


Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

# Information Processing and Management

journal homepage: [www.elsevier.com/locate/infoproman](http://www.elsevier.com/locate/infoproman)

## Semantics-enabled query performance prediction for ad hoc table retrieval

Maryam Khodabakhsh<sup>a</sup>, Ebrahim Bagheri<sup>\*,b</sup><sup>a</sup> Computer Engineering Department, Shahrood University of Technology, Shahrood, Iran<sup>b</sup> Laboratory for Systems, Software and Semantics (LS3), Ryerson University, Canada

### ARTICLE INFO

#### Keywords:

Query performance prediction  
Ad hoc table retrieval  
Neural embeddings

### ABSTRACT

Predicting the performance of a retrieval method for a given query is a highly important and challenging problem in information retrieval. Accurate Query Performance Prediction (QPP) plays an important role in real time handling of queries with varying levels of difficulty. While there have been several successful query performance predictors, no predictors have yet been introduced within the context of the ad hoc table retrieval task, which is concerned with answering a query with a ranked list of tables. In this paper, we propose to perform query performance prediction based on neural embedding techniques for ad hoc table retrieval and introduce three neural features. The neural features are based on neural embedding techniques and leverage the distance between tokens in the embedding space in order to capture their semantic similarity. We evaluate our proposed work based on a gold standard test collection and compare it with the state-of-the-art post-retrieval query performance prediction methods. We find that our neural features (1) are effective for predicting the performance of content-based ranking functions; and not as effective for feature-based ranking functions, and (2) show a synergistic impact on existing QPP methods and hence are able to increase their performance in practice.

### 1. Introduction

Within the context of ad hoc retrieval, the effectiveness of information retrieval (IR) methods can significantly vary across different sets of queries and hence the quality of the results returned for some queries can be poorer compared to others. Such queries are often known as ‘difficult’ queries. Thus, early and real-time determination of those queries that are more difficult than others can help IR systems choose appropriate strategies for dealing with them through techniques such as query reformulation (Diaz, 2016).

Query Performance Prediction (QPP) methods are specifically designed to identify the difficulty of a query by estimating the performance of a retrieval system for that given query in the absence of relevance judgments. Query performance prediction has been studied well in the past few years most specifically for the task of ad hoc document retrieval (Arabzadeh, Zarrinkalam, Jovanovic, Al-Obeidat, & Bagheri, 2020; Shtok, Kurland, Carmel, Raiber, & Markovits, 2012; Zhou & Croft, 2007). Empirical studies have shown that the quality of existing predictors can still be improved specially when considering factors such as query ambiguity, missing content, and vocabulary mismatch (Carmel & Kurland, 2012), which are often hard to address because most predictor models are based on features such as the frequency of query terms (He & Ounis, 2004) or the average IDF of query terms in the collection (Tomlinson, 2004). These features are primarily frequency-based term-dependent features and cannot address issues such as vocabulary mismatch.

\* Corresponding author.

E-mail addresses: [m\\_khodabakhsh@shahroodut.ac.ir](mailto:m_khodabakhsh@shahroodut.ac.ir) (M. Khodabakhsh), [bagheri@ryerson.ca](mailto:bagheri@ryerson.ca) (E. Bagheri).<https://doi.org/10.1016/j.ipm.2020.102399>

Received 29 April 2020; Received in revised form 9 September 2020; Accepted 24 September 2020

Available online 06 October 2020

0306-4573/ © 2020 Elsevier Ltd. All rights reserved.

While existing work in query performance prediction has been predominantly focused on various statistical measures based on term frequency and lexical matching, the information retrieval community has recently embarked on exploring the impact and importance of neural information retrieval techniques (Bagheri, Ensan, & Al-Obeidat, 2018; Mitra, Craswell et al., 2018; Zuccon, Koopman, Bruza, & Azzopardi, 2015) or latent models (Furnas et al., 2017; Shen, He, Gao, Deng, & Mesnil, 2014) where a document that does not overlap in terminology with the query can still be retrieved. The benefit of such methods is that they consider semantic association between the document and query spaces and hence are able to bridge the gap between these two spaces when the vocabulary mismatch problem is observed (Furnas et al., 2017; Mitra et al., 2018; Shen et al., 2014; Zuccon et al., 2015). Empirical experimentation has shown that such features are strong complementary features for term-based features and can lead to noticeable improvement in ad hoc retrieval (Ensan & Bagheri, 2017). While such features have been widely adopted in the recent years for performing retrieval, they have only been investigated in the context of ad hoc *document* retrieval or QPP for ad hoc *document* retrieval.

In this paper, we are interested specifically in query performance prediction in ad hoc *table* retrieval (Bagheri & Al-Obeidat, 2020; Zhang & Balog, 2018a). The task of ad hoc table retrieval is concerned with the identification of the most relevant full tables to a given input query. The effective retrieval of relevant tables is important as it facilitates the identification of information that are structured in tabular form. These include tabular data from the financial sector, demographic information, catalogues and even database relational tables, just to name a few. As argued by Zhang and Balog (2018a), the effective retrieval of tables has strong impact on other dependent tasks such as table completion and table mining. While ad hoc table retrieval is a subset of the ad hoc retrieval task, it can be considered to be more challenging than ad hoc document retrieval, as unlike documents, data in a table lack continuity and context that would be helpful for measuring relevance. Therefore, ad hoc table retrieval has been gaining increasing attention from the research community in the recent years (Bagheri & Al-Obeidat, 2020; Deng, Zhang, & Balog, 2019; Gentile, Ristoski, Eckel, Ritze, & Paulheim, 2017; Herzig, Nowak, Muller, Piccinno, & Eischenschlos, 2020; Zhang & Balog, 2018a; 2020).

Recent work in ad hoc table retrieval (Bagheri & Al-Obeidat, 2020) has shown that various retrieval methods can have complementary behavior. In other words, the best table retrieval methods in the literature do not necessarily perform better on all range of queries. Therefore, it is essential to be able to predict which method has the best performance on any given input query to be able to effectively route the query to the best retrieval method. As such, our work in this paper is focused on query performance prediction in the context of ad hoc table retrieval, which has not been explored in the literature. We will investigate how existing post-retrieval query performance predictors perform in the context of ad hoc table retrieval and subsequently systematically propose to integrate semantic neural information into performance prediction for building stronger predictors.

To this end, our work rests on and explores a fundamental idea that given the user's query and a retrieved document set, features that are based on some notion of frequency would not be sufficient features for predicting the performance of the retrieval method. It is our hypothesis that performance is appropriately estimated if semantic features are also taken into consideration. For this purpose, we predict query performance based on well-known neural embedding techniques (Mikolov, Chen, Corrado, & Dean, 2013a; Mikolov, Sutskever, Chen, Corrado, & Dean, 2013b; Pennington, Socher, & Manning, 2014) to measure the neural characteristics of the retrieved document set and the query.

Our proposed neural predictors have two main components: 1) in the first component, we define a representation function to transform the raw content of queries and tables into a bag of tokens (tokens can be keywords or entities), and 2) in the second component, various neural features based on neural embedding techniques are proposed and integrated into the state-of-the-art post-retrieval predictor models. The key contributions of our work can be summarized as follows:

1. We systematically investigate the impact of the state-of-the-art post-retrieval query performance prediction methods in the context of ad hoc table retrieval;
2. We show how neural features derived from neural embedding techniques can be integrated into existing query performance predictors, leading to neural predictors;
3. We evaluate our proposed neural predictors on a gold standard test collection derived from Wikipedia tables. We further analyze our findings from various perspectives and identify areas where neural predictors are able to provide effective performance improvements.

In our experiments, we will explore and systematically answer four research questions, including (1) whether the representation of table content in the form of keywords or entities have an impact on determining query difficulty; (2) would it be possible to improve the estimation power of existing post-retrieval query performance predictors using neural predictors that are proposed in this paper; (3) Would the proposed neural predictors have a complementary impact on existing QPP methods; and (4) Whether the consideration of table specific structural characteristics such as table cells, rows and columns have any impact on the performance of our proposed neural predictors.

The rest of this paper is organized as follows. In the next section, we review the related work in two subsections covering pertinent work in ad hoc table retrieval and query performance prediction techniques. Section 3 provides the technical details of our proposed neural predictors. The details of our experiments, gold standard test collection, baselines and our findings are presented in Section 4. Finally, the paper is concluded in Section 5 with a summary of our findings.

## 2. Related work

In this section, we first review prior work on ad hoc table retrieval, and then present the related literature on query performance prediction.

## 2.1. Ad hoc table retrieval

Tables are suitable tools for grouping large volumes of information in one simplified structure. They are widely adopted for reporting data on the Web and have found a special place on Wikipedia. Various researchers have already explored various aspects of table analytics such as mining (Bhagavatula, Noraset, & Downey, 2013), searching (Cafarella, Halevy, Wang, Wu, & Zhang, 2008; Pimplikar & Sarawagi, 2012; Zhang & Balog, 2018a), extracting (Cafarella et al., 2008), completing (Zhang & Balog, 2019) and generating (Zhang & Balog, 2018b) tables in Web documents. The focus of our paper is on the task of searching for and retrieving tables in response to a given unstructured query, often known as ad hoc table retrieval. A full account of the relevant state of the art can be found in a very recent paper on this topic (Zhang & Balog, 2020). Existing work in ad hoc table retrieval can be categorized from two perspectives: (1) the types of the *ranking function* used for ranking tables, and (2) the *features* used in the retrieval methods. We review the related literature from these two views:

**Ranking Functions.** The objective of a ranking function is to rank a set of documents based on their relevance to the query, which can be in the form of feature-based learn to rank methods (Capannini et al., 2016; Tax, Bockting, & Hiemstra, 2015), statistical and probability techniques such as language models (Metzler & Croft, 2004), entity linking-based techniques (Das Sarma et al., 2012; Limaye, Sarawagi, & Chakrabarti, 2010), models that rely on diversity of table content (Nguyen, Nguyen, Weidlich, & Aberer, 2015), network properties (Yakout, Ganjam, Chakrabarti, & Chaudhuri, 2012) or knowledge graph-based similarity (Zhang & Chakrabarti, 2013). More recent techniques capture the semantic aspect of the tables and queries in the process of ranking. For instance, Zhang and Balog (2018a) were the first to develop a ranking method to match queries and tables based on their semantic association where both queries and tables were represented using semantic concepts (bag-of-entities and bag-of-categories) as well as continuous dense vectors (word and graph embeddings). Additional work was then subsequently developed to learn neural representations for tables for the purpose of retrieval (Deng et al., 2019).

**Features.** Ranking functions rely on a set of features extracted from the query and table spaces. Some of these features are *table-specific*, which means that they represent the specific properties of a table such as the number of empty table cells in the table or the number of times a table has been viewed in the past (Bhagavatula et al., 2013). Another group of features are *table-query interdependent* features, which measure the association of query and table pairs, e.g., total query term frequency in the table body (Cafarella et al., 2008). The third type of feature focuses on measuring *query-specific* characteristics such as the length of the query (Tyree, Weinberger, Agrawal, & Paykin, 2011). This third type of feature is often the least discriminative form of features for ranking.

## 2.2. Query performance prediction

Query performance prediction models can be classified into pre-retrieval and post-retrieval models (Carmel & Yom-Tov, 2010). Pre-retrieval models predict the quality of the search results using only the raw query, and its descriptive features gathered at indexing time, e.g., the frequency of individual query terms in the target corpus (He & Ounis, 2004). Unlike pre-retrieval models where the predictions take place before the retrieval process, post-retrieval models analyze the search results of a query after retrieval in order to make a judgement about the difficulty of the query. Extensive research in the literature has already shown that post-retrieval methods for QPP are much stronger compared to pre-retrieval methods (Cummins, Lalmas, O’Riordan, & Jose, 2011b; Déjean, Ionescu, Mothe, & Ullah, 2020). As such, the focus of our work in this paper will be on post-retrieval models.

There are three broad categories of post-retrieval QPP models in the literature, namely Clarity-oriented models, Robustness-oriented models and Score-oriented models.

Clarity-oriented models (Cronen-Townsend, Zhou, & Croft, 2002) predict query performance by measuring the clarity of the returned results with respect to the corpus. The hypothesis of these models is that users expect the results to focus on the specific topics covered by the query; thus, the more related keywords to the query are observed in the result set, the less difficult the query would be. To this end, Clarity computes the difference between the likelihood of the most frequent keywords in the results and their likelihood in the whole corpus using measures such as the Kullback-Leibler divergence. If the language model of the results are similar to the language model of the whole corpus, then one can conclude that the query is difficult as it is not discriminatory enough.

