From Scattered Sources to Comprehensive Technology Landscape : A Recommendation-based Retrieval Approach

Duong Chi Thang[1], Percia David Dimitri[2,3], Dolamic Ljiljana[3], Mermoud Alain[3], Lenders Vincent[3], and Aberer Karl[1]

[1] Distributed Information Systems Laboratory, EPFL
[2] Information Science Institute, University of Geneva
[3] Cyber-Defence Campus, armasuisse Science and Technology, EPFL

Abstract. Mapping the technology landscape is crucial for market actors to take informed investment decisions. However, given the large amount of data on the Web and its subsequent information overload, manually retrieving information is a seemingly ineffective and incomplete approach. In this work, we propose an end-to-end recommendation based retrieval approach to support automatic retrieval of technologies and their associated companies from raw Web data. This is a two-task setup involving (i) technology classification of entities extracted from company corpus, and (ii) technology and company retrieval based on classified technologies. Our proposed framework approaches the first task by leveraging DistilBERT which is a state-of-the-art language model. For the retrieval task, we introduce a recommendation-based retrieval technique to simultaneously support retrieving related companies, technologies related to a specific company and companies relevant to a technology. To evaluate these tasks, we also construct a data set that includes company documents and entities extracted from these documents together with company categories and technology labels. Experiments show that our approach is able to return 4 times more relevant companies while outperforming traditional retrieval baseline in retrieving technologies.

Keywords: Technology monitoring · Information retrieval · Entity-based retrieval · Technology classifier · Recommender system

1 Introduction

The expanding and accelerating pace of technology development continuously reshapes the technological landscape [33]. Depicting an up-to-date and holistic map of organizations that develop, implement or sell a given technology is an important business challenge, should a technology be novel or established [28]. In such a dynamic environment, market actors are increasingly confronted by an information overload issue, as an ever raising amount of heterogeneous and unstructured market information needs to be collected, stored, cleaned, structured and analyzed [29]. Hence, the automated analysis of the complex network of organizations and technologies is a key business intelligence necessity not only for public entities, but also for private investors [13].

Such an information retrieval necessity has triggered various business intelligence and technology monitoring procedures, which have been either developed by in-house R&D efforts of organizations or by academic actors (*e.g.*, [9]). Often, such procedures consist of non-automated approaches that struggle to tackle the information overload challenge, as they do not provide a reliable, systematic and scalable information retrieval methodology for mapping the technology market and determine which companies are developing/commercializing which technology [20]. These extant procedures are mainly based on frameworks that find documents matching query terms instead of finding entities *per se* [10]. Yet, finding entities is central when it comes to investigate the technological landscape and finding new technologies. Even though related works have partially investigated such an issue, to the best of our knowledge, no work has provided an end-to-end automated information retrieval framework for finding technology related entities and mapping the technological landscape through a recommendation based perspective. In this work, we develop such a framework in order to map the industrial technology landscape, which would be highly applicable in several domains where a comprehensive understanding of the landscape is required. For instance, in cybersecurity, by mapping the technological landscape of cybersecurity, decision-makers will be provided relevant information for taking more informed purchase decisions[22]. Similarly, in the stock market, having a comparative analysis on technological advances of competing companies is crucial for making correct investment decisions.

Our framework is a two step approach that first classifies the entities into technologies before performing company and technology retrieval. We use DistilBERT, which is a state-of-the-art language model, to construct the entity embeddings allowing us to achieve high accuracy in technology classification. For retrieval, we propose a recommendation based approach that takes into account both the technologies, companies and their relationship. This recommendation approach enables better retrieval results in comparison with state-of-the-art learning based recommendation approaches such as GMF [14,27], MLP [14] and NCF [14]. Our results show that our approach is able to return 4 times more relevant companies in company to company retrieval in comparison with traditional tf-idf retrieval approach.

The remainder of this article proceeds as follows. Section 2 grounds this research by emphasizing different methodologies of information retrieval for technology monitoring, and the research gaps. Section 3 details the data collection. Section 4 presents the information retrieval model and framework. Section 5 presents the entity extraction methodology and the technology classification we use. Section 6 presents the technology retrieval method we used. Section 7 presents the empirical evaluation, while Section 8 concludes and set the path for further research.

2   Related Works

The information overload triggered by the big data era has motivated researchers and practitioners to develop numerous automated information retrieval methods by using different yet often complementary approaches [25,18,32]. Such methods have been widely used in fields as digital libraries [16], information filtering and recommender systems [6,15], media search [26] and search engines [2].

In the field of technology monitoring and forecasting – and more specifically in the specific context of technology landscape monitoring –, numerous works have been published, involving the extraction of patents [19], scientific articles [12] and social media analysis [7]. These technology retrieval methods can be classified into either a *keyword-based* or a *entity-based* approach. Most of the existing methods are keyword-based in which the queries and the data are mostly plain texts [4].

