

## Article

# Building Knowledge Graphs and Recommender Systems for Suggesting Reskilling and Upskilling Options from the Web

Albert Weichselbraun <sup>1,2,\*</sup> , Roger Waldvogel <sup>1</sup> , Andreas Fraefel <sup>1</sup> , Alexander van Schie <sup>1</sup>   
and Philipp Kuntschik <sup>1</sup>

<sup>1</sup> Institute for Information Research, University of Applied Sciences of the Grisons, 7000 Chur, Switzerland

<sup>2</sup> webLyzard Technology, 1090 Vienna, Austria

\* Correspondence: albert.weichselbraun@fhgr.ch; Tel.: +41-81-286-3727

**Abstract:** As advances in science and technology, crisis, and increased competition impact labor markets, reskilling and upskilling programs emerged to mitigate their effects. Since information on continuing education is highly distributed across websites, choosing career paths and suitable upskilling options is currently considered a challenging and cumbersome task. This article, therefore, introduces a method for building a comprehensive knowledge graph from the education providers' Web pages. We collect educational programs from 488 providers and leverage entity recognition and entity linking methods in conjunction with contextualization to extract knowledge on entities such as prerequisites, skills, learning objectives, and course content. Slot filling then integrates these entities into an extensive knowledge graph that contains close to 74,000 nodes and over 734,000 edges. A recommender system leverages the created graph, and background knowledge on occupations to provide a career path and upskilling suggestions. Finally, we evaluate the knowledge extraction approach on the CareerCoach 2022 gold standard and draw upon domain experts for judging the career paths and upskilling suggestions provided by the recommender system.

**Keywords:** knowledge extraction; knowledge base population; entity recognition; entity classification; entity linking; slot filling; knowledge graph; recommender system



**Citation:** Weichselbraun, A.; Waldvogel, R.; Fraefel, A.; van Schie, A.; Kuntschik, P. Building Knowledge Graphs and Recommender Systems for Suggesting Reskilling and Upskilling Options from the Web. *Information* **2022**, *13*, 510. <https://doi.org/10.3390/info13110510>

Academic Editor: Pierpaolo Basile

Received: 7 September 2022

Accepted: 17 October 2022

Published: 25 October 2022

**Publisher's Note:** MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Copyright:** © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

## 1. Introduction

Over the last decade, vast amounts of open knowledge have been published as public knowledge graphs. These graphs allow for the integration of knowledge from heterogeneous sources (e.g., via links to other datasets), and are fully interoperable with the Semantic Web technology stack defining standards for storing, querying, and even reasoning upon these graphs. The increasing popularity of knowledge graphs also led to considerable growth of available datasets. As of May 2020, the Linked Open Data (LOD) initiative counted 1255 public knowledge graphs compared to only twelve in 2007 (<https://www.lod-cloud.net>, accessed on 17 October 2022). Public knowledge graphs have shaped research in many domains, where tools and methods that draw upon graphs such as DBpedia Spotlight [1] and Recognyze lite [2] have emerged. These developments have been complemented by growing business interest in maintaining private graphs of enterprise knowledge which are often tightly integrated with public knowledge graphs [3].

The research presented in this paper aims at developing automatic knowledge graph construction methods and recommender systems for the human resources and continuing education domain. Both employers and employees are seriously affected by labor market disruptions caused by shifting in-demand skills due to the pace of technology adoption, automation, and crisis. Companies surveyed for the World Economic Forum estimate that by 2025 6.4% of their workforce may be displaced by shifts in the division of labor between humans and machines. In turn, new professions are expected to grow from 7.8% to 13.5% as roles better adapted to these markets emerge [4]. Continuing education, particularly

reskilling and upskilling provides feasible mitigation strategies by enabling employees to obtain in-demand skills which increase their chances in the labor market. Companies estimate that around 40% of their workforce will require reskilling of six months or less, plan to offer reskilling and upskilling to 70% of their employees by 2025, and expect employees to pick up new skills on the job [4].

Despite a growing pressure to attend continuing education programs, choosing feasible programs is currently a challenging and cumbersome task. Identifying suitable career paths and short-term upskilling opportunities requires insights on relevant occupations and continuing education programs, information that is distributed across multiple websites and providers.

This article addresses these challenges by developing automatic knowledge graph construction components that populate and update a continuing education knowledge graph based on the websites of 488 different education providers. The created graph enables semantic search and supports sophisticated queries (e.g., for education that provides specific skill sets) over a SPARQL interface. It, therefore, provides a single point of entry for searches that aim at locating suitable continuing education programs and publishes a query interface that offers support for more sophisticated applications that build upon this graph. The created knowledge graph and the presented methods have been developed within the CareerCoach project (<https://www.fhgr.ch/careercoach>, accessed on 17 October 2022), and are already used in industry to power data analytics applications and a platform that provides semantic search for reskilling and upskilling options. Furthermore, we develop recommender systems that draw upon the knowledge graph to help users in selecting suitable career paths and education programs.

The main contributions of this paper are

1. the adaptation of knowledge extraction methods to the human resources and continuing education domain;
2. applying these methods to a complex industry-driven setting that requires robust methods capable of operating on content retrieved from a multitude of education providers for automatic knowledge graph construction;
3. the creation of a continuing education knowledge graph that comprises 73,969 nodes and 734,447 edges;
4. developing a knowledge-driven recommender system that draws upon this background knowledge to support users in identifying useful reskilling and upskilling options;
5. evaluating the created systems based on a slot-filling benchmark and domain expert assessments.