Robustness-oriented models estimate query performance by measuring robustness with respect to perturbations in the query (Zhou & Croft, 2007), in the documents (Zhou & Croft, 2006), and in the retrieval method (Aslam & Pavlu, 2007). Intuitively, the more robust the result list is, the higher the performance of the query would be. One of the most popular robustness models is Query Feedback (QF) (Zhou & Croft, 2007) that rebuilds a new query based on the result list. The correlation between the previous result list and the retrieved result list for the new query is an indication of robustness, which is then used to estimate query performance.

Measuring clarity and robustness is often a time-consuming process, which is due to the need to analyze the content of the result list. Less expensive alternatives include Score-oriented models, which use the score distribution of the retrieved documents. The basic idea behind these models is that the more a document is similar to a query, the higher the score of the document would be. Therefore, high scores in the top ranked documents are an indication of better retrieval performance and lower query difficulty. Weighted Information Gain (WIG) (Zhou & Croft, 2007) and Normalized Query Commitment (NQC) (Shtok et al., 2012) are two popular measures based on retrieval score. WIG essentially measures the divergence between the mean retrieval score of the top ranked documents and that of the whole corpus. It was shown that WIG is very effective for performing QPP on an MRF-based retrieval framework (Metzler & Croft, 2005). On the other hand, NQC measures the standard deviation of the retrieval scores normalized by the corpus score. Another Score-oriented model is Score Magnitude and Variance (SMV) (Tao & Wu, 2014) that considers both magnitude and variance of scores of the ranked list of results to measure the performance of a query.

In addition to the above three types of predictors, there are also other works, such as fusion based retrieval (Shaw & Fox, 1994),

wherein several ranked lists, each one retrieved by a different ranking function, are combined into a single fused ranked list by a fusion method (Markovits, Shtok, Kurland, & Carmel, 2012; Roitman, 2018). NeuralQPP, from Zamani, Croft, and Culpepper (2018), integrates three neural components, the component for top-k retrieval score, another for term distribution in the top retrieved documents and the last for representation of documents in the semantic space.

We would like to note that our work distinguishes itself from the existing literature as it (1) considers predicting query performance for ad hoc table retrieval, which has not been investigated in the past; and (2) it integrates information from neural embeddings and integrates it with existing QPP predictors, which is also a novel aspect of our work that has not been explored in the past.

### 3. Proposed neural query performance predictors

Let us begin by setting out the notation. Let  $q$ ,  $C$  and  $R$  denote a query, a corpus, and a ranking function, respectively. A ranking function maps each table in  $C$  and query  $q$  into a real number, which represents the relevance degree. We use  $\pi_R^k(q; C)$  to denote the list of the top  $k$  tables ranked by  $R$  in response to the query  $q$ . Our goal is to establish a neural predictor,  $NP$ , to quantify the effectiveness of  $\pi_R^k(q; C)$  when no relevance judgments are available:

$$NP(\pi_R^k(q; C)) \stackrel{\text{def}}{=} P_{\text{base}}^{NF} \circ \varphi(\pi_R^k(q; C)) = P_{\text{base}}^{NF}(\varphi(\pi_R^k(q; C))) \quad (1)$$

$NP(\pi_R^k(q; C))$  is based on a function composition and quantifies the effectiveness of  $\pi_R^k(q; C)$  based on some form of semantic information that can give way to inexact matching. The outer function, denoted as  $P_{\text{base}}^{NF}$ , can be any state-of-the-art post-retrieval query performance predictor. The predictor  $P_{\text{base}}^{NF}$  operates upon  $\varphi(\pi_R^k(q; C))$  by employing semantic features. In our work, semantic features are implemented as *neural features* (NF) based on neural embedding techniques (Mikolov et al., 2013a). The inner function, denoted by  $\varphi$ , is a representation function that transforms query  $q$  and the tables in  $\pi_R^k(q; C)$  to tokens such as keywords or entities. In the following subsections, we describe (1) the representation function  $\varphi$ , and (2) the connection between existing query performance predictors and neural features in more detail.

#### 3.1. Representation function $\varphi$

In this subsection, we introduce the representation function  $\varphi$  to model the raw content of queries and the corpus as a bag of tokens where each token can be either a keyword or an entity. In the context of ad hoc retrieval, corpus  $C$  is a collection of documents, where documents can be web pages (Cummins, Jose, & O’Riordan, 2011a), news articles (Cummins et al., 2011a), web tables (Cafarella et al., 2008), or images (Siddiquie, Feris, & Davis, 2011), among others. In this work, the documents are tables (Bhagavatula, Noraset, & Downey, 2015). Also, query  $q$  represents the information need of the user.

*Keyword-based representation.* Given  $\pi_R^k(q; C)$ , for each table  $T \in \pi_R^k(q; C)$ , function  $\varphi$  takes the keywords in  $T$  and represents them as a bag of words (BOW).

*Entity-based representation.* In contrast to keyword-based representation, it is also possible to represent a table as a set of entities. Zhang and Balog show that some tables have an entity focus (Zhang & Balog, 2017), which means that significant portions of the table have corresponding entity representations. Therefore, function  $\varphi$  uses entities of each table  $T \in \pi_R^k(q; C)$  to represent the content of  $T$  as a bag of entities (BOE). The function  $\varphi$  can also transform query  $q$  to a bag of entities.

#### 3.2. The connection between QPP and neural features

As mentioned earlier, query performance predictors, such as WIG and Clarity, utilize term probability, which we denote as  $p(\omega)$ , where  $\omega$  is considered in the context of a query, table or corpus in a strict term matching sense and disregard semantically related terms when computing term probability. As a result such predictors can suffer from not paying attention to the semantic association between terms, queries and tables. The objective of the outer function in Eq. (1), i.e.,  $P_{\text{base}}^{NF}$ , is to enable the inclusion of such semantic associations through measuring term similarities based on neural features.

In this paper, we introduce three main neural features (NFs) by using neural embedding techniques: (1) Neural Matching (NM), (2) Neural Aggregated Matching (NAM) and (3) Neural Distance (ND) and explain them in more detail in the following.

##### 3.2.1. The neural matching (NM) feature

One of the important applications of embedding techniques in IR is to employ the dense vector representations that they offer to model queries and documents (tables in our case) from the embeddings of the individual tokens that appear in each of them. This requires a systematic way to aggregate the embeddings of the individual tokens in the queries or documents (tables) to represent the whole query or document. Average Word (or token) Embedding (AWE) is a popular way for performing the aggregation (Le & Mikolov, 2014; Nalisnick, Mitra, Craswell, & Caruana, 2016; Vulić & Moens, 2015). Given table  $T$  and query  $q$ , AWE defines their aggregated representations as follows:

$$\vec{T} = \frac{1}{|\varphi(T)|} \sum_{t_T \in \varphi(T)} \frac{\vec{t}_T}{\|\vec{t}_T\|} \quad (2)$$

$$\vec{q} = \frac{1}{|\varphi(q)|} \sum_{t_q \in \varphi(q)} \frac{\vec{t}_q}{\|\vec{t}_q\|} \quad (3)$$

where  $\varphi(T)$  and  $\varphi(q)$  are the output of the representation function  $\varphi$  that can be either bag of words or bag of entities for table  $T$  and query  $q$ , respectively. The table vector, denoted by  $\vec{T}$ , is the centroid of all the normalized vectors for token  $t_T$  in the table and serves as a single embedding for the whole Table. Similarly,  $\vec{q}$  is the query vector.

Inspired by Eq. (2), we define the corpus vector as  $\vec{C} = \frac{1}{|\varphi(C)|} \sum_{t_C \in \varphi(C)} \frac{\vec{t}_C}{\|\vec{t}_C\|}$  where  $\varphi(C)$  is the tokens of corpus  $C$  denoted as  $\varphi(C) = \cup_{T \in C} \varphi(T)$ . In other words,  $\varphi(C)$  is a bag of all of the table tokens.

In ad hoc table retrieval, the relevance of the retrieved tables to the input query is of foremost importance. Now, given that we have built vectors for both tables as well as queries, our objective is to measure the degree of relevance between them. Therefore, it is possible to directly calculate the similarity between a given table and the query vector representation. We assume table  $T$  is related to query  $q$  if the table vector is semantically similar to the query vector. Our proposed Neural Matching (NM) feature measures this semantic similarity using a vector similarity metric, such as cosine similarity or dot-product.

The intuition behind this feature is based on how tokens are adopted to express the content of a table. The assumption is that if the set of tokens representing a table is accurately selected, then any other query submitted from the users will use tokens that will be semantically close to the table representation in the embedding space. Therefore, given the fact that the position of a table can be determined based on the centroid of the position of its constituent tokens within the embedding space and also the position of a query can be determined in the same way in the same space, it is possible to measure the relevance of the tables to the queries using our proposed Neural Matching (NM) feature as follows:

**Definition 1 (Neural Matching Feature).** Given  $\vec{q}$  and  $\vec{T}$ , neural matching is equivalent to the cosine similarity between  $\vec{q}$  and  $\vec{T}$  and formally denoted as  $NM(q, T)$ .

It is worth mentioning that the neural matching feature can be used for measuring the relevance of corpus  $C$  to query  $q$  given  $\vec{q}$  and  $\vec{C}$ , denoted as  $NM(q, C)$  or any two arbitrary vectors.

Now, given the neural matching feature, we are interested in using it in post-retrieval predictors in order to define  $P_{base}^{NM}$ . For this purpose, we adopt the widely used post-retrieval predictors reported in the literature as shown in Table 1 where  $NM$  has been integrated into the base post-retrieval predictors.

As shown in Table 1, Clarity (Cronen-Townsend et al., 2002) and WIG (Zhou & Croft, 2007) require the probability of the query or corpus tokens in the tables or the corpus. Our proposed Neural Matching feature is an appropriate alternative for this probability as it measures the degree of relevance of each term to the table or the corpus. Other predictors based on standard deviation, such as NQC (Shtok et al., 2012) and SMV (Tao & Wu, 2014), employ a Score-oriented approach  $Score(q, T)$ , which indicates the score of table  $T$  for query  $q$ . By considering the fact that if a table achieves high relevance score for a given query, then the keywords or entities of the table will be semantically close to the keywords or entities used in the query in the embedding space, we revise these predictors by using  $NM(q, T)$  instead of  $Score(q, T)$ .

It is worth noting that the vectors of the tables and the corpus can be calculated offline. Therefore, we will only need to measure the neural matching feature at runtime.

### 3.2.2. Neural aggregated matching (NAM) feature

In this feature, we model each table as a bag of token vectors where vectors of tokens incorporate both syntactic and semantic aspects of the table content and therefore show to be effective for measuring the relevance degree. We adopt the same consideration for queries and the corpus and model them as a bag of token vectors. This is unlike the proposed NM feature where the queries, tables and the corpus are represented by a vector in embedding space. Now, based on the bag of token vectors, we formally define the

**Table 1**

The neural matching feature integrated into post-retrieval query performance predictors.  $T$  is a table in corpus  $C$  and  $V(\varphi(C))$  is the vocabulary of the unique tokens in the corpus  $C$ .  $\lambda(t_q)$  reflects the relative weight of the tokens type  $t_q$  and is inversely related to the square root of the number of query tokens of that type.  $\mu$  is the mean score of tables in  $\pi_R^k(q; C)$ . In the main Clarity formula Cronen-Townsend et al. (2002),  $\Pr(term | \pi_R^k(q; C))$  is the language model induced from the result list  $\pi_R^k(q; C)$  and equals the sum over all tables in the result list. Here, it is replaced with  $NM(t_C, \pi_R^k(q; C))$  in  $P_{Clarity}^{NM}$ .  $NM(t_C, \pi_R^k(q; C)) = \sum_{T \in \pi_R^k(q; C)} NM(t_C, T) (\prod_{t_q \in \varphi(q)} NM(t_q, T))$ .