Hossari et al. (2019) proposed an automated framework for detecting the existence of new technologies in texts, and extract terms used to describe new technologies [17]. Aharonson & Schilling developed a framework that captures the distance between patents, and a company's technological footprint. Their framework also enables to measure the proximity of technological footprints between organizations [3]. Tang and Liu (2008) presented a three-layer model for technology forecasting based on text mining techniques, incorporating a collection layer, an analysis layer and a representation layer, before overlapping a semantic web based approach in order to map the industrial technology landscape [35]. On the other hand, entity-based techniques [37] require linking a piece of text to an entity in a knowledge base such that retrieval is done on the entities instead of the raw texts. Woon and Madnick (2008) developed an information retrieval framework for visualizing the technology landscape by exploring the use of term co-occurence frequencies as an indicator of semantic closeness between pairs of entities [21]. In the field of energy related technologies, Mikheev (2018) developed an ontology based data access framework under a semantic approach to query complex datasets, creating an automated mapping procedure to connect data to ontology entities [24]. By applying a semantic approach, Sitarz et al. (2012) developed an automated framework for identifying thematic groups of scientific publications based on clustering of sets of co-occurence words and and financial-analysis techniques for trends detection and forecasting [34].

In our setting, we opt for the entity-based approach as it allows us to handle different mentions of the same entity while enabling better retrieval accuracy due to external information from the knowledge base. However, to the best of our knowledge, little has been done when it comes to apply information retrieval methodologies with the aim of presenting a holistic and comprehensive monitoring and mapping of the industrial technology landscape. In order to do so, a technique needs to consider both the technologies and companies at the same time. The following steps have to be undertaken: (i) an entity fishing approach needs to be applied for extracting and classifying technology entities; (ii) then, these technology entities need to be linked to specific companies; (iii) and finally, these entities must be ranked according to their level of relevance. Demartini et al. (2009) provided a formal model for entity fishing and ranking [10]. Yet, to the best of our knowledge, no such work has been deployed in the context of technology landscape monitoring. Similarly, Balog et al. (2012) developed a framework for assessing the strength of association between a topic and a person (*i.e.*, expertise retrieval) [5]. Yet, to the best of our knowledge, no such framework has been applied in the context of technology landscape monitoring for linking technology entities and company entities.

Table 1. Dataset statistics.

Dataset	#comps.	#terms
Website	7907	22104
Patent	6814	7796
Jobs	3894	2532
Twitter	790	3236
Total	18339	27977

3   Dataset construction

To the best of our knowledge, there is no public dataset available for the technology retrieval and classification task. As a result, we aim to create and publish such a dataset to further research in this field. We have defined the following requirements for our dataset. First, the dataset should be multilingual as the technology retrieval task should be language independent. Second, the dataset should be realistic and coming from real-world data as this would enable objective evaluation of any proposed approach.

In the following, we discuss our data construction process. The dataset is constructed by first crawling different data sources that are publicly available on the Internet. As we aim to develop a language-agnostic framework, we decide to construct a dataset based on Swiss companies as Switzerland is a multilingual country with French, German and Italian as official languages while English is a working language. For each company, we intend to collect all possible documents that are related to the company's actual activities. This involves the following data sources: the company's *website*, the company's *job postings*, the *patents* and the company's *tweets*. In addition, to maintain an up-to-date dataset about the companies, we periodically crawl data from the above data sources. The above data sources are selected as they could provide different perspectives regarding a company's technology offering. The statistics of the dataset[4] is shown in Table 1.

3.1   Collecting data

To collect the data, we first need a list of Swiss companies. We obtain the list of companies registered in Switzerland from the federal Central Business Name Index (Zefix)[5]. Additional information regarding given company is then extracted from the corresponding cantonal commercial register record. These records provide, among others, information on location, people, type of company.

Websites. As commercial registers do not provide the information on regarding websites of the registered companies, we need to find the company's website based on the company's name. We first clean the company name (e.g. removing Ltd.) before performing Google search. The information extracted from commercial registers in combination with the results of the previously mentioned search are then put through the

---

[4]Note that there are overlapping companies and terms among different data sources.
[5]https://zefix.ch

classifier to link the company to a correct website. We then crawl the pages from the detected website, using their text for entity extraction.

Patents. We use Patents data from United States Patent and Trademarks Office (USPTO). PatentView[6] provides half yearly data dump of the USPTO database, which we inject directly into our system. The patents are linked to the companies based on the location and assignee information, while the title and the abstract of the patents are used to extract the entities.

Jobs. The Indeed[7] is used as a source of information related to jobs. We perform weekly per canton search for jobs, retrieving for each job the following information: title, description, company name and the original posting. Job's title and description are used for technology annotation, while the linking to the company is based on it's name and location.

Tweets. For each company, we look up its Twitter handle from Crunchbase which is a commercial database of company information. From the handle, we collect the latest 3000 tweets for each company using sempi.tech which is a social listening framework. We use these tweets for entity extraction.

