjective databases as a key enabling technology for supporting experiential search and built OPINEDB, a first such system. OPINEDB has also been used to power Voyageur, our experiential travel search engine [12]. OPINEDB is based on a new data model that incorporates user-generated data into a database system that can support complex queries, but also gives the designer flexibility to tune the schema for the application needs. We described how OPINEDB processes queries that require semantic interpretation and demonstrated that OPINEDB outperforms alternative approaches.

Subjective databases introduce several new future research challenges. There are many improvements that can be made to how such a system interprets queries specified using natural language. Similarly, a subjective database system should be able to take into consideration a user profile to provide better search results in case the user chooses to share such a profile. In the longer run, the system should be able to suggest queries to the user based on their profile and based on what may be unusual in the domain. For example, if there are reviews claiming that an expensive hotel has dirty rooms, that would be important to point out to the user because it contradicts their expectations. More generally, the challenge is to model the user’s expectations and point out the unexpected experiential aspects. Finally, the topic of bias on the Web is a very timely one [4], and review data clearly contains biases. One of the interesting areas for future research is to use the expressive query capabilities of a system like OPINEDB to uncover biases with the goal of helping users make better decisions about their purchases.

## 8. REFERENCES

- [1] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, et al. Tensorflow: A system for large-scale machine learning. In *OSDI*, pages 265–283, 2016.
- [2] I. Androutsopoulos, G. D. Ritchie, and P. Thanisch. Natural language interfaces to databases - an introduction. *Natural Language Engineering*, 1(1):29–81, 1995.
- [3] L. Aroyo and C. Welty. Truth is a lie: Crowd truth and the seven myths of human annotation. *AI Magazine*, 36(1):15–24, 2015.
- [4] R. A. Baeza-Yates. Bias on the web. *Commun. ACM*, 61(6):54–61, 2018.
- [5] S. Bird and E. Loper. Nltk: the natural language toolkit. In *Proceedings of the ACL 2004 on Interactive poster and demonstration sessions*, page 31, 2004.
- [6] S. Brody and N. Elhadad. An unsupervised aspect-sentiment model for online reviews. In *NAACL HLT*, pages 804–812, 2010.
- [7] M. Buhrmester, T. Kwang, and S. D. Gosling. Amazon’s mechanical turk: A new source of inexpensive, yet high-quality, data? *Perspectives on psychological science*, 6(1):3–5, 2011.
- [8] B. Chen, B. An, L. Sun, and X. Han. Semi-supervised lexicon learning for wide-coverage semantic parsing. In *COLING*, pages 892–904, 2018.
- [9] D. M. Christopher, R. Prabhakar, and S. Hinrich. Introduction to information retrieval. *An Introduction To Information Retrieval*, 151(177):5, 2008.
- [10] A. Conneau, D. Kiela, H. Schwenk, L. Barrault, and A. Bordes. Supervised learning of universal sentence representations from natural language inference data. In *EMNLP*, pages 670–680, 2017.
- [11] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [12] S. Evensen, A. Feng, A. Halevy, J. Li, V. Li, Y. Li, H. Liu, G. Mihaila, J. Morales, N. Nuno, E. Pavlovic, W.-C. Tan, and X. Wang. Voyageur: An experiential travel search engine. In *The World Wide Web Conference, WWW ’19*, pages 3511–5, 2019.
- [13] R. Fagin. Combining fuzzy information from multiple systems. In *PODS*, pages 216–226. ACM, 1996.
- [14] R. Fagin, A. Lotem, and M. Naor. Optimal aggregation algorithms for middleware. *Journal of computer and system sciences*, 66(4):614–656, 2003.
- [15] E. Fast, B. Chen, and M. S. Bernstein. Empath: Understanding Topic Signals in Large-Scale Text. In *CHI*, pages 4647–4657, 2016.
- [16] K. Ganesan and C. Zhai. Opinion-based entity ranking. *Information retrieval*, 15(2):116–150, 2012.
- [17] GitHub. BERT-BiLSTM-CRF-NER. <https://github.com/macanv/BERT-BiLSTM-CRF-NER>, 2018.
- [18] GitHub. OpineDB. [https://github.com/rit-git/opinedb\\_public](https://github.com/rit-git/opinedb_public), 2018.
- [19] C. Gormley and Z. Tong. *Elasticsearch: The Definitive Guide: A Distributed Real-Time Search and Analytics Engine*. " O’Reilly Media, Inc.", 2015.
- [20] W. L. Hamilton, K. Clark, J. Leskovec, and D. Jurafsky. Inducing domain-specific sentiment lexicons from unlabeled corpora. In *EMNLP*, pages 595–605, 2016.
- [21] R. He, W. S. Lee, H. T. Ng, and D. Dahlmeier. An unsupervised neural attention model for aspect extraction. In *ACL*, pages 388–397, 2017.
- [22] M. Hu and B. Liu. Mining and summarizing customer reviews. In *SIGKDD*, pages 168–177, 2004.
- [23] M. Hu and B. Liu. Mining opinion features in customer reviews. In *AAAI*, pages 755–760, 2004.
- [24] I. F. Ilyas, G. Beskales, and M. A. Soliman. A survey of top-k query processing techniques in relational database systems. *ACM Computing Surveys (CSUR)*, 40(4):11, 2008.
- [25] S. Iyer, I. Konstas, A. Cheung, J. Krishnamurthy, and L. Zettlemoyer. Learning a neural semantic parser from user feedback. In *ACL*, pages 963–973, 2017.
- [26] R. Kiros, Y. Zhu, R. R. Salakhutdinov, R. Zemel, R. Urtasun, A. Torralba, and S. Fidler. Skip-thought vectors. In *NIPS*, pages 3294–3302, 2015.
- [27] E. P. Klement, R. Mesiar, and E. Pap. Book review: "triangular norms". *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 11(02):257–259, 2003.
- [28] G. Klir and B. Yuan. *Fuzzy sets and fuzzy logic*, volume 4. Prentice hall New Jersey, 1995.
- [29] F. Li and H. V. Jagadish. Understanding natural language queries over relational databases. *SIGMOD Record*, 45(1):6–13, 2016.
- [30] Y. Li, A. Feng, J. Li, S. Mumick, A. Y. Halevy, V. Li, and W. Tan. Subjective databases. *arXiv preprint arXiv:1902.09661*, 2019.
- [31] B. Liu. *Sentiment Analysis and Opinion Mining*. Morgan Claypool, 2012.
- [32] C. Makris and P. Panagopoulos. Improving opinion-based entity ranking. In *WEBIST*, pages 223–230, 2014.
- [33] T. Mikolov, K. Chen, G. Corrado, and J. Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013.
- [34] M. Pontiki, D. Galanis, H. Papageorgiou, I. Androutsopoulos, S. Manandhar, A.-S. Mohammad, M. Al-Ayyoub, Y. Zhao, B. Qin, O. De Clercq, et al. Semeval-2016 task 5: Aspect based sentiment analysis. In *SemEval-2016*, pages 19–30, 2016.
- [35] M. Pontiki, D. Galanis, H. Papageorgiou, S. Manandhar, and I. Androutsopoulos. Semeval-2015 task 12: Aspect based sentiment analysis. In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pages 486–495, 2015.
- [36] M. Pontiki, D. Galanis, J. Pavlopoulos, H. Papageorgiou, I. Androutsopoulos, and S. Manandhar. Semeval-2014 task 4: Aspect based sentiment analysis. In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, pages 27–35, 2014.
- [37] H. Poon. Grounded unsupervised semantic parsing. In *ACL*, pages 933–943, 2013.
- [38] A. Popescu, O. Etzioni, and H. A. Kautz. Towards a theory of natural language interfaces to databases. In *IUI*, pages 149–157, 2003.
- [39] G. Qiu, B. Liu, J. Bu, and C. Chen. Opinion word expansion and target extraction through double propagation. *COLING*, 37(1):9–27, 2011.
- [40] R. Řehřek and P. Sojka. Gensim-statistical semantics in python. 2011.