Base QPP Method	$P_{base}^{NM}$ : Revised Post-retrieval Query Performance Predictors with Neural Matching Feature (NM)
Clarity Cronen-Townsend et al. (2002)	$P_{Clarity}^{NM}(\varphi(\pi_R^k(q; C))) = \sum_{t_C \in V(\varphi(C))} NM(t_C, \pi_R^k(q; C)) \log \frac{NM(t_C, \pi_R^k(q; C))}{NM(t_C, C)}$
WIG Zhou and Croft (2007)	$P_{WIG}^{NM}(\varphi(\pi_R^k(q; C))) = \frac{1}{k} \sum_{T \in \pi_R^k(q; C)} \sum_{t_q \in \varphi(q)} \lambda(t_q) \log \frac{NM(t_q, T)}{NM(t_q, C)}$
NQC Shtok et al. (2012)	$P_{NQC}^{NM}(\varphi(\pi_R^k(q; C))) = \frac{\frac{1}{k} \sum_{T \in \pi_R^k(q; C)} (NM(q, T) - \mu)^2}{INM(q, C)}$
SMV Tao and Wu (2014)	$P_{SMV}^{NM}(\varphi(\pi_R^k(q; C))) = \frac{\frac{1}{k} \sum_{T \in \pi_R^k(q; C)} (NM(q, T)   \ln \frac{NM(q, T)}{\mu})}{NM(q, C)}$



**Table 2**

The neural aggregated matching feature integrated into post-retrieval query performance predictors. Table format is similar to Table 1. Also,  $NAM(t_C, \pi_R^k(q; C)) = \sum_{T \in \pi_R^k(q; C)} NAM(t_C, T) (\prod_{t_q \in \varphi(q)} NAM(t_q, T))$ .

Base QPP Method	$P_{base}^{NAM}$ : Revised Post-retrieval Query Performance Predictors with Neural Aggregated Matching Feature (NAM)
Clarity Cronen-Townsend et al. (2002)	$P_{Clarity}^{NAM}(\varphi(\pi_R^k(q; C))) = \sum_{t_C \in V(\varphi(C))} NAM(t_C, \pi_R^k(q; C)) \log \frac{NAM(t_C, \pi_R^k(q; C))}{NAM(t_C, C)}$
WIG Zhou and Croft (2007)	$P_{WIG}^{NAM}(\varphi(\pi_R^k(q; C))) = \frac{1}{k} \sum_{T \in \pi_R^k(q; C)} \sum_{t_q \in \varphi(q)} \lambda(t_q) \log \frac{NAM(t_q, T)}{NAM(t_q, C)}$
NQC Shtok et al. (2012)	$P_{NQC}^{NAM}(\varphi(\pi_R^k(q; C))) = \frac{\sqrt{\frac{1}{k} \sum_{T \in \pi_R^k(q; C)} (NAM(q, T) - \mu)^2}}{ NAM(q, C) }$
SMV Tao and Wu (2014)	$P_{SMV}^{NAM}(\varphi(\pi_R^k(q; C))) = \frac{\sqrt{\frac{1}{k} \sum_{T \in \pi_R^k(q; C)} (NAM(q, T) + \ln \frac{NAM(q, T)}{\mu})}}{NAM(q, C)}$

Neural Aggregated Matching (NAM) feature as follows:

**Definition 2 (Neural Aggregated Matching Feature).** Given  $\varphi(q)$  and  $\varphi(T)$ , neural aggregated matching is the maximum cosine of the token vector pairs in the  $\varphi(q)$  and  $\varphi(T)$  denoted as  $NAM(q, T)$ :

$$NAM(q, T) = \max_{t_q \in \varphi(q), t_T \in \varphi(T)} (\{\cos(\vec{t}_q, \vec{t}_T)\}) \quad (4)$$

In this feature, the pairwise similarities between all table and query token vectors are computed and aggregated by a maximum function, which is a common approach to predict similarity (Kenter & De Rijke, 2015). Our assumption is that the maximum similarity captures the minimum cost required to match the query to the table.

The adopted formulation of the post-retrieval predictors based on NAM are provided in Table 2 where NAM is placed in the base post-retrieval predictors. This feature can determine the relevance between a query  $q$  and corpus  $C$  by giving  $\varphi(q)$  and  $\varphi(C)$ , as  $NAM(q, C)$ . In some cases, e.g., Clarity (Cronen-Townsend et al., 2002) and WIG (Zhou & Croft, 2007), where the calculation of the probability of a token in the table or corpus is required, neural aggregated matching feature computes the cosine similarity between the token vector and all the table or corpus token vectors first, and then aggregates those through a maximum function.

### 3.2.3. Neural distance (ND) feature

An alternative approach to estimate the relevance of a table and a query is based on the distance between them (Bagheri et al., 2018; Galke, Saleh, & Scherp, 2017) by considering the fact that the lower the distance between a query and a table is, the greater their similarity would be and the more relevant they would be to each other. A recent method, based on Earth Mover's Distance (EMD) (Rubner, Tomasi, & Guibas, 1998), involves estimating the similarity between pairs of documents by minimizing the amount of distance that each embedded term in the first document needs to 'travel' to reach the embedded terms in the second document. This measure, referred to as the Word Mover's Distance (WMD), was proposed by Kusner, Sun, Kolkin, and Weinberger (2015). In this feature, given  $T$  and  $q$ , the distance between them will be based on tokens traveling from  $q$  to tokens in  $T$ . As defined in Kusner et al. (2015), the transportation matrix  $\theta$  is a flow matrix in which  $\theta_{i,j}$  shows to what degree token  $i$  in  $q$  is transported to token  $j$  in  $T$ . Matrix  $\theta_{i,j}$ , which essentially determines what token pairs from the two documents should be connected to each other, needs to be learnt based on a linear optimization program. The distance between two documents can be calculated by minimizing the following linear optimization function:

$$WMD(q, T) = \sum_{i=1}^{|\varphi(q)|} \sum_{j=1}^{|\varphi(T)|} \theta_{i,j} d(i, j) \quad (5)$$

where  $d(i, j)$  is the distance between the vector of token  $i$  in  $q$  and the vector of token  $j$  in  $T$  in the embedding space. Our proposed Neural Distance (ND) feature estimates the relevance degree of table  $T$  to query  $q$  based on the minimum value computed from optimizing Eq. (5). As such, we formally define the Neural Distance (ND) feature as follows:

**Definition 3 (Neural Distance Feature).** Given  $q$  and  $T$ , neural distance feature is equivalent to the Word Mover's Distance between the vectors of tokens in  $q$  and the vectors of tokens in  $T$ , denoted as  $ND(q, T)$ .

Table 3 shows how to integrate the Neural Distance feature in post-retrieval predictors.

## 4. Experiments

The objective of our experiments is to evaluate the effectiveness of our proposed neural predictors on the task of query performance prediction for ad hoc table retrieval. To this end, we define four research questions that will be answered through performing empirical experimentation. More specifically, we address the following research questions in our experiments:

**RQ1** Given the *Representation function*  $\varphi$ , which can provide tokens in the form of either keywords or entities to represent queries and tables, which type of token would be most effective for determining query difficulty?

**Table 3**

The neural distance feature integrated into post-retrieval query performance predictors. Table format is similar to Table 1. Also,  $ND(t_C, \pi_R^k(q; C)) = \sum_{T \in \pi_R^k(q; C)} ND(t_C, T) (\prod_{t_q \in \varphi(q)} ND(t_q, T))$ .

Base QPP Method	$P_{base}^{ND}$ : Revised Post-retrieval Query Performance Predictors with Neural Distance Feature (ND)
Clarity Cronen-Townsend et al. (2002)	$P_{Clarity}^{ND}(\varphi(\pi_R^k(q; C))) = \sum_{t_C \in V(\varphi(C))} ND(t_C, \pi_R^k(q; C)) \log \frac{ND(t_C, \pi_R^k(q; C))}{ND(t_C, C)}$
WIG Zhou and Croft (2007)	$P_{WIG}^{ND}(\varphi(\pi_R^k(q; C))) = \frac{1}{k} \sum_{T \in \pi_R^k(q; C)} \sum_{t_q \in \varphi(q)} \lambda(t_q) \log \frac{ND(t_q, T)}{ND(t_q, C)}$
NQC Shtok et al. (2012)	$P_{NQC}^{ND}(\varphi(\pi_R^k(q; C))) = \frac{\sqrt{\frac{1}{k} \sum_{T \in \pi_R^k(q; C)} (ND(q, T) - \mu)^2}}{IND(q, C)}$
SMV Tao and Wu (2014)	$P_{SMV}^{ND}(\varphi(\pi_R^k(q; C))) = \frac{\sqrt{\frac{1}{k} \sum_{T \in \pi_R^k(q; C)} (ND(q, T) + \ln \frac{ND(q, T)}{\mu})}}{ND(q, C)}$

**RQ2** Based on the three proposed neural features, would it be possible to improve the estimation power of existing post-retrieval query performance predictors?

**RQ3** Would the proposed neural predictors have a complementary impact on existing QPP methods and hence lead to improved performance when systematically interpolated with baseline methods?

**RQ4** Whether the consideration of table specific structural characteristics such as table cells, rows and columns have any impact on the performance of our proposed neural predictors?

Summarily, RQ1 will be investigating which form of query, table and corpus representation, through either keywords or entities, be more effective in the context of the QPP task. RQ2 explores whether the integration of the proposed neural features into existing QPP methods leads to improved performance prediction. RQ3 will explore the synergistic nature of the neural features and existing keyword-based QPP methods to see whether their interpolation will lead to improved performance and RQ4 investigates whether the explicit consideration of table characteristics will lead to improved performance by the QPP method.

In the following, we first introduce the experimental settings, including the ranking functions, the test collection and the evaluation metrics. We then report and analyze our results and present the most important findings in order to answer the research questions.

#### 4.1. Test collection

We employed the test collection introduced in Zhang and Balog (2018a) in our experiments. This test collection includes three main components, namely the table corpus, the query set and the relevance judgments. The table corpus is composed of the Wiki-Tables corpus (Galke et al., 2017) containing over 1.6M tables. For each table, five information fields are provided: table caption, column headings, data rows, page title, and section title. To represent the tables as the bag of entities, we use the same entities that were extracted by Zhang and Balog (2018a). The query set consists of 60 queries, which were sampled from two independent sources. We also used query entities as provided in Bagheri and Al-Obeidat (2020). The relevance judgments consist of 3120 query-table pairs. Out of these, 377 are labeled as highly relevant, 474 as relevant, and 2269 as non-relevant.

#### 4.2. Ranking functions

Given a query  $q$ , a ranking function  $R$  assigns a real number to each table  $T \in C$  that represents the relative relevance of  $T$  with respect to  $q$ . In the experiments, we use seven different ranking functions to serve as  $R$ . These ranking functions can be broadly classified as (1) *content-based ranking functions* that focus on table content for measuring relevance between the query and table spaces; and (2) *feature-based ranking functions* that define table and query characteristics that take the structural characteristics of tables into account.

##### 1. Content-based Ranking Functions:

- STR ranking (Zhang & Balog, 2018a) models both the table and the query as sets of semantic vectors and then uses two general strategies (early and late fusion), for computing the similarity between queries and tables based on their semantic representations.
- Single Field ranking (Cafarella, Halevy, & Khoussainova, 2009) represents tables as an ordinary document and adopts ad hoc document retrieval, in the form of a language model with Dirichlet smoothing, for ranking tables.
- Table2VecW and Table2VecE (Deng et al., 2019) are content-based ranking functions that learn neural representations for tables. The difference between the two ranking functions is based on the content they consider in the table to learn the neural representation. Table2VecW benefits from title, section title, table caption, table headings, and all table cells, while Table2VecE employs the sequence of all entities that appear within table cells.

##### 2. Feature-based Ranking Functions:

- LTR ranking (Zhang & Balog, 2018a) uses the full set of features listed in (Bhagavatula et al., 2013; Cafarella et al., 2008) and trains a Random Forest (Breiman, 2001) regression as the learner with 1000 trees.

- (b) WebTable ranking (Cafarella et al., 2008) employs the features mentioned in Cafarella et al. (2008) for each table and trains linear regression as the ranking function with 5-fold cross validation.
- (c) WikiTable ranking (Bhagavatula et al., 2013) considers a Lasso model (Tibshirani, 1996) with coordinate ascent as the ranking function, which is trained on the features from Bhagavatula et al. (2013) with 5-fold cross validation.

We note that the selected ranking functions are either (1) state of the art (SOTA) methods in ad hoc table retrieval such as STR and Table2VecW and Table2VecE, (2) have been used as baseline in several recent publications such as WikiTable and WebTable or (3) are standard information retrieval methods that have shown robust retrieval performance across different tasks such as Single Field. The run files for all the ranking functions, including content-based and feature-based functions are readily available<sup>1</sup> Accordingly, the result list length per query provided by the run files is 20. We report the performance of the ranking functions based on the run files in Table 5.

#### 4.3. Evaluation metrics

The prediction quality of a QPP method is evaluated by measuring Pearson and Kendall Correlations between the predicted ranking of the queries against the actual ranking of the queries according to their observed performance measured on NDCG and AP Järvelin and Kekäläinen (2002).