3.2 Entity Extraction

From the documents collected in the previous step, we use DBPedia Spotlight (DBPS)[8], an open source tool to automatically annotate the mentions of the DBPedia entities within the text content. Dbpedia is a multi-domain ontology derived from Wikipedia. Each DBPedia entity is an URI with the prefix *http://dbpedia.org/resource/* followed by the identifier of the corresponding Wikipedia article. DBPS allows us to not only extract entities but also link the entities to their corresponding DBPedia entities. DBPS is also capable of handling multiple languages, which is important in our setting. For each data source, we perform entity extraction independently. The output of this step is a list of entities and their corresponding number of appearance in the data source of a company. We store the output in JSON line format where each line is a triple of company, DBPedia entity and number of entity occurrences. Note that while our data sources are heterogeneous as they come from different domains, DBPS allows us to standardize them. Whether it is a website, a patent, a job posting or a tweet, DBPS extracts all the candidate technology mentions while ignoring irrelevant data. These technology mentions are the input to any technology classification or retrieval technique.

3.3 Data labeling

Technology labeling. To support the evaluation of a technology classifier, we also provide a list of Wikipedia terms and their labels whether a term is about a technology or not. Even though Technology is one of Wikipedia's Main Topic Classification (MTC) categories, one can not rely on this concept to extract all technology related articles, as the categories within the Wikipedia graph are very loosely related("is related to").

---

[6]https://patentsview.org
[7]http://indeed.ch

We approach the labeling in a top-down approach where we aim to label the MTC categories. First, Wikipedia directed categories graph was cleaned by removing hidden categories, admin and user pages, followed by regular expression filters removing categories referring to companies, brands, currencies etc. We than calculate the shortest path to each of the 28 MTC, and retain the categories having the shortest path to Technology, Science, or Engineering topics. This process resulted in the list of 7876 categories. These categories were than manually labeled as technology or non technology, resulting with 1356 categories being labeled an technology. This is the only manual step in our data construction process. An article is then considered to be a technology if it is directly connected to a category labeled as such.

Company categorization. Crunchbase maintains a database of companies and their detail information where the data are manually curated by Crunchbase staff and online contributors [8]. Each company is associated with several categories describing its main activities. For instance, Roche which is a pharmaceutical company based in Switzerland is categorized as Biotechnology, Health Care, Health Diagnostics and Pharmaceutical. From the companies collected in the above steps, we crawl Crunchbase to obtain their categories. The categories can be considered as pseudo-labels for our com-com and tech-com retrieval tasks.

4 Model and Approach

We develop a unified framework to classify entities into technologies and to perform technology related retrieval. This requires solving two tasks of technology classification and technology retrieval.

4.1 Model

Our framework considers a set of companies $C = \{c_1, \cdots, c_n\}$ and a set of entities $E = \{e_1, \cdots, e_m\}$. The connection between an entity e_i and a company c_j can be observed through several data sources. For each data source, we measure this connection by the number of times the entity e_i is mentioned by the company c_j. We can present these connections for a specific data source by an *interaction matrix* $\boldsymbol{M}$ where M_{ij} is the occurrence frequency of the entity e_i in the corpus of company c_j. We also denote these matrices by $M = \{\boldsymbol{M}_1, \cdots, \boldsymbol{M}_k\}$ where $\boldsymbol{M}_i$ represents an interaction matrix.

Tasks. We have 3 retrieval tasks: company to technology (com-tech) retrieval, company to company (com-com) retrieval and technology to company (tech-com) retrieval. Com-tech and com-com retrieval are connected as both require an accurate representation of a company by its technologies. Two companies are similar if they have similar technology representation. Similarly, for tech-com retrieval, we need a good representation of a technology by its companies. There is a mutual reinforcing relationship between company and technology representation. A good company representation requires knowing similar technologies while knowing company similarity is helpful in constructing a good technology representation. This means we need a common model to approach these three different retrieval tasks.

---

[8]crunchbase.com

4.2   General approach

Our framework takes as input the entities extracted from the company corpora. These entities are identified using entity extraction frameworks such as DBPedia-spotlight. Each entity is associated with its description. This information is used in the second step to classify the entities into technologies or not. Our technology classifiers are constructed using BERT as a feature extractor. The technologies obtained from the previous step are used as input for the retrieval step. We reformulate the retrieval problem as a recommendation problem based on collaborative filtering where the technologies are "recommended" to a company if the technologies are considered to be related to the company's activities. Casting this as a recommendation problem has several benefits. The relevancy of a technology to a company can be measured more accurately if similar companies in the same domain are considered. This also means that missing technologies in a company corpus can be recovered by considering similar companies. We propose a technology recommendation model that extends traditional matrix factorization by integrating both the technology meaning and the interactions between technologies and companies.


Fig. 1. From data sources to company and technology retrieval

5   Technology Classification

As the input entities to our system are extracted using an entity extraction framework such as DBPedia-spotlight [23], they cover all possible domains while we are interested only in technologies. To this end, we develop a technology classifier to filer out unrelated entities. Note that DBPedia-spotlight also performs entity linking where each entity is linked to a DBPedia page or a Wikipedia article describing this entity. We leverage this description to construct our technology classifier. For each entity, we extract its abstract which we define to be the description from DBPedia or the first paragraph of its Wikipedia article whichever available.