- [41] S. Rothe, S. Ebert, and H. Schütze. Ultradense word embeddings by orthogonal transformation. In *NAACL HLT*, pages 767–777, 2016.
- [42] E. F. Sang and F. De Meulder. Introduction to the conll-2003 shared task: Language-independent named entity recognition. *arXiv preprint cs/0306050*, 2003.
- [43] M. Stonebraker and L. A. Rowe. *The design of Postgres*, volume 15. ACM, 1986.
- [44] D. Suciu, D. Olteanu, C. Ré, and C. Koch. *Probabilistic Databases*. Synthesis Lectures on Data Management. Morgan & Claypool Publishers, 2011.
- [45] Y. Tai and H. Kao. Automatic domain-specific sentiment lexicon generation with label propagation. In *IIWAS*, page 53, 2013.
- [46] The Booking.com Dataset. <https://www.kaggle.com/jiashenliu/515k-hotel-reviews-data-in-europe>.
- [47] The Yelp Dataset. <https://www.yelp.com/dataset>.
- [48] I. Trummer, A. Y. Halevy, H. Lee, S. Sarawagi, and R. Gupta. Mining subjective properties on the web. In *SIGMOD*, pages 1745–1760, 2015.
- [49] I. S. Vicente, R. Agerri, and G. Rigau. Simple, Robust and (almost) Unsupervised Generation of Polarity Lexicons for Multiple Languages. In *EACL*, pages 88–97, 2014.
- [50] W. Wang, S. J. Pan, D. Dahlmeier, and X. Xiao. Recursive neural conditional random fields for aspect-based sentiment analysis. In *EMNLP*, pages 616–626, 2016.
- [51] W. Wang, S. J. Pan, D. Dahlmeier, and X. Xiao. Coupled multi-layer attentions for co-extraction of aspect and opinion terms. In *AAAI*, pages 3316–3322, 2017.
- [52] T. Wilson, J. Wiebe, and P. Hoffmann. Recognizing contextual polarity in phrase-level sentiment analysis. In *HLT/EMNLP*, pages 347–354, 2005.
- [53] H. Xin, R. Meng, and L. Chen. Subjective knowledge base construction powered by crowdsourcing and knowledge base. In *SIGMOD*, pages 1349–1361. ACM, 2018.
- [54] X. Yan, J. Guo, Y. Lan, and X. Cheng. A biterm topic model for short texts. In *WWW*, pages 1445–1456, 2013.
- [55] L. A. Zadeh. Fuzzy logic= computing with words. *IEEE transactions on fuzzy systems*, 4(2):103–111, 1996.
- [56] L. Zhang, S. Wang, and B. Liu. Deep learning for sentiment analysis: A survey. *Wiley Interdiscip. Rev. Data Min. Knowl. Discov.*, 8(4), 2018.
- [57] V. Zhong, C. Xiong, and R. Socher. Seq2sql: Generating structured queries from natural language using reinforcement learning. *arXiv preprint arXiv:1709.00103*, 2017.