#### 4.4. Baselines

For comparative analytics, we adopt Clarity (Cronen-Townsend et al., 2002), WIG (Zhou & Croft, 2007), NQC (Shtok et al., 2012) and SMV (Tao & Wu, 2014), which are among the most widely used query performance predictors. The detailed formulation of these predictors is provided in Table 4. We note that these metrics are implemented according to the detailed description provided in Carmel and Yom-Tov (2010) and do not include any relation to neural embeddings.

Unlike these four baseline predictors that are based on a completely unsupervised formulation of the predictor, the most recent post-retrieval query performance predictor is work by Zamani et al. (2018) that uses a weak supervision strategy to learn query rankings. As such, while we will compare our work with this baseline QPP method, we cannot use it to integrate our approach with it.

It should be noted that given our approach requires neural embeddings, in our experiments, we use the pre-trained embedding vectors from Li et al. (2016) for neural embeddings of entities and the Google News embeddings<sup>2</sup> for neural embeddings of keywords. For BERT, we use the pretrained model offered by HuggingFace with 12 layers, 768 hidden layers, 12 heads, and 110M parameters trained on cased English text, often known as BERT-Base Cased. For ELMo, we have used the model provided by Flair Akbik, Blythe, and Vollgraf (2018).

#### 4.5. Analysis

We present the findings of our experiments based on the four research questions introduced earlier.

##### 4.5.1. RQ1. Impact of the representation function $\varphi$

The objective of the first research question is to determine which token representation is most suitable for modeling tables and queries when performing query performance prediction in the context of ad hoc table retrieval. The representation function  $\varphi$  can represent tables and queries in the form of keywords or entities. The idea behind a bag of words representation is that lexical similarity between tables and queries can be considered to be an indication of relevance. On the other hand, the intuition behind the bag of entities representation is that the relevance of a table to a query can be captured beyond mere keywords and can depend on the semantic association between the two, which might not necessarily have much lexical overlap. Bagheri and Al-Obeidat have already shown this by providing examples of where queries and tables that do not have any lexical overlap are considered to be relevant as assessed in human judgements (Bagheri & Al-Obeidat, 2020). For instance, for Query 30, ‘pain medication’, in the query set of our test collection, the human judgements had determined that Table 0520-188 ‘diseases and condition’ was relevant but the query and table did not have any shared keywords. Therefore, there is existing evidence to support that the representation of tables as bag of entities can be beneficial. In this research question, we are interested in understanding the impact of each representation on query performance prediction.

To this end, we compare the performance of our neural predictors on different ranking functions by using two different representations: bag of words and bag of entities. The results are shown in Tables 6–9. The tables present the results for NDCG and AP for both Kendall and Pearson correlations. In these tables, the best predictor for the combination of each QPP method and ranking function is shown in bold. As seen in the table, performances for both measures, i.e., NDCG and AP, as well as correlations, i.e., Kendall and Pearson, are consistent.

First, we compare the quality of our proposed neural predictors for both bag of words and bag of entities representations. The

<sup>1</sup> The runs for STR, LTR, WikiTable, WebTable and Single Field ranking functions can be obtained from <https://github.com/iai-group/www2018-table> while the runs for Table2VecW and Table2VecE can be accessed at <https://github.com/iai-group/sigir2019-table2vec>.

<sup>2</sup> <https://code.google.com/archive/p/word2vec/>



**Table 4**

The state-of-the-art post-retrieval query performance predictors used as baselines in our work.  $T$  is a table in corpus  $C$  and  $V$  ( $\varphi(C)$ ) is the vocabulary of the unique tokens in the corpus  $C$ .  $\lambda(t_q)$  reflects the relative weight of the tokens type  $t_q$  and is inversely related to the square root of the number of query tokens of that type.  $\mu$  is the mean score of tables in  $\pi_R^k(q; C)$ .  $\Pr(t_C | \pi_R^k(q; C)) = \sum_{T \in \pi_R^k(q; C)} \Pr(t_C | T) (\prod_{t_q \in \varphi(q)} \Pr(t_q | T))$ .

Baseline Predictors	Formula
Clarity Cronen-Townsend et al. (2002)	$Clarity(\varphi(\pi_R^k(q; C))) = \sum_{t_C \in V(\varphi(C))} \Pr(t_C   \pi_R^k(q; C)) \log \frac{\Pr(t_C   \pi_R^k(q; C))}{\Pr(t_C   C)}$
WIG Zhou and Croft (2007)	$WIG(\varphi(\pi_R^k(q; C))) = \frac{1}{k} \sum_{T \in \pi_R^k(q; C)} \sum_{t_q \in \varphi(q)} \lambda(t_q) \log \frac{\Pr(t_q   T)}{\Pr(t_q   C)}$
NQC Shtok et al. (2012)	$NQC(\varphi(\pi_R^k(q; C))) = \frac{\sqrt{\frac{1}{k} \sum_{T \in \pi_R^k(q; C)} (Score(q, T) - \mu)^2}}{ Score(q, C) }$
SMV Tao and Wu (2014)	$SMV(\varphi(\pi_R^k(q; C))) = \frac{\sqrt{\frac{1}{k} \sum_{T \in \pi_R^k(q; C)} (Score(q, T)   \ln \frac{Score(q, T)}{\mu})}}{Score(q, C)}$

**Table 5**

Performance of the ranking function on the ad hoc table retrieval task in terms of NDCG@10 and @20.

	Method	NDCG@10	NDCG@20
Content-based	STR	0.6293	0.6825
	Single Field	0.4860	0.5473
	Table2VecW	0.6096	0.6505
	Table2VecE	0.5569	0.6161
	LTR	0.5456	0.6031
Feature-based	WikiTable	0.4766	0.5206
	WebTable	0.2992	0.3726

results reported in Tables 6–9 show that while the bag of entities representation is computationally expensive, it results in weaker quality compared to the bag of words representation for most neural predictors and ranking functions, regardless of being content-based or feature-based (except for the WebTable ranking function). The main reason that we find for this is that it may be due to the fact that many of the tables in the test collection do not contain a noticeable number of entities. This was a finding by Bagheri and Al-Obeidat (2020) as well who reported that an entity-based retrieval method is only effective for those tables that consist of a sizeable number of entities; otherwise, a sparse entity representation for the tables, resulting from the lack of a sufficient number of entities, would lead to poor performance. Similarly to Bagheri and Al-Obeidat (2020), we find that the lack of entities to represent tables leads

**Table 6**

The results of various neural predictors with different ranking functions for both bag of words and bag of entities based representations in terms of Pearson Correlation. The bold values show the largest value in each rectangle related to a ranking function and a QPP method. For instance, 0.172 is largest value related to STR and Clarity.

NDCG@20 Pearson		Ranking Function													
		Content-based								Feature-based					
		STR		Single Field		Table2VecW		Table2VecE		WikiTable		WebTable		LTR	
		Word	Entity	Word	Entity	Word	Entity	Word	Entity	Word	Entity	Word	Entity	Word	Entity
Clarity	$P_{Clarity}^{NM}$	0.138	0.114	<b>0.237</b>	0.106	0.154	<b>0.204</b>	0.129	0.168	0.169	<b>0.228</b>	0.159	0.244	0.034	<b>0.167</b>
	$P_{Clarity}^{NAM}$	<b>0.172</b>	0.115	0.2	0.097	0.168	0.188	0.032	<b>0.18</b>	0.025	0.186	0.039	<b>0.25</b>	0.098	0.142
	$P_{Clarity}^{ND}$	0.051	0.068	0.115	0.094	0.059	0.099	0.059	0.05	0.031	0.026	0.079	0.036	0.018	0.069
WIG	$P_{WIG}^{NM}$	0.085	0.087	0.009	0.106	0.009	0.126	0.043	0.119	0.025	0.136	0.049	0.131	0.003	0.053
	$P_{WIG}^{NAM}$	<b>0.387</b>	0.021	<b>0.457</b>	0.101	<b>0.339</b>	0.046	<b>0.389</b>	0.031	<b>0.42</b>	0.034	0.111	<b>0.177</b>	<b>0.341</b>	0.008
	$P_{WIG}^{ND}$	0.032	0.016	0.103	0.056	0.009	0.169	0.013	0.119	0.182	0.042	0.037	0.032	0.06	0.087
NQC	$P_{NQC}^{NM}$	<b>0.177</b>	0.007	0.107	0.172	0.032	0.071	0.017	0.031	0.063	0.059	0.003	0.097	0.114	0.062
	$P_{NQC}^{NAM}$	0.015	0.025	<b>0.173</b>	0.16	0.173	0.056	<b>0.109</b>	0.012	<b>0.285</b>	0.059	0.089	<b>0.116</b>	<b>0.163</b>	0.06
	$P_{NQC}^{ND}$	0.08	0.021	0.08	0.112	<b>0.228</b>	0.173	0.01	0.01	0.006	0.081	0.1	0.075	0.021	0.062
SMV	$P_{SMV}^{NM}$	<b>0.161</b>	0.006	0.115	0.203	0.006	0.055	0.013	0.01	0.025	0.091	0.009	0.088	0.087	0.039
	$P_{SMV}^{NAM}$	0.042	0.029	0.136	<b>0.244</b>	0.175	0.085	<b>0.063</b>	0.012	<b>0.266</b>	0.11	<b>0.131</b>	0.06	<b>0.164</b>	0.011
	$P_{SMV}^{ND}$	0.053	0.02	0.085	0.063	<b>0.306</b>	0.161	0.017	0.007	0.029	0.124	0.079	0.112	0.003	0.02

**Table 7**

The results of various neural predictors with different ranking functions for both bag of words and bag of entities based representations in terms of Pearson Correlation. The bold values show the largest value in each rectangle related to a ranking function and a QPP method.

AP@20 Pearson		Ranking Function													
		Content-based								Feature-based					
		STR		Single Field		Table2VecW		Table2VecE		WikiTable		WebTable		LTR	
		Word	Entity	Word	Entity	Word	Entity	Word	Entity	Word	Entity	Word	Entity	Word	Entity
Clarity	$P_{Clarity}^{NM}$	0.156	0.114	<b>0.242</b>	0.115	0.182	<b>0.257</b>	0.17	<b>0.196</b>	0.157	<b>0.22</b>	0.199	<b>0.278</b>	0.013	<b>0.207</b>
	$P_{Clarity}^{NAM}$	<b>0.186</b>	0.104	0.189	0.094	0.172	0.188	0.067	0.163	0.01	0.196	0.066	0.272	0.049	0.157
	$P_{Clarity}^{ND}$	0.131	0.102	0.152	0.218	0.136	0.031	0.167	0.022	0.036	0.079	0.14	0.078	0.004	0.095
WIG	$P_{WIG}^{NM}$	0.03	0.096	0.066	0.114	0.042	0.153	0.002	0.168	0.008	0.136	0.064	0.162	0.021	0.068
	$P_{WIG}^{NAM}$	<b>0.362</b>	0.012	<b>0.4</b>	0.055	<b>0.32</b>	0.015	<b>0.365</b>	0.036	<b>0.388</b>	0.047	0.097	0.23	<b>0.327</b>	0.017
	$P_{WIG}^{ND}$	0.065	0.049	0.151	0.01	0.018	0.187	0.024	0.152	0.137	0.039	0.073	<b>0.245</b>	0.037	0.116
NQC	$P_{NQC}^{NM}$	<b>0.19</b>	0.003	0.134	0.152	0.005	0.037	0.007	0.04	0.022	0.013	0.056	0.092	0.047	0.035
	$P_{NQC}^{NAM}$	0.113	0.028	<b>0.162</b>	0.128	<b>0.306</b>	0.042	<b>0.192</b>	0.009	<b>0.317</b>	0.033	<b>0.17</b>	0.095	<b>0.277</b>	0.022
	$P_{NQC}^{ND}$	0.103	0.013	0.087	0.085	0.209	0.15	0.007	0.006	0.051	0.039	0.152	0.084	0.009	0.025
SMV	$P_{SMV}^{NM}$	<b>0.171</b>	0.0001	0.143	0.185	0.021	0.019	0.018	0.021	0.07	0.051	0.058	0.089	0.015	0.01
	$P_{SMV}^{NAM}$	0.074	0.039	0.1	<b>0.218</b>	<b>0.288</b>	0.067	<b>0.12</b>	0.014	<b>0.301</b>	0.081	<b>0.221</b>	0.054	<b>0.26</b>	0.021
	$P_{SMV}^{ND}$	0.108	0.005	0.11	0.12	0.274	0.143	0.006	0.019	0.084	0.083	0.134	0.126	0.018	0.008

**Table 8**

The results of various neural predictors with different ranking functions for both bag of words and bag of entities based representations in terms of Kendall Correlation. The bold values show the largest value in each rectangle related to a ranking function and a QPP method.