BERT-based encoder. We use DistilBERT [31] which is a state-of-the-art language model while being light-weight and fast to train to construct abstract embeddings. We use these abstract embeddings as representations for the entities. We pass each abstract

through DistilBERT to obtain the token embedding of the [CLS] token, which is commonly used as the sentence or paragraph embedding. We denote this token embedding as z_e where e denotes an entity.

Although DistilBERT is a distilled version of BERT which was trained on English wikipedia, distillation may incur information lost. To this end, we propose to fine-tune DistilBERT to better capture the abstract meaning which usually contain several scientific terminology, therein, obtaining better entity representation. We fine-tune DistilBERT for our technology classification in an end-to-end manner where we feed the abstracts through the model to obtain the abstract embeddings. These abstract embeddings are then fed through a linear classifier to get the technology predictions. The parameters of DistilBERT and the linear classifier are updated together in an end-to-end manner using SGD [30] with the binary cross entropy loss as the loss function. We observe that with finetuning, we are able to achieve better accuracy.

Embedding refinement. To further increase the capacity of our classifier, we pass z_e through several layers including a dropout layer to reduce overfitting. More precisely, z_e is passed through a multi-layer neural network with 2 blocks where each block is a linear layer followed by a BatchNorm layer and a sigmoid non-linearity. A block can be formulated as follows:

$$t_e = \sigma(BatchNorm(\boldsymbol{W} z_e + b_1))$$

Between 2 blocks, we also use a dropout layer to reduce overfitting. We observe that be refining the embedding this way, we can achieve better accuracy than non-refinement.

6   Technology Retrieval

The case for retrieving as recommending. There are several requirements for the retrieval model of this task. First, the technologies retrieved for a company must be derived from the technologies mentioned in the company's corpus as these technologies are the most likely ones to reflect the company's actual activities. As each technology has a different level of relevancy to the company, com-tech retrieval considering only observed technologies is a *technology ranking* problem. Second, it is safe to assume that a company's corpus may not contain all the technologies the company is working on. Companies may not publicly mention a technology to keep a competitive advantage or it could simply be due missing data. For these reasons, the com-tech retrieval is also a *technology discovery* problem. To solve both problems at the same time, we propose a recommendation model that identifies potentially related technologies of a company by looking at similar companies while measuring the relevancy between every pair of technology and company for ranking technologies.

Recommendation model. Given a set of companies $\mathbb{C} = \{c_1, \cdots, c_n\}$ and the set of technologies $\mathbb{T} = \{t_1, \cdots, t_m\}$ obtained after the technology classification step, our recommendation model takes as input a com-tech interaction matrix $\boldsymbol{M} \in \mathbb{R}^{n \times m}$ as

$$M_{ij} = \begin{cases} f(c_i, t_j), \text{if there is a mention of } t_j \text{ in any data source of } \mathbb{C}_i \\ \emptyset, \text{if there is no mention of } t_j \text{ in } \mathbb{C}_i \end{cases}$$

The function $f(c_i, t_j)$ captures the importance of t_j according to c_i. This importance can be measured by the number of times t_j occurs in the corpus of c_i rescaled by some weighting scheme. In our setting, we measure the importance per data source using tf-idf before combining all data sources:

$$f(c_i, t_j) = \sum_k w_k f_k(c_i, t_j)$$

where $f_k(c_i, t_j)$ is the importance function of data source k which is measured by the tf-idf value of t_j considering each company as a "document". w_k is its associated weight which captures our perceived relevance of the data source to our retrieval tasks.

A high value of $f(c_i, t_j)$ does not mean the company c_i is actually working on technology t_j. For instance, during the pandemic, there could be several mentions of the word "vaccine" but it does not necessarily mean that a certain company is developing a vaccine. This example shows that answering com-tech retrieval by only looking at the company corpus could be problematic. We can have the same argument for $M_{ij} = \emptyset$. This does not mean the company c_i has no activity related to t_j. It could be the case that the company is working on this technology but it has not mentioned it *yet* in the corpus.

The above com-tech interaction model is akin to collaborative filtering with implicit feedback [27,14]. To answer the com-tech retrieval problem, we first need to solve the recommendation problem where we need to estimate the unobserved entries of M. The estimation is usually done by learning a model $\hat{M}_{i,j} = f(c_i, t_j | \Theta)$ where $\hat{M}_{i,j}$ is the estimated score of M_{ij}, f is a parameterized function that predicts the interaction score between c_i and t_j, Θ denotes the parameters of f.