NDCG@20 Kendall		Ranking Function													
		Content-based								Feature-based					
		STR		Single Field		Table2VecW		Table2VecE		WikiTable		WebTable		LTR	
		Word	Entity	Word	Entity	Word	Entity	Word	Entity	Word	Entity	Word	Entity	Word	Entity
Clarity	$P_{Clarity}^{NM}$	0.094	0.127	0.153	0.018	<b>0.188</b>	0.134	<b>0.177</b>	0.165	<b>0.134</b>	0.111	0.078	0.187	0.084	0.123
	$P_{Clarity}^{NAM}$	0.087	0.117	0.156	0.047	0.139	0.058	0.06	0.064	0.005	0.099	0.127	<b>0.21</b>	0.067	0.097
	$P_{Clarity}^{ND}$	<b>0.144</b>	0.045	<b>0.208</b>	0.017	0.144	0.012	0.133	0.001	0.107	0.003	0.078	0.053	<b>0.158</b>	0.106
WIG	$P_{WIG}^{NM}$	0.109	0.043	0.009	0.006	0.036	0.115	0.016	0.037	0.019	0.0001	0.003	0.119	0.046	0.056
	$P_{WIG}^{NAM}$	<b>0.234</b>	0.048	<b>0.28</b>	0.082	<b>0.26</b>	0.081	<b>0.264</b>	0.043	<b>0.2</b>	0.072	0.019	0.122	<b>0.212</b>	0.042
	$P_{WIG}^{ND}$	0.055	0.0001	0.1	0.004	0.005	0.074	0.024	0.045	0.146	0.064	0.033	<b>0.142</b>	0.03	0.004
NQC	$P_{NQC}^{NM}$	<b>0.081</b>	<b>0.081</b>	0.036	<b>0.207</b>	0.034	0.09	0.008	0.009	0.046	0.034	0.041	0.068	0.077	0.036
	$P_{NQC}^{NAM}$	0.06	0.069	0.137	0.2	<b>0.13</b>	0.067	<b>0.148</b>	0.063	<b>0.153</b>	0.075	0.088	0.081	<b>0.16</b>	0.024
	$P_{NQC}^{ND}$	0.014	0.036	0.004	0.134	0.061	0.104	0.031	0.008	0.037	0.038	0.095	<b>0.1</b>	0.031	0.067
SMV	$P_{SMV}^{NM}$	0.042	0.063	0.029	0.214	0.006	0.077	0.01	0.008	0.028	0.072	0.054	0.067	0.068	0.033
	$P_{SMV}^{NAM}$	0.049	<b>0.071</b>	0.13	<b>0.242</b>	<b>0.136</b>	0.108	<b>0.114</b>	0.037	<b>0.167</b>	0.088	0.112	0.032	<b>0.161</b>	0.014
	$P_{SMV}^{ND}$	0.024	0.041	0.006	0.181	0.1	0.108	0.041	0.006	0.032	0.08	0.079	<b>0.125</b>	0.027	0.026

to poorer performance by our proposed neural predictors when using the bag of entities representation.

We further study the effectiveness of bag of words and bag of entities based representations on two state-of-the-art post-retrieval QPP methods in Table 10, namely Clarity and WIG. We cannot include SMV and NQC in this table as they rely on table relevance scores and do not depend on whether the table representation is based on entities or keywords. As such, they are not applicable to RQ1. As seen in this table, in the majority of the cases, the bag of words based representation shows a superior performance compared to the bag of entities based representation on both NDCG and AP for both Kendall and Pearson correlations. We again note consistent observations across NDCG and AP for both correlation measures. Therefore, we conclude that in the context of ad hoc table retrieval, and for both when our neural predictors are present and otherwise, using a bag of words representation for both tables and queries is sufficient for performing the QPP task.

**Finding 1.** In the context of ad hoc table retrieval, a bag of words representation of tables and queries is the more efficient form of representation compared to the bag of entities representation as it provides stronger performance for the QPP task.

**Table 9**

The results of various neural predictors with different ranking functions for both bag of words and bag of entities based representations in terms of Kendall Correlation. The bold values show the largest value in each rectangle related to a ranking function and a QPP method.

AP@20 Kendall		Ranking Function													
		Content-based								Feature-based					
		STR		Single Field		Table2VecW		Table2VecE		WikiTable		WebTable		LTR	
		Word	Entity	Word	Entity	Word	Entity	Word	Entity	Word	Entity	Word	Entity	Word	Entity
Clarity	$P_{Clarity}^{NM}$	0.087	<b>0.144</b>	0.166	0.049	<b>0.206</b>	0.162	<b>0.195</b>	0.174	0.104	0.17	0.101	0.17	0.087	0.135
	$P_{Clarity}^{NAM}$	0.079	0.048	0.151	0.023	0.126	0.087	0.053	0.025	0.007	<b>0.184</b>	0.115	<b>0.184</b>	0.078	0.056
	$P_{Clarity}^{ND}$	0.128	0.086	<b>0.202</b>	0.001	0.146	0.068	0.166	0.058	0.173	0.038	0.117	0.038	<b>0.187</b>	0.099
WIG	$P_{WIG}^{NM}$	0.016	0.056	0.041	0.0001	0.028	0.159	0.001	0.086	0.026	0.119	0.007	0.119	0.018	0.06
	$P_{WIG}^{NAM}$	<b>0.225</b>	0.004	<b>0.232</b>	0.059	<b>0.252</b>	0.063	<b>0.251</b>	0.005	<b>0.258</b>	0.133	0.033	0.133	<b>0.235</b>	0.028
	$P_{WIG}^{ND}$	0.048	0.03	0.117	0.007	0.019	0.099	0.022	0.058	0.091	0.176	0.065	<b>0.176</b>	0.019	0.037
NQC	$P_{NQC}^{NM}$	<b>0.143</b>	0.032	0.093	<b>0.166</b>	0.013	0.034	0.002	0.003	0.02	0.082	0.038	0.082	0.049	0.02
	$P_{NQC}^{NAM}$	0.127	0.041	0.095	0.143	<b>0.224</b>	0.016	<b>0.145</b>	0.05	<b>0.193</b>	0.083	<b>0.124</b>	0.083	<b>0.205</b>	0.007
	$P_{NQC}^{ND}$	0.063	0.002	0.028	0.108	0.083	0.082	0.038	0.04	0.054	0.079	0.094	0.079	0.018	0.034
SMV	$P_{SMV}^{NM}$	0.112	0.027	0.094	0.175	0.032	0.026	0.011	0.011	0.051	0.081	0.052	0.081	0.034	0.018
	$P_{SMV}^{NAM}$	<b>0.119</b>	0.06	0.074	<b>0.197</b>	<b>0.225</b>	0.057	<b>0.116</b>	0.047	<b>0.209</b>	0.042	<b>0.149</b>	0.042	<b>0.199</b>	0.034
	$P_{SMV}^{ND}$	0.083	0.03	0.05	0.124	0.103	0.098	0.04	0.017	0.058	0.102	0.078	0.102	0.029	0.002

**Table 10**

The results of baseline predictors with different ranking functions for both bag of words and bag of entities based representations in terms of Pearson and Kendall Correlation on both NDCG and AP. Bold values show the larger value based on the pair of ranking function and QPP method for word vs entity.

		Ranking Function														
		Content-based								Feature-based						
		STR		Single Field		Table2VecW		Table2VecE		WikiTable		WebTable		LTR		
		Word	Entity	Word	Entity	Word	Entity	Word	Entity	Word	Entity	Word	Entity	Word	Entity	
Pearson	NDCG@20	Clarity	<b>0.174</b>	0.006	<b>0.14</b>	0.063	0.051	<b>0.072</b>	<b>0.285</b>	0.037	0.004	<b>0.108</b>	<b>0.17</b>	0.003	<b>0.293</b>	0.157
		WIG	<b>0.209</b>	0.004	<b>0.275</b>	0.0001	<b>0.149</b>	0.075	<b>0.283</b>	0.08	<b>0.205</b>	0.053	0.2	<b>0.224</b>	<b>0.272</b>	0.119
		AP@20	Clarity	<b>0.202</b>	0.027	<b>0.157</b>	0.06	0.029	<b>0.098</b>	<b>0.244</b>	0.06	0.082	<b>0.141</b>	<b>0.191</b>	0.022	<b>0.263</b>
Kendall	NDCG@20	WIG	<b>0.231</b>	0.032	<b>0.353</b>	0.071	<b>0.197</b>	0.143	<b>0.283</b>	0.089	<b>0.232</b>	0.028	0.194	<b>0.211</b>	<b>0.291</b>	0.137
		Clarity	<b>0.079</b>	0.055	<b>0.041</b>	0.01	0.055	<b>0.094</b>	<b>0.119</b>	0.045	0.003	<b>0.074</b>	<b>0.094</b>	0.05	<b>0.166</b>	0.031
		WIG	<b>0.098</b>	0.056	<b>0.176</b>	0.057	<b>0.089</b>	0.015	<b>0.171</b>	0.024	<b>0.133</b>	0.028	<b>0.106</b>	0.067	<b>0.16</b>	0.04
AP@20	Clarity	<b>0.08</b>	0.015	<b>0.02</b>	0.011	0.062	<b>0.063</b>	<b>0.131</b>	0.033	0.011	<b>0.047</b>	<b>0.106</b>	0.069	<b>0.153</b>	0.013	
	WIG	<b>0.148</b>	0.047	<b>0.239</b>	0.002	<b>0.158</b>	0.043	<b>0.176</b>	0.011	<b>0.169</b>	0.02	<b>0.107</b>	0.051	<b>0.196</b>	0.027	

#### 4.5.2. RQ2. Performance of the neural predictors

In the second research question, we aim at exploring the impact of the proposed neural features on query performance prediction. We expect that the impact of the neural predictors would be different for the various ranking functions especially because the ranking functions consider a different set of features when determining relevance and performing ranking, i.e., content-based ranking functions compared to feature-based ranking functions. For instance, a content-based ranking function such as STR (Zhang & Balog, 2018a) measures the similarity between tables and queries based on their semantic associations, which means that such a method would be able to retrieve relevant yet lexically non-overlapping tables for an input query. In contrast, a feature-based ranking method such as LTR (Zhang & Balog, 2018a) is based on a set of structural features and hence would be sensitive to how information is spread within the table structure. As such, given the fact that the performance of the ranking functions are different, we expect that the performance of the QPP methods would also be different.

In our experiments, we compare the performance of the four baseline QPP methods and our proposed neural predictors over all ranking functions. The results are reported in Table 11. In this table, the best performing QPP method for each of the ranking functions has been shown in Bold. Furthermore, given the findings from RQ1, we report the results of the QPP methods based on the bag of words representation noting the weaker performance of the bag of entities representation for this task. We also note that we resort to reporting NDCG using Kendall correlation as we already showed in Tables 6–9 that AP and NDCG values with both Kendall and Pearson correlations have consistent behavior.

Our first observation is that our proposed  $P_{WIG}^{NAM}$  predictor has the best performance for content-based ranking functions, i.e., the

**Table 11**

The quality of different query performance predictors for each ranking function in terms of Kendall's correlation. The best value for each ranking function is specified in boldface.