Semantic-aware matrix factorization. We propose a semantic-aware recommendation model extending traditional matrix factorization (MF) approach. In MF, each company and each technology is represented by an embedding in a shared latent space. The technology and the company embeddings are learned such that they can reconstruct the interaction matrix M. More precisely, let $c_i \in \mathbb{R}^d$ and $e_j \in \mathbb{R}^d$ be the embeddings of company c_i and technology t_j respectively. Then, MF aims to estimate the relevancy score M_{ij} by:

$$\hat{M}_{ij} = f(c_i, t_j | c_i, e_j) = c_i^T e_j$$

However, MF considers the technologies to be independent even if they are semantically related such as "deep learning" and "machine learning". To this end, we propose to incorporate the meaning of the technologies into MF while extending the model capacity by passing the technology embedding through several linear layers. In the following, we describe in detail the layers of our architecture.

Semantic embedding: We capture the technology meaning using BERT [31,11] as a feature extractor over the technology abstract:

$$s_i^{(0)} = f_{BERT}(a_i)$$

where s_i is the semantic technology embedding and a_i is the abstract of technology t_i.

MLP layers: The semantic technology embedding is passed through several MLP layers to further reduce the size of the embedding while increasing the model capacity.

$$\boldsymbol{s}_i^{(1)} = \boldsymbol{W}_1 \boldsymbol{s}_i^{(0)} + b_1$$
$$\cdots \tag{1}$$
$$\boldsymbol{s}_i^{(k)} = \boldsymbol{W}_k \boldsymbol{s}_i^{(k-1)} + b_k$$

where $\boldsymbol{W}_i, b_i, \sigma$ are the weight matrix, the bias and the non-linearity.

Combination layer: To obtain the final technology embedding $\boldsymbol{t}_i$, we combine the semantic technology embedding $\boldsymbol{s}^{(k)}$ with the raw technology embedding from MF $\boldsymbol{e}_i$ by summing them: $\boldsymbol{t}_i = \boldsymbol{s}^{(k)} + \boldsymbol{e}_i$. The summation allows us to save model's parameters in comparison with concatenation while it is also inspired by transformer architecture [36,11] where positional encodings are added to the word embeddings.

Model learning. To learn the parameters, traditional MF uses the squared loss between the predicted and actual interaction score: $\mathcal{L}_{sqr}(\Theta) = \sum_{c_i, t_j \in \boldsymbol{M}} (M_{ij} - \hat{M}_{ij})$. However, such a method does not consider the unobserved entries directly. To this end, we follow the pairwise learning approach that aims to optimize the relative ranking between technologies. We use the margin hinge loss which is defined as follows:

$$\mathcal{L}_{hinge}(\Theta, c_i, t_j, t_k) = max(0, m + M_{ij} - \hat{M}_{ik})$$

where c_i, t_j is an observed pair of company and technology while t_k is a negative sample meaning c_i, t_k is an unobserved entry of the interaction matrix $\boldsymbol{M}$.

Recommendation-based Retrieval. Then, the com-tech retrieval problem can be answered by ranking the technologies $\mathbb{T}$ with respect to a company c based on their interaction scores. More precisely, let $\boldsymbol{N}$ to be the com-tech interaction matrix after the unobserved entries are estimated. The top-k com-tech retrieval result for a company c_i is a list of technologies ordered by their interaction scores $\boldsymbol{N}$.

7   Empirical Evaluation

7.1   Experimental setup

Datasets. We evaluate our model on the constructed dataset described in Section 3. As our dataset is crawl constantly which makes it difficult for evaluation, we fix the dataset used in the experiments to be the data collected before 01/04/2020. This snapshot and the code are publicly available at https://figshare.com/s/c014bb8565705e74dd1b.

Metrics. To evaluate the results of com-com retrieval, we leverage the company categorisation from Crunchbase. We measure the quality of com-com retrieval by the number of overlapping categories between the query company and the results. More precisely, let $\mathbb{C}(c)$ denote the set of categories of company c. We define the retrieval accuracy for a company c considering top-k most relevant results as follows:

$$P@k(c) = \frac{\sum_{i=\overline{1,k}} |\mathbb{C}(c) \cup \mathbb{C}(c_i)|}{k} \tag{2}$$

This metric can be extended to a set of companies C as $P@k(C) = \frac{\sum_{c \in C} P@k(c)}{|C|}$.

While it is "straightforward" to evaluate the search results for com-com, the evaluation for com-tech and tech-com is more challenging as there is no available ground truth. For tech-com search, we follow the approach of com-com retrieval where we label each technology by the Crunchbase categories. This is akin to consider each technology as a "company". The number of technologies to be labelled is usually small as we are interested in only important technologies. To this end, we have labeled 119 technologies which are considered important in the cybersecurity domain [1]. The retrieval accuracy $P@k(t)$ for a technology t is defined similarly as in Equation 2. On the other hand, for com-tech retrieval, this approach is not practical as for each company, the list of retrieved technologies is very large. To this end, we opt for a qualitative evaluation.

Baselines. For technology classification, we construct a baseline using SVM on tf-idf featurization. More precisely, we construct a vector representing a Wikipedia category by combining the vector distances of a category to each of the Wikipedia's MTC and its TF-IDF weighted bag of words (BOW) representation. The weighted BOW representation of the given category is created from the stemmed text obtained by concatenating the abstracts of all Wikipedia articles directly connected to it. Mutual information based feature reduction than resulted in a vector of the length 1000. These vectors are used as the input features for the classifier. For retrieval, we first compare with a tf-idf retrieval approach where the tf-idf values are also the relevancy of the technologies in com-tech retrieval. For com-com retrieval, it is an tf-idf weighted version of Jaccard similarity where each company is represented by its set of technologies. We also compare our recommendation-based retrieval approach with other recommendation models including GMF [14,27], MLP [14] and NCF [14]. The above baselines are selected as to the best of our knowledge, there is no public implementation of technology classification and retrieval techniques. The above baselines represent the best starting points for these tasks.

Environments. Our experiments ran on an Intel Xeon CPU E5-2620 v4 @ 2.10GHz server with a Titan V GPU with 12GB VRAM and 128GB RAM. Our model was implemented using Pytorch 1.7.1 and Spotlight as the recommendation framework and DistilBERT from HuggingFace as the language model.

7.2 Effectiveness of technology classification

Quality of technology classifiers. In this experiment, we analyze the correctness of our technology classifiers. We compare our proposed classifier using BERT with the baseline classifiers based on tf-idf and other BERT models. For this experiment, we compare these approaches on three metrics: accuracy, f1-score and AUC. We use k-fold cross validation with a 80-20 split.

The results show in Table 2 confirms the benefit of fine-tuning and our refinement step. Our proposed approach outperforms the baselines on all metrics. The difference between using tf-idf as feature and BERT is 0.1, 0.02 and 0.04 for F1-score, accuracy and AUC respectively. This is expected as large language models trained on large text corpora are able to capture word meanings better. We also observe that fine-tuning improves accuracy in comparison with using BERT without fine-tuning.

Table 2. Comparison of technology classifiers

	F1-score	Accuracy	AUC
tf-idf	0.531	0.781	0.819
DistilBERT w/o refine w/o finetune	0.499	0.771	0.811
DistilBERT + refine + w/o finetune	0.597	0.734	0.806
FDistilBERT + refine + w/ finetune (Ours)	**0.639**	**0.799**	**0.857**

7.3 Effectiveness of technology retrieval

In this experiment, we analyze our proposed recommendation-based retrieval model. We compare our model with a tf-idf based retrieval where each technology is associated with a tf-idf value while each company is represented by a tf-idf vector. We also compare our approach with several recommendation models including 1) Generalized MF [27] which is a generalized version of MF, 2) MLP [14] which is a multi-layer recommendation model starting from random vectors and 3) NCF [14] or Neural Collaborative Filtering which is a recommendation model based on deep learning.

Table 3. Effects of retrieval model

	Company-company retrieval				Technology-company retrieval			
	top-5	top-10	top-15	top-20	top-5	top-10	top-15	top-20
MF	3.121	3.556	3.568	3.568	1.078	1.882	2.816	3.461
MLP	3.215	3.819	3.836	3.836	1.065	2.053	3.079	3.539
NCF	2.905	3.103	3.105	3.105	1.118	2.211	2.947	3.5
BERT	4.155	5.986	6.636	6.722	1.986	3.211	3.882	5.066
MF+BERT	**4.458**	**7.004**	**8.447**	**8.845**	**2.197**	**3.684**	**5.105**	**6.289**

The experimental results shown in Table 3 show that our proposed model based on BERT embeddings as initial technology embeddings are better than the baselines. The difference between our worst model and the best baseline is 0.6 at top-5 for com-com retrieval and 1.8 at top-5 for tech-com retrieval. This can be explained by the fact that our models can capture the meaning of the technologies while the baselines consider the technologies to be independent. Among our proposed models, adding the raw technology embedding from MF with the BERT technology embedding is better than using the BERT embedding alone. We can attribute this to the larger capacity of our model which helps in capture the interaction matrix better.

7.4 Parameter sensitivity

Effects of margin. We vary the margin of the hinge loss from 0.001 to 0.1 to analyze its effects on the retrieval results. The experimental results are shown in Figure 2. We observe that the number of overlapping categories tends to increase with the margin. For instance, the number of overlapping categories for top-5 com-com retrieval is 3.73 when the margin is 0.001 but it increases to 4.95 when the margin is 0.1. This is expected

Fig. 2. Effects of margin

as with the larger margin, our model tends to generalize as it aims to capture common technologies between the companies. We observe this phenomenon clearly from Table 5 that we obtain more specific technologies with smaller margin. With larger margins, generic technologies that are shared among different companies are more representative than more specific ones. This experiment confirms our ability to control the specificity of the retrieval results by changing the margin of the hinge loss.