NDCG@20 Kendall		Ranking Function						
		Content-based				Feature-based		
		STR	Single Field	Table2VecW	Table2VecE	WikiTable	WebTable	LTR
Clarity	<i>Clarity</i>	0.079	0.041	0.055	0.119	0.003	0.094	0.166
	$P_{Clarity}^{NM}$	0.094	0.153	0.188	0.177	0.134	0.078	0.084
	$P_{Clarity}^{NAM}$	0.087	0.156	0.139	0.06	0.005	0.127	0.067
	$P_{Clarity}^{ND}$	0.144	0.208	0.144	0.133	0.107	0.078	0.158
WIG	<i>WIG</i>	0.098	0.176	0.089	0.171	0.133	0.106	0.16
	$P_{WIG}^{NM}$	0.109	0.009	0.036	0.016	0.019	0.003	0.046
	$P_{WIG}^{NAM}$	<b>0.234</b>	<b>0.28</b>	<b>0.26</b>	<b>0.264</b>	0.2	0.019	0.212
	$P_{WIG}^{ND}$	0.055	0.1	0.005	0.024	0.146	0.033	0.03
NQC	<i>NQC</i>	0.076	0.041	0.229	0.252	<b>0.318</b>	0.165	<b>0.338</b>
	$P_{NQC}^{NM}$	0.081	0.036	0.034	0.008	0.046	0.041	0.077
	$P_{NQC}^{NAM}$	0.06	0.137	0.13	0.148	0.153	0.088	0.16
	$P_{NQC}^{ND}$	0.014	0.004	0.061	0.031	0.037	0.095	0.031
SMV	<i>SMV</i>	0.084	0.055	0.219	0.239	0.303	<b>0.17</b>	0.308
	$P_{SMV}^{NM}$	0.042	0.029	0.006	0.01	0.028	0.054	0.068
	$P_{SMV}^{NAM}$	0.049	0.13	0.136	0.114	0.167	0.112	0.161
	$P_{SMV}^{ND}$	0.024	0.006	0.1	0.041	0.032	0.079	0.027

STR, Single Field, Table2VecW and Table2VecE ranking functions. In contrast, in the feature-based ranking functions, the baseline predictors without the integration of our proposed neural features are more successful. This observation confirms the relationship between the query performance predictors and the characteristics of the ranking functions for bag of word based representation in the table retrieval task. The impact of our proposed neural predictors are most visibly observed when the ranking functions consider the similarity between the *contents* of tables and the queries. Given our proposed neural predictors employ neural embeddings to capture the semantic association between keyword vectors and the distances between the embedded keyword vectors, they are able to model and capture semantic relationships, which are primarily content-based in nature. This reflects itself in developing stronger QPP methods for content-based ranking functions that are based on query and table semantic associations or similarity through language models.

**Finding 2.** For the purpose of predicting the performance of ranking functions that are primarily based on the *content value* of the tables (content-based ranking functions), while the neural predictors do not always improve the performance of the baselines, the integration of one of our proposed neural predictors with the WIG metric leads to consistent and noticeable performance improvement over all baselines.

For the feature-based ranking functions such as LTR, WikiTable and WebTable that are primarily based on structural features of the table, the base QPP methods without the integration of our proposed neural features show better performance. The structural features of the table include features, such as the number of rows, columns, and empty cells, and the query word frequency in specific parts of the table, such as the leftmost column, or second-to-left column, that cannot be effectively captured by our proposed neural predictors. As such, the baseline QPP methods show better performance compared to our proposed neural predictors.

**Finding 3.** Traditional QPP methods that do not account for the semantic association between the table and query spaces are more effective for determining query difficulty of feature-based ranking functions that consider table-specific structural features.

As mentioned earlier, one of the strongest state of the art post-retrieval query performance prediction methods in the literature for document retrieval is the work by Zamani et al. (2018), which uses weak supervision to rank queries. Given this method requires training and is not fully unsupervised, our proposed predictors cannot be integrated into its formulation. However, we compare our work against this baseline as well in Table 12. For easier comparison to the performance of our own work, we have included  $P_{WIG}^{NAM}$  in this table; however, comparison against other variations of our work can be done according to performance values reported in Tables 6–9. As shown in Table 12, the only case where the weakly supervised method by Zamani et al. shows better performance compared to  $P_{WIG}^{NAM}$  is on the STR method (and only on Pearson correlation). However, for all other ranking functions, it shows a consistently weaker performance. We also note that while all the baseline QPP methods as well as our proposed neural predictors are fully unsupervised, the work proposed in Zamani et al. (2018) requires a training process but even so, while shown to be effective for QPP in document retrieval, it is not as effective for the context of table retrieval.

#### 4.5.3. RQ3. Synergistic impact of the neural predictors

The main objective of the third research question is to investigate whether the interpolation of neural predictors with base QPP

**Table 12**

Comparative performance of the work by [Zamani et al. \(2018\)](#) with  $P_{WIG}^{NAM}$ . Please note P stands for Pearson and K stands for Kendall correlation. The only places where the baseline is better than our proposed approach is marked with boldface.

		$P_{WIG}^{NAM}$				Zamani et al			
		NDCG@20		AP@20		NDCG@20		AP@20	
		P	K	P	K	P	K	P	K
Content-based	STR	0.387	0.234	0.362	0.225	0.322	<b>0.295</b>	0.338	<b>0.271</b>
	Single Field	0.457	0.280	0.400	0.232	0.289	0.209	0.307	0.228
	Table2VecW	0.339	0.260	0.320	0.252	0.104	0.089	0.011	0.001
	Table2VecE	0.389	0.264	0.365	0.251	0.037	0.009	0.087	0.041
Feature-based	WikiTable	0.420	0.200	0.388	0.258	0.060	0.086	0.010	0.016
	WebTable	0.111	0.019	0.097	0.033	0.124	0.085	0.113	0.073
	LTR	0.341	0.212	0.327	0.235	0.110	0.116	0.150	0.138

methods will enhance the quality of the state-of-the-art predictors or not. Our hypothesis is that the neural and baseline predictors can produce overlapping results while in many cases a subset of their correct results is distinct and non-overlapping. For this reason, we believe that the interpolation of these predictors can result in synergy between the correctly identified results of each predictor and hence lead to improved quality. For this purpose, we employ the CombSum approach [Shaw and Fox \(1994\)](#) and form three variations, referred to as (1) *Best Baseline + Best Neural*, (2) *Best of each Predictor - Best Neural + Baseline* and (3) *Best Word + Best Entity*. In all of three variations, first, the scores generated by each predictor is normalized and then, linear interpolation is performed.

**Best Baseline + Best Neural.** In this variation of our interpolations, for each ranking function, the best neural predictor and the best baseline QPP method are selected to be interpolated through CombSum. The quality of the final CombSum predictor in terms of Kendall correlation and the amount of achieved improvement over the baseline ( $\Delta_{Baseline}$ ) and neural predictors ( $\Delta_{Neural}$ ) are reported in [Table 13](#). The first observation based on  $\Delta_{Baseline}$  is that the best neural predictors have been able to improve the best baselines that they have been interpolated with in content-based ranking functions. Another observation here is with regards to the performance of the interpolated QPP method for content-based ranking functions. The observed performance is consistent for these ranking functions where the best neural predictor is  $P_{WIG}^{NAM}$  for all content-based methods. Furthermore, the interpolation has been able to consistently improve the best baseline QPP method.

This indicates that:

1. the best non-neural baseline QPP methods cannot provide synergistic support for our proposed neural predictors when applied to methods that are solely based on table and query content (content-based ranking);
2. the proposed neural predictors in this paper are able to provide additional complementary performance over baseline QPP methods when dealing with feature-based ranking functions.

**Finding 4.** The linear interpolation of the best baseline QPP method with our best neural predictor leads to improved performance over the best baseline QPP method for content-based ranking functions. In contrast, improvements are observed over the best neural predictors when the ranking functions are feature-based in nature (LTR, WebTable, WikiTable).

An important point to mention here based on  $\Delta_{Baseline}$  is that the degree of improvement over the baselines in content-based ranking function is more than other ranking functions and we conclude that the neural predictor is more effective in increasing the quality of the baselines when dealing with content-based ranking functions. This confirms our findings with regards to the relationship between the ranking functions and query performance predictors concluded in RQ2.

**Best of each Predictor - Best Neural + Baseline.** In this variation of our interpolations, we study the performance of interpolation for each predictor individually. For this purpose, in each ranking function, the best neural predictor is selected to be interpolated with the baseline of the same predictor through the CombSum approach. [Table 14](#) shows the interpolation performance and the improvement percentages over the baseline and neural predictor for Clarity. As shown in [Table 14](#), interpolating the best neural

**Table 13**

The evaluation results for the interpolation of the best neural predictor with the best baseline for each ranking function individually.

NDCG@20 Kendall	Ranking Function	Best Baseline	Best Neural Predictors	CumbSum	$\Delta_{Baseline}$	$\Delta_{Neural}$
Content-based	STR	WIG	$P_{WIG}^{NAM}$	0.169	7.1%	-6.5%
	Single Field	WIG	$P_{WIG}^{NAM}$	0.269	9.3%	-1.1%
	Table2VecW	NQC	$P_{WIG}^{NAM}$	0.257	2.8%	-0.3%
	Table2VecE	NQC	$P_{WIG}^{NAM}$	0.298	4.6%	3.4%
Feature-based	WikiTable	NQC	$P_{WIG}^{NAM}$	0.311	-0.7%	11.1%
	WebTable	SMV	$P_{Clarity}^{NAM}$	0.1273	-4.27%	0.03%
	LTR	NQC	$P_{WIG}^{NAM}$	0.311	-2.7%	9.9%



**Table 14**  
Evaluation results for clarity.

$NDCG@20$ Kendall	Ranking Function	Best Neural Predictors	CumbSum	$\Delta_{Baseline}$	$\Delta_{Neural}$
Content-based	STR	$P_{Clarity}^{ND}$	0.245	16.6%	10.1%
	Single Field	$P_{Clarity}^{ND}$	0.197	15.6%	-1.1%
	Table2VecW	$P_{Clarity}^{NM}$	0.147	9.2%	-4.1%
	Table2VecE	$P_{Clarity}^{NM}$	0.136	1.7%	-4.1%
Feature-based	WikiTable	$P_{Clarity}^{NM}$	0.057	5.4%	-7.7%
	WebTable	$P_{Clarity}^{NAM}$	0.058	-3.6%	-6.9%
	LTR	$P_{Clarity}^{ND}$	0.261	9.5%	10.3%

predictor with Clarity will not lead to improved performance over the Clarity method itself in WebTable. The results from  $\Delta_{Neural}$  show that the quality of the best neural predictors are better than the interpolated methods (except on the STR and LTR ranking functions). Therefore, we find that the interpolation of our proposed neural predictors with Clarity can potentially lead to improvements over Clarity (except WebTable) but not over the neural predictors; therefore, showing non-synergistic overlapping behavior between them.

We report the quality of CombSum approach and the improvements over the best neural predictors and baseline for WIG in Table 15. The main observation here is that the neural aggregated matching feature when applied to WIG is the best neural predictor for all ranking functions (except WebTable). Another observation is that the interpolation of  $P_{WIG}^{NAM}$  with WIG improves the performance of WIG in all ranking functions except for WebTable. Also, the results show that when  $P_{WIG}^{NAM}$  is interpolated with WIG, the estimation quality of  $P_{WIG}^{NAM}$  is increased in feature-based ranking functions while this increase cannot be observed for content-based ranking functions.

NQC and SMV are score-oriented models and consider the variance of the returned table scores for measuring the performance of a query as described earlier in Table 4. Tables 16 and 17 show the quality of CombSum after interpolation with the best neural predictor with NQC and SMV, respectively. For both NQC and SMV, the overall performance of CombSum is not comparable to our proposed neural baselines nor to other CombSums such as with WIG. As such, we conclude that the interpolation of the best neural predictors with NQC and SMV is not an effective strategy.

**Finding 5.** The linear interpolation of each base QPP method with its best neural predictor does not necessarily lead to improved performance over the base QPP method. However, the WIG method shows to be the strongest baseline method over all ranking functions, which is significantly improved as a result of interpolation with its best neural predictor (except WebTable).

**Best Word + Best Entity.** In the last variation of our interpolations, we investigate, while we already found that bag of words based representation is superior to bag of entities representation, whether the interpolation of the two representations leads to increased performance for QPP. To this end, the results of the best neural predictor for the bag of words and bag of entities representations are selected and interpolated through CombSum for each ranking function separately. Table 18 reports the degree of improvement over the best bag of words and bag of entities representations.

We observe that in most ranking functions, there are no consistent improvements observed over the bag of words based representation or bag of entities based representation. We find there were no performance changes in LTR and WikiTable, which is due to the fact that the best predictors in the bag of words and bag of entities based representations are the same and therefore no improvements could be observed. Overall, when considering improvements obtained through the interpolation of the best predictors from bag of words and bag of entities representations in comparison with the best predictor from the bag of words representation, described by  $\Delta_{Word}$  in Table 18, we conclude that the bag of entities representation does not contribute to meaningful improvements for the QPP task. This reinforces our findings in RQ1.

**Finding 6.** The best predictors for each of the ranking functions obtained through bag of words and bag of entities representations exhibit overlapping behavior and hence their interpolation does not lead to meaningful improvements over the best predictor from the bag of words representation.