Fig. 3. Effects of embedding size

Embedding size. In this experiment, we analyse the effects of the embedding size on the retrieval results. We vary the embedding size from 32 to 512. Results in Figure 3 shows that as the embedding size increases, we can retrieve companies better for both tech-com and com-com retrieval tasks. This is expected as increasing embedding size also improves the model capacity. However, there is a trade-off in increasing embedding size as it incurs longer training time as shown in Figure 4. The difference in training time between embedding size of 32 and 512 is 3 times. However, even with the largest embedding size, the training time per epoch is still very fast - only around 1.5 second.

7.5 End-to-end comparison

Having evaluated the individual components of our solution, we turn to its end-to-end performance in comparison with the baseline. Table 4 compares the performance of our
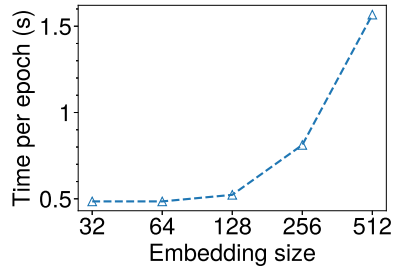
Fig. 4. Training time per epoch.

	com-com retrieval		tech-com retrieval	
	tf-idf	Ours	tf-idf	Ours
top-5	1.2916	**4.458**	2.0161	**2.197**
top-10	1.3022	**7.004**	3.4838	**3.648**
top-15	1.3022	**8.447**	4.3710	**5.105**
top-20	1.3022	**8.845**	4.7419	**6.289**

Table 4. End-to-end evaluation

approach with a tf-idf based retrieval approach which uses tf-idf as feature for technology classifier and technology retrieval. Our approach leads to significantly better retrieval results in both retrieval tasks. Our model is nearly 4 times better than the baseline in the com-com retrieval at top-5 while the difference is 0.18 for tech-com retrieval. First, this can be attributed to our approach's better technology classifier by using language model. Second, our recommendation retrieval model considers both the companies, technologies and their relationships as the same time, while the technologies and companies in the tf-idf model are handled independently. This enables our model to leverage the similarity of companies to support technology retrieval and vice versa.

Table 5. Com-tech retrieval

Acronis AG			InterHype SARL		
Ours-0.01	Ours-0.1	tf-idf	Ours-0.01	Ours-0.1	tf-idf
CyberTruck	Encryption	Cloud computing	Computer security	Cloud computing	Nous
Virtual machine	Communication	Backup	Automatic train protection	Computer science	Sand
Cloud storage	Virtualization	Disaster recovery	SMS	Computer security	Antiseptic
Off-site data protection	Internet	Web server	Off-site data protection	Digital transformation	Habitat
Encryption	Personal firewall	Ransomware	Backup	Information security	Glass

7.6 Qualitative analysis

In this experiment, we analyze the retrieval results qualitatively. Table 5 shows the com-tech retrieval results where we search for cyber security companies. For com-tech retrieval, as discussed above, we are able to control the specificity of the results by changing the margin. In addition, our proposed model is able to return less noise in comparison with tf-idf one as tf-idf model may return non-technological terms due to the quality of its classifier. This phenomenon can be seen for instance in the search for InterHype Sarl which is a cyber security company. For com-com and tech-com retrieval, due to space constraint, we do not include them. However, we observe that our model can return companies that are in the same domains as the queried company or technology. This is in line with the quantitative result observed in previous experiments.

8   Conclusion and Future works

In this paper, we propose an end-to-end framework to first extract and classify technological mentions from company corpuses and then, retrieve related technologies and

companies. Our technology classifier is based on DistilBERT model with finetuning and refinement to achieve better accuracy while our recommendation-based retrieval model enables more relevant results. In the future works, we aim to allow users to reformulate queries to better capture users intentions. Another possible research direction is to combine technology classification and entity extraction for more accurate results.

References

1. Security-relevant technology and industry base. https://www.ar.admin.ch/en/beschaffung/ruestungspolitik-des-bundesrates/offset. accessed: 2021-06-16
2. Aggarwal, C.C.: Information retrieval and search engines. In: Machine Learning for Text, pp. 259–304. Springer (2018)
3. Aharonson, B.S., Schilling, M.A.: Mapping the technological landscape: Measuring technology distance, technological footprints, and technology evolution. Research Policy **45**(1), 81–96 (2016)
4. Arai, K.: Extraction of keywords for retrieval from paper documents and drawings based on the method of determining the importance of knowledge by the analytic hierarchy process: Ahp. International Journal of Advanced Computer Science and Applications **11**(10) (2020)
5. Balog, K., Fang, Y., de Rijke, M., Serdyukov, P., Si, L.: Expertise retrieval **6**(2-3), 127–256 (Feb 2012)
6. Belkin, N.J., Croft, W.B.: Information filtering and information retrieval: Two sides of the same coin? Communications of the ACM **35**(12), 29–38 (1992)
7. Chen, X., Han, T.: Disruptive technology forecasting based on gartner hype cycle. In: TEMSCON. pp. 1–6. IEEE (2019)
8. Daiber, J., Jakob, M., Hokamp, C., Mendes, P.N.: Improving efficiency and accuracy in multilingual entity extraction. In: ICSS. pp. 121–124. New York, NY, USA (2013)
9. Daim, T.U., Chiavetta, D., Porter, A.L., Saritas, O.: Anticipating future innovation pathways through large data analysis. Springer (2016)
10. Demartini, G., Gaugaz, J., Nejdl, W.: A vector space model for ranking entities and its application to expert search. In: ECIR. pp. 189–201. Springer (2009)
11. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805 (2018)
12. Dotsika, F., Watkins, A.: Identifying potentially disruptive trends by means of keyword network analysis. Technological Forecasting and Social Change **119**, 114–127 (2017)
13. Durak, U.: Flight 4.0: The changing technology landscape of aeronautics. In: Advances in Aeronautical Informatics, pp. 3–13. Springer (2018)
14. He, X., Liao, L., Zhang, H., Nie, L., Hu, X., Chua, T.S.: Neural collaborative filtering. In: WWW. pp. 173–182 (2017)
15. Heggo, I.A., Abdelbaki, N.: Data-driven information filtering framework for dynamically hybrid job recommendation. In: Intelligent Systems in Big Data, Semantic Web and Machine Learning, pp. 23–49. Springer (2021)
16. Hersh, W.R.: Information retrieval and digital libraries. In: Biomedical Informatics, pp. 613–641. Springer (2014)
17. Hossari, M., Dev, S., Kelleher, J.D.: Test: A terminology extraction system for technology related terms. In: Proceedings of the 2019 11th International Conference on Computer and Automation Engineering. pp. 78–81 (2019)
18. Kaur, P., Pannu, H.S., Malhi, A.K.: Comparative analysis on cross-modal information retrieval: a review. Computer Science Review **39**, 100336 (2021)

19. Kim, M., Park, Y., Yoon, J.: Generating patent development maps for technology monitoring using semantic patent-topic analysis. Computers & Industrial Engineering **98**, 289–299 (2016)
20. Loveridge, D., Cagnin, C.: Fta as due diligence for an era of accelerated interdiction by an algorithm-big data duo. In: Anticipating Future Innovation Pathways Through Large Data Analysis, pp. 3–23. Springer (2016)
21. Madnick, S., Woon, W.L.: Semantic distances for technology landscape visualization (2008)
22. Mahmood, T., Afzal, U.: Security analytics: Big data analytics for cybersecurity: A review of trends, techniques and tools. In: 2013 2nd national conference on Information assurance (ncia). pp. 129–134. IEEE (2013)
23. Mendes, P.N., Jakob, M., et al.: Dbpedia spotlight: shedding light on the web of documents. In: Proceedings of the 7th international conference on semantic systems. pp. 1–8 (2011)
24. Mikheev, A.: Ontology-based data access for energy technology forecasting. In: IWCI 2018. pp. 147–151. Atlantis Press (2018)
25. Munir, K., Anjum, M.S.: The use of ontologies for effective knowledge modelling and information retrieval. Applied Computing and Informatics **14**(2), 116–126 (2018)
26. Rao, J., Yang, W., Zhang, Y., Ture, F., Lin, J.: Multi-perspective relevance matching with hierarchical convnets for social media search. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 33, pp. 232–240 (2019)
27. Rendle, S., Freudenthaler, C., Gantner, Z., Schmidt-Thieme, L.: Bpr: Bayesian personalized ranking from implicit feedback. In: UAI. pp. 452–461 (2009)
28. Rikhardsson, P., Yigitbasioglu, O.: Business intelligence & analytics in management accounting research: Status and future focus. International Journal of Accounting Information Systems **29**, 37–58 (2018)
29. Roetzel, P.G.: Information overload in the information age: a review of the literature from business administration, business psychology, and related disciplines with a bibliometric approach and framework development. Business research **12**(2), 479–522 (2019)
30. Ruder, S.: An overview of gradient descent optimization algorithms. arXiv preprint arXiv:1609.04747 (2016)
31. Sanh, V., Debut, L., Chaumond, J., Wolf, T.: Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. arXiv preprint arXiv:1910.01108 (2019)
32. Shalaby, W., Zadrozny, W.: Patent retrieval: a literature review. Knowledge and Information Systems pp. 1–30 (2019)
33. Shalf, J.: The future of computing beyond Moore's Law. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences **378**(2166) (2020)
34. Sitarz, R., Kraslawski, A.: Application of semantic and lexical analysis to technology forecasting by trend analysis-thematic clusters in separation processes. In: Computer aided chemical engineering, vol. 30, pp. 437–441. Elsevier (2012)
35. Tang, F., Liu, Y.: Applying semantic web into technology forecasting in enterprises. In: International Conference on Service Operations and Logistics, and Informatics. vol. 1, pp. 135–138. IEEE (2008)
36. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. NeurIPS **30**, 5998–6008 (2017)
37. Wang, Y., Rastegar-Mojarad, M., Komandur-Elayavilli, R., Liu, H.: Leveraging word embeddings and medical entity extraction for biomedical dataset retrieval using unstructured texts. Database **2017** (2017)