**Table 15**  
Evaluation results for WIG.

$NDCG@20$ Kendall	Ranking Function	Best Neural Predictors	CumbSum	$\Delta_{Baseline}$	$\Delta_{Neural}$
Content-based	STR	$P_{WIG}^{NAM}$	0.169	7.1%	-6.5%
	Single Field	$P_{WIG}^{NAM}$	0.269	9.3%	-1.1%
	Table2VecW	$P_{WIG}^{NAM}$	0.198	10.9%	-6.2%
	Table2VecE	$P_{WIG}^{NAM}$	0.248	7.7%	-1.6%
Feature-based	WikiTable	$P_{WIG}^{NAM}$	0.209	7.6%	0.9%
	WebTable	$P_{WIG}^{ND}$	0.049	-5.7%	1.6%
	LTR	$P_{WIG}^{NAM}$	0.229	6.9%	1.7%

**Table 16**  
Evaluation results for NQC.

NDCG@20 Kendall	Ranking Function	Best Neural Predictors	CumbSum	$\Delta_{Baseline}$	$\Delta_{Neural}$
Content-based	STR	$P_{NQC}^{NM}$	0.08	0.4%	-0.1%
	Single Field	$P_{NQC}^{NAM}$	0.085	4.4%	-5.2%
	Table2VecW	$P_{NQC}^{NAM}$	0.017	-21.2%	-11.3%
	Table2VecE	$P_{NQC}^{NAM}$	0.015	-23.7%	-13.3%
Feature-based	WikiTable	$P_{NQC}^{NAM}$	0.11	-20.8%	-4.3%
	WebTable	$P_{NQC}^{ND}$	0.107	-5.8%	1.2%
	LTR	$P_{NQC}^{NAM}$	0.048	-29%	-11.2%

**Table 17**  
Evaluation results for SMV.

NDCG@20 Kendall	Ranking Function	Best Neural Predictors	CumbSum	$\Delta_{Baseline}$	$\Delta_{Neural}$
Content-based	STR	$P_{SMV}^{NAM}$	0.041	-4.3%	-0.8%
	Single Field	$P_{SMV}^{NAM}$	0.116	6.1%	-1.4%
	Table2VecW	$P_{SMV}^{NAM}$	0.02	-19.9%	-11.6%
	Table2VecE	$P_{SMV}^{NAM}$	0.009	-23%	-10.5%
Feature-based	WikiTable	$P_{SMV}^{NAM}$	0.086	-21.7%	-8.1%
	WebTable	$P_{SMV}^{NAM}$	0.186	1.6%	7.4%
	LTR	$P_{SMV}^{NAM}$	0.014	-29.4%	-14.7%

**Table 18**  
Evaluation results for the interpolation of the best predictors in the both word and entity representations with each other.

NDCG@20 Kendall	Ranking Function	Best Predictor in Word based Representation	Best Predictor in Entity based Representation	CumbSum	$\Delta_{Word}$	$\Delta_{Entity}$
Content-based	STR	$P_{WIG}^{NAM}$	$P_{Clarity}^{NM}$	0.23	-0.4%	10.3%
	Single Field	$P_{WIG}^{NAM}$	$P_{SMV}^{NAM}$	0.162	-11.8%	-4.5%
	Table2VecW	$P_{WIG}^{NAM}$	NQC	0.257	-0.3%	2.8%
	Table2VecE	$P_{WIG}^{NAM}$	NQC	0.298	3.4%	4.6%
Feature-based	WikiTable	NQC	NQC	0.318	-	-
	WebTable	SMV	$P_{Clarity}^{NAM}$	0.112	-5.8%	-9.8%
	LTR	NQC	NQC	0.338	-	-

We would like to report further on our findings by reviewing our observations across the three types of interpolations reported in this section. We find that the neural aggregated matching (NAM) feature achieves noticeable performance improvements when interpolated with baseline QPP methods. Based on the results of the first variation of our interpolations shown in Table 13, we can observe that the best neural predictor is built using NAM in seven of the ranking functions. This observation is confirmed based on the second variation of our interpolation that in most cases, NAM has been selected as the best neural predictor for the different base QPP methods. Finally, the results of the last variation show that NAM is the variation of our proposed neural predictors to appear the most in Table 18.

**Finding 7.** From amongst the three variations of our proposed neural predictors, the Neural Aggregated Matching (NAM) feature, which considers the maximum similarity obtained from the neural representations of the tokens in the table and query spaces, has the best overall performance.

Furthermore, as neural embeddings play a central role in our proposed neural predictors, it is important to also understand whether the choice of neural embedding representations impact the performance of our neural predictors. There have recently been several variations of embeddings such as BERT Devlin, Chang, Lee, and Toutanova (2019) and ELMo Peters et al. (2018), that have shown impressive performance for tasks such as question answering and language understanding. For the sake of comparison, we compare the performance of the best interpolated variation of our neural predictor, i.e.,  $P_{WIG}^{NAM}$ , for cases when Word2vec-based Mikolov et al. (2013b) Google News word embeddings are adopted to when embeddings from BERT and ELMo are used. The results are shown in Table 19.

As seen in the table, the obtained results do not show one particular type of embedding outperforming the other by a significant

**Table 19**

Impact of embedding type on the performance of  $P_{WIG}^{NAM}$  for each ranking function in terms of Kendall's correlation. Bold values show the largest value per ranking function.

NDCG@20 Kendall	Ranking Function						
	Content-based				Feature-based		
	STR	Single Field	Table2VecW	Table2VecE	WikiTable	WebTable	LTR
Word2vec	0.234	0.28	<b>0.26</b>	0.264	<b>0.2</b>	0.019	0.212
BERT	<b>0.268</b>	0.262	0.208	<b>0.289</b>	0.19	<b>0.388</b>	0.23
ELMo	0.246	<b>0.296</b>	0.25	0.246	0.157	0.231	<b>0.25</b>

margin. While the Word2vec embeddings show favorable performance on the Table2VecW and WikiTable, it is BERT that shows a better performance on STR, Table2VecE, and WebTable. Also, ELMo shows the best performance on Single Field and LTR ranking functions. We note that except for WebTable where the performance of BERT and ELMo is substantially better than word2vec, the results on other ranking functions are comparable.

**Finding 8.** The performance of the best interpolated variation of our neural predictor when customized based on either Word2vec, BERT or ELMo neural embeddings does not seem to depend on the type of the ranking function and in principal one embedding type does not outperform the other (except by BERT on WebTable).

#### 4.5.4. RQ4. Impact of table characteristics on the neural predictors

The objective the last research question is to understand whether the consideration of table-specific characteristics such as table rows, columns, and cells have any impact on the quality of query performance prediction. In the previous research questions, we have considered each table to be a single document and hence all the content in the table cells were concatenated for the sake of query performance prediction. In this research question, we investigate whether the consideration of each of these table elements leads to any impact on QPP performance.

In order to consider table elements for performing QPP, we develop embedding representations for each table element, i.e., cell, column or row, by taking the average of the embeddings of the content observed in that element. For instance, the representation for a table cell is the average of the embedding representations of the content in that cell. We adopt these table element representations to compute our proposed neural predictors by averaging over table element embeddings. Without loss of generality, we report our findings on the best interpolated variation of our neural predictor, i.e.,  $P_{WIG}^{NAM}$  in Table 20. As seen in the table, the last row is the situation where the whole table is considered to be a single document. We find that the consideration of the whole table as a single document produces the best performance for the QPP task on all ranking functions except for WebTable. This seems to be inline with the findings reported by Deng et al. (2019) who found that the consideration of a table as a document leads the improved retrieval performance for the ad hoc table retrieval task as well.

**Finding 9.** While table-specific characteristics such as rows, columns and cells provide richer content presentation and better interpretability, they do not necessarily lead to improved QPP performance. The consideration of the whole table as a document shows the best QPP performance over all ranking functions (except WebTable).

**Table 20**

The impact of table characteristics on the performance of  $P_{WIG}^{NAM}$  for each ranking function in terms of Kendall's correlation. Bold values show the largest value per ranking function.

NDCG@20 Kendall	Ranking Function						
	Content-based				Feature-based		
	STR	Single Field	Table2VecW	Table2VecE	WikiTable	WebTable	LTR
Row-based	0.188	0.194	0.158	0.114	0.028	<b>0.118</b>	0.095
Column-based	0.183	0.202	0.165	0.092	0.07	0.081	0.108
Cell-based	0.157	0.165	0.169	0.097	0.021	0.05	0.069
Unstructured	<b>0.234</b>	<b>0.28</b>	<b>0.26</b>	<b>0.264</b>	<b>0.2</b>	0.019	<b>0.212</b>

## 5. Concluding remarks

The main objective of our work in this paper has been to explore how neural embedding techniques can be exploited to predict query performance in the ad hoc table retrieval task. Our motivation is that neural embedding techniques enable us to estimate the performance of a table ranking function based on the semantic content of the query and the retrieved table list. Neural embedding techniques assign a low-dimensional vector to each token in a new space such that the properties of, and the relationships between, the tokens are preserved. In this space, the vectors of tokens that are semantically or syntactically similar tend to be close.

We proposed a set of neural predictors that represent the content of tables and queries as bag of tokens, which are then systematically integrated into existing state-of-the-art post-retrieval query-performance predictors. The proposed neural features employ neural embeddings to capture the semantic relationship between the representation of tables and queries based on their relationship in the embedding space.

We evaluated the effectiveness of the neural predictor through a complete set of experiments on a test collection that consists of 1.6M WikiTables and 60 queries. We compared the ranking of queries by the neural predictor and actual performance measured according to their true NDCG using Pearson's correlation to assess the neural predictor quality. We summarize our key findings from the results of the experiments as follows:

- A bag of words representation is the more suitable representation for tables and queries in the context of QPP for ad hoc table retrieval. A bag of entity representation does not show equally effective performance especially because of the sparsity of table entity representations;
- Our proposed neural predictors are effective in improving the performance of existing QPP methods when predicting query difficulty for content-based ranking function. The proposed predictors do not contribute to better QPP performance for feature-based ranking functions;
- From among the proposed neural features, NAM, which captures the maximum similarity between the query and table tokens in the embedding space, leads to a more accurate estimation of query performance compared to other proposed neural features.

Finally, we would like to highlight that as this work is among the very first attempts to perform query performance prediction for ad hoc table retrieval, the prediction performances would not be considered suitable to be used in production. Our work shows that while the consideration of neural predictors do provide better performance prediction results compared to existing state of the art QPP methods, the absolute correlation values are still low (for both our approach and existing QPP methods in the literature). As such, we believe there is optimistically room for much better work to build on the foundations of our work.

## CRedit authorship contribution statement

**Maryam Khodabakhsh:** Conceptualization, Software, Validation, Investigation, Formal analysis, Writing - original draft.  
**Ebrahim Bagheri:** Conceptualization, Methodology, Investigation, Writing - review & editing.

## References

- Akbik, A., Blythe, D., & Vollgraf, R. (2018). *Contextual string embeddings for sequence labeling*. COLING 2018, 27th international conference on computational linguistics, 1638–1649.
- Arabzadeh, N., Zarrinkalam, F., Jovanovic, J., Al-Obeidat, F., & Bagheri, E. (2020). Neural embedding-based specificity metrics for pre-retrieval query performance prediction. *Information Processing & Management*, 57(4), 102248. <https://doi.org/10.1016/j.ipm.2020.102248>.
- Aslam, J. A., & Pavlu, V. (2007). *Query hardness estimation using Jensen-Shannon divergence among multiple scoring functions*. European conference on information retrieval. Springer, 198–209.
- Bagheri, E., & Al-Obeidat, F. (2020). *A latent model for ad hoc table retrieval*. Advances in information retrieval. Springer International Publishing, 86–93.
- Bagheri, E., Ensan, F., & Al-Obeidat, F. (2018). Neural word and entity embeddings for ad hoc retrieval. *Information Processing & Management*, 54(4), 657–673.
- Bhagavatula, C. S., Noraset, T., & Downey, D. (2013). *Methods for exploring and mining tables on wikipedia*. Proceedings of the ACM SIGKDD workshop on interactive data exploration and analytics, 18–26.
- Bhagavatula, C. S., Noraset, T., & Downey, D. (2015). *Tabel: entity linking in web tables*. International semantic web conference. Springer, 425–441.
- Breiman, L. (2001). Random forests. *Machine learning*, 45(1), 5–32.
- Cafarella, M. J., Halevy, A., & Khoussainova, N. (2009). Data integration for the relational web. *Proceedings of the VLDB Endowment*, 2(1), 1090–1101.
- Cafarella, M. J., Halevy, A., Wang, D. Z., Wu, E., & Zhang, Y. (2008). Webtables: Exploring the power of tables on the web. *Proceedings of the VLDB Endowment*, 1(1), 538–549.
- Capannini, G., Lucchese, C., Nardini, F. M., Orlando, S., Perego, R., & Tonellotto, N. (2016). Quality versus efficiency in document scoring with learning-to-rank models. *Information Processing & Management*, 52(6), 1161–1177.
- Carmel, D., & Kurland, O. (2012). *Query performance prediction for IR*. Proceedings of the 35th international ACM SIGIR conference on research and development in information retrieval, 1196–1197.
- Carmel, D., & Yom-Tov, E. (2010). Estimating the query difficulty for information retrieval. *Synthesis Lectures on Information Concepts, Retrieval, and Services*, 2(1), 1–89.
- Cronen-Townsend, S., Zhou, Y., & Croft, W. B. (2002). *Predicting query performance*. Proceedings of the 25th annual international ACM SIGIR conference on research and development in information retrieval, 299–306.
- Cummins, R., Jose, J., & O'Riordan, C. (Jose, O'Riordan, 2011a). *Improved query performance prediction using standard deviation*. Proceedings of the 34th international ACM SIGIR conference on research and development in information retrieval, 1089–1090.
- Cummins, R., Lalmas, M., O'Riordan, C., & Jose, J. M. (Lalmas, O'Riordan, Jose, 2011b). *Navigating the user query space*. International symposium on string processing and information retrieval. Springer, 380–385.
- Das Sarma, A., Fang, L., Gupta, N., Halevy, A., Lee, H., Wu, F., ... Yu, C. (2012). *Finding related tables*. Proceedings of the 2012 ACM SIGMOD international conference on management of data, 817–828.
- Déjean, S., Ionescu, R. T., Mothe, J., & Ullah, M. Z. (2020). *Forward and backward feature selection for query performance prediction*. Proceedings of the 35th annual ACM

- symposium on applied computing690–697.
- Deng, L., Zhang, S., & Balog, K. (2019). *Table2vec: Neural word and entity embeddings for table population and retrieval*. *Proceedings of the 42nd international ACM SIGIR conference on research and development in information retrieval*SIGIR19New York, NY, USA: Association for Computing Machinery1029–1032. <https://doi.org/10.1145/3331184.3331333>.
- Devlin, J., Chang, M., Lee, K., & Toutanova, K. (2019). BERT: pre-training of deep bidirectional transformers for language understanding. In J. Burstein, C. Doran, & T. Solorio (Eds.). *Proceedings of the 2019 conference of the north American chapter of the association for computational linguistics: Human language technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2–7, 2019, volume 1 (long and short papers)* (pp. 4171–4186). Association for Computational Linguistics. <https://doi.org/10.18653/v1/n19-1423>.
- Diaz, F. (2016). *Pseudo-query reformulation*. *European conference on information retrieval*. Springer521–532.
- Ensan, F., & Bagheri, E. (2017). *Document retrieval model through semantic linking*. *Proceedings of the tenth ACM international conference on web search and data mining, WSDM 2017, Cambridge, United Kingdom, February 6–10, 2017*181–190. <https://doi.org/10.1145/3018661.3018692>.
- Furnas, G. W., Deerwester, S., Durnais, S. T., Landauer, T. K., Harshman, R. A., Streeter, L. A., & Lochbaum, K. E. (2017). *Information retrieval using a singular value decomposition model of latent semantic structure*. *ACM SIGIR forum*51. *ACM SIGIR forum* ACM New York, NY, USA90–105.
- Galke, L., Saleh, A., & Scherp, A. (2017). *Word embeddings for practical information retrieval*. *Informatik 2017. Gesellschaft für Informatik, Bonn*2155–2167.
- Gentile, A. L., Ristoski, P., Eckel, S., Ritze, D., & Paulheim, H. (2017). Entity matching on web tables: a table embeddings approach for blocking. In V. Markl, S. Orlando, B. Mitschang, P. Andritsos, K. Sattler, & S. Breß (Eds.). *Proceedings of the 20th international conference on extending database technology, EDBT 2017, Venice, Italy, March 21–24, 2017* (pp. 510–513). OpenProceedings.org. <https://doi.org/10.5441/002/edbt.2017.57>.
- He, B., & Ounis, I. (2004). *Inferring query performance using pre-retrieval predictors*. *International symposium on string processing and information retrieval*. Springer43–54.
- Herzig, J., Nowak, P. K., Müller, T., Piccinno, F., & Eisenschlos, J. (2020). TAPAS: Weakly supervised table parsing via pre-training. *Proceedings of the 58th annual meeting of the association for computational linguistics*4320–4333.
- Järvelin, K., & Kekäläinen, J. (2002). Cumulated gain-based evaluation of ir techniques. *ACM Transactions on Information Systems (TOIS)*, 20(4), 422–446.
- Kenter, T., & De Rijke, M. (2015). *Short text similarity with word embeddings*. *Proceedings of the 24th ACM international conference on information and knowledge management*1411–1420.
- Kusner, M., Sun, Y., Kolkin, N., & Weinberger, K. (2015). *From word embeddings to document distances*. *International conference on machine learning*957–966.
- Le, Q., & Mikolov, T. (2014). *Distributed representations of sentences and documents*. *International conference on machine learning*1188–1196.
- Li, Y., Zheng, R., Tian, T., Hu, Z., Iyer, R., & Sycara, K. (2016). *Joint embedding of hierarchical categories and entities for concept categorization and dataless classification*. *Proceedings of COLING 2016, the 26th international conference on computational linguistics: Technical papers*2678–2688.
- Limaye, G., Sarawagi, S., & Chakrabarti, S. (2010). Annotating and searching web tables using entities, types and relationships. *Proceedings of the VLDB Endowment*, 3(1–2), 1338–1347. <https://doi.org/10.14778/1920841.1921005>.
- Markovits, G., Shtok, A., Kurland, O., & Carmel, D. (2012). *Predicting query performance for fusion-based retrieval*. *Proceedings of the 21st ACM international conference on information and knowledge management*813–822.
- Metzler, D., & Croft, W. B. (2004). Combining the language model and inference network approaches to retrieval. *Information Processing & Management*, 40(5), 735–750.
- Metzler, D., & Croft, W. B. (2005). *A markov random field model for term dependencies*. *Proceedings of the 28th annual international ACM SIGIR conference on research and development in information retrieval*472–479.
- Mikolov, T., Chen, K., Corrado, G., & Dean, J. (Chen, Corrado, Dean, 2013a). *Efficient estimation of word representations in vector space*. *Proceedings of workshop at ICLR*.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (Sutskever, Chen, Corrado, Dean, 2013b). *Distributed representations of words and phrases and their compositionality*. *Advances in neural information processing systems*3111–3119.
- Mitra, B., Craswell, N., et al. (2018). An introduction to neural information retrieval. *Foundations and Trends® in Information Retrieval*, 13(1), 1–126.
- Nalisnick, E., Mitra, B., Craswell, N., & Caruana, R. (2016). *Improving document ranking with dual word embeddings*. *Proceedings of the 25th international conference companion on world wide web*83–84.
- Nguyen, T. T., Nguyen, Q. V. H., Weidlich, M., & Aberer, K. (2015). *Result selection and summarization for web table search*. *2015 IEEE 31st international conference on data engineering*. IEEE231–242.
- Pennington, J., Socher, R., & Manning, C. D. (2014). *Glove: Global vectors for word representation*. *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*1532–1543.
- Peters, M. E., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., & Zettlemoyer, L. (2018). Deep contextualized word representations. In M. A. Walker, H. Ji, & A. Stent (Eds.). *Proceedings of the 2018 conference of the north American chapter of the association for computational linguistics: Human language technologies, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1–6, 2018, volume 1 (long papers)* (pp. 2227–2237). Association for Computational Linguistics. <https://doi.org/10.18653/v1/n18-1202>.
- Pimplikar, R., & Sarawagi, S. (2012). *Answering table queries on the web using column keywords*. *Proceedings of the VLDB endowment*908–919.
- Roitman, H. (2018). *Enhanced performance prediction of fusion-based retrieval*. *Proceedings of the 2018 ACM SIGIR international conference on theory of information retrieval*195–198.
- Rubner, Y., Tomasi, C., & Guibas, L. J. (1998). *A metric for distributions with applications to image databases*. *Sixth international conference on computer vision (IEEE Cat. No. 98CH36271)*. IEEE59–66.
- Shaw, J. A., & Fox, E. A. (1994). *Combination of multiple searches*. *Proceedings of the third text retrieval conference, TREC 1994*. National Institute of Standards and Technology (NIST)105–108.
- Shen, Y., He, X., Gao, J., Deng, L., & Mesnil, G. (2014). *A latent semantic model with convolutional-pooling structure for information retrieval*. *Proceedings of the 23rd ACM international conference on conference on information and knowledge management*101–110.
- Shtok, A., Kurland, O., Carmel, D., Raiber, F., & Markovits, G. (2012). Predicting query performance by query-drift estimation. *ACM Transactions on Information Systems (TOIS)*, 30(2), 1–35.
- Siddiquie, B., Feris, R. S., & Davis, L. S. (2011). *Image ranking and retrieval based on multi-attribute queries*. *CVPR 2011*. IEEE801–808.
- Tao, Y., & Wu, S. (2014). *Query performance prediction by considering score magnitude and variance together*. *Proceedings of the 23rd ACM international conference on conference on information and knowledge management*1891–1894.
- Tax, N., Bockting, S., & Hiemstra, D. (2015). A cross-benchmark comparison of 87 learning to rank methods. *Information Processing & Management*, 51(6), 757–772.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1), 267–288.
- Tomlinson, S. (2004). *Robust, web and terabyte retrieval with hummingbird SearchServer™ at TREC 2004*. *Proceedings of the thirteenth text retrieval conference, TREC 2004*. National Institute of Standards and Technology (NIST).
- Tyree, S., Weinberger, K. Q., Agrawal, K., & Paykin, J. (2011). *Parallel boosted regression trees for web search ranking*. *Proceedings of the 20th international conference on world wide web*387–396.
- Vulić, I., & Moens, M.-F. (2015). *Monolingual and cross-lingual information retrieval models based on (bilingual) word embeddings*. *Proceedings of the 38th international ACM SIGIR conference on research and development in information retrieval*363–372.
- Yakout, M., Ganjam, K., Chakrabarti, K., & Chaudhuri, S. (2012). *InfoGather: entity augmentation and attribute discovery by holistic matching with web tables*. *Proceedings of the 2012 ACM SIGMOD international conference on management of dataSIGMOD 12*New York, NY, USA: Association for Computing Machinery97–108. <https://doi.org/10.1145/2213836.2213848>.
- Zamani, H., Croft, W. B., & Culpepper, J. S. (2018). *Neural query performance prediction using weak supervision from multiple signals*. *The 41st international ACM SIGIR conference on research & development in information retrieval*105–114.
- Zhang, M., & Chakrabarti, K. (2013). *InfoGather+: Semantic matching and annotation of numeric and time-varying attributes in web tables*. *Proceedings of the 2013 ACM SIGMOD international conference on management of dataSIGMOD 13*New York, NY, USA: Association for Computing Machinery145–156. <https://doi.org/10.1145/2463676.2465276>.



- Zhang, S., & Balog, K. (2017). *Entitables: Smart assistance for entity-focused tables*. *Proceedings of the 40th international ACM SIGIR conference on research and development in information retrieval*255–264.
- Zhang, S., & Balog, K. (Balog, 2018a). *Ad hoc table retrieval using semantic similarity*. *Proceedings of the 2018 world wide web conference*1553–1562.
- Zhang, S., & Balog, K. (Balog, 2018b). *On-the-fly table generation*. *The 41st international ACM SIGIR conference on research & development in information retrieval*595–604.
- Zhang, S., & Balog, K. (2019). *Auto-completion for data cells in relational tables*. *Proceedings of the 28th ACM international conference on information and knowledge management*761–770.
- Zhang, S., & Balog, K. (2020). *Web table extraction, retrieval, and augmentation: a survey*. *ACM Transactions on Intelligent Systems and Technology*, 11(2), 13:1–13:35. <https://doi.org/10.1145/3372117>.
- Zhou, Y., & Croft, W. B. (2006). *Ranking robustness: a novel framework to predict query performance*. *Proceedings of the 15th ACM international conference on information and knowledge management*567–574.
- Zhou, Y., & Croft, W. B. (2007). *Query performance prediction in web search environments*. *Proceedings of the 30th annual international ACM SIGIR conference on research and development in information retrieval*543–550.
- Zuccen, G., Koopman, B., Bruza, P., & Azzopardi, L. (2015). *Integrating and evaluating neural word embeddings in information retrieval*. *Proceedings of the 20th australasian document computing symposium*1–8.