The rest of the paper is organized as follows: Section 2 discusses related work. Section 3 presents an overview of the developed knowledge graph construction and recommender systems, introduces the methods used by the automatic knowledge graph construction process (Section 3.1), and then describes the recommender systems used for suggesting career paths and the corresponding continuing educations (Section 3.2). Afterward, Section 4 provides a comprehensive evaluation of these systems, which is followed by a discussion in Section 5. The paper concludes with Section 6, which summarizes the presented work, and provides insights into planned future research.

## 2. Related Work

Industry-scale knowledge graphs such as the Google Knowledge Graph, Microsoft's Bing Knowledge Graph, and the knowledge graphs deployed by Facebook, eBay, and IBM Watson have gained in importance in recent years [5]. In contrast to Linked Open Data sources such as DBpedia [6] and Wikidata [7], these knowledge graphs are proprietary and tailored towards specific use cases within the companies' product portfolios. Nevertheless, they still comprise a considerable number of concepts. The LinkedIn economic graph (<https://economicgraph.linkedin.com/the-future-of-work>, accessed on 17 October

2022), for instance, contains information on over 774 million members, 50 million companies, 36,000 skills and 90,000 schools.

Creating and maintaining such comprehensive knowledge graphs requires significant resources, which has accelerated research in automated knowledge extraction and knowledge base population methods.

### 2.1. Knowledge Extraction and Knowledge Base Population

Knowledge base population, for instance, applies knowledge extraction techniques towards discovering facts in unstructured textual resources that are then integrated into a knowledge base or knowledge graph.

DBpedia, which is one of the most popular and influential knowledge graphs, is constructed by extracting facts from Wikipedia Web pages and storing them in the form of (subject, predicate, object) triples [6]. While DBpedia is solely populated based on knowledge extracted from Wikipedia, many other approaches operate on considerably more heterogeneous document collections as reflected in the composition of evaluation datasets which cover News articles [8], question answering [9–11], and even general Web documents [12,13].

Lin et al. [14] distinguish between two approaches towards knowledge base population: (i) methods that draw upon entity linking and slot filling to extract information on predefined slots (e.g., a person's occupation, age, etc.) and (ii) open information extraction (Open IE) that identifies arbitrary (subject, predicate, object) triples but also requires a subsequent linking step for disambiguation and integration into a knowledge base. Slot filling, relation linking, and entity linking are, therefore, important subtasks required for the knowledge base population.

Entity recognition (ER) identifies mentions of entities in text documents, entity classification (EC) also determines the entity type (e.g., person, organization, skill, etc.), and entity linking (EL) links mentions to a knowledge graph such as DBpedia and Wikidata.

Rule-based approaches, machine learning, and deep learning have been successfully used for ER and EC [15] with the latter group providing state-of-the-art results. Another interesting development has been the introduction of approaches such as SpanNER [16] that consider ER as a token span detection problem. These methods tend to perform particularly well on out-of-vocabulary words, while token labeling provides better results for use cases with long entities and low label consistency [16].

Knowledge graph disambiguation and neural models are currently among the most effective approaches toward EL. Systems such as AIDA [17], HITS [18], Babelify [19], AGDIS-TIS [20] and its multilingual versions MAG [21] have been among the top performers [2]. Recently, neural models such as NCEL [22], Dynamic Graph Convolutional Networks [23], and models using minimal context information [24] have gained in importance and show superior performance in many settings.

### 2.2. Slot Filling

Slot filling combines multiple pieces of information (e.g., skills, learning outcomes, etc.) into a single knowledge base entry. When applied to open-world scenarios, slot filling is very challenging, as demonstrated by competitions such as the TAC 2017 Cold Start Slot Filling Task in which even the winning systems only obtained F-measures below 20% [25]. Restricting the task to a single language yields considerably better results. At the TAC KBP 2013 English Slot Filling evaluation, for instance, the top-ranked system developed by Roth et al. [26] yielded an F1 score of 37.3%. If applied to a single domain, even complex temporal tasks such as temporal slot filling yield good results with F1 scores of over 76.5% [27]. This assessment is also confirmed by Ritze et al. [28] who demonstrate that slot filling is well suited for augmenting knowledge bases.

Siddique et al. [29] present a zero-shot slot-filling model capable of adapting effectively to new domains by using pre-trained natural language processing models that provide domain-independent word representations, propagating insights between tasks, using a

generalized similarity function, and contextualizing word embeddings. In evaluations on the SNIPS [30], ATIS [31], MultiWOZ [32] and SGD [33] datasets their system outperformed state-of-the-art slot filling systems such as Coach [34], RZS[35] and CT [36].

### 2.3. Open Knowledge Extraction

At the European Semantic Web Conference (ESWC), the Open Knowledge Extraction (OKE) challenges explored the use of Open Knowledge Extraction in the context of the Semantic Web. The challenges comprised entity recognition, entity linking, and entity typing as well as class induction and relation extraction tasks. Cerezo-Costas and Martin-Vicente [37] introduced a deep learning engine that extracts relations between entity pairs and applies expert knowledge to filter illogical relations. Chabchoub et al. [38], who won the 2016 OKE challenge Task 1, combined entity recognition, and heuristics for merging and filtering the detected mentions with a disambiguation technique that uses previous entities for disambiguation.

Recently, end-to-end systems such as KBPearl [14] and Falcon [39] have gained in importance. KBPearl, for instance, combines open information extraction, entity linking, and relation linking tasks [14]. Falcon also jointly performs relation and entity linking tasks, linking entities and relations to DBpedia and Wikidata [39].

### 2.4. Recommender Systems

Recommender systems aim at meeting the user's personalized interests and overcoming information overload by suggesting relevant items [40]. They are mainly categorized into collaborative filtering, content-based recommender systems, and hybrid methods combining these approaches [41]. Standard methods used in knowledge graph aware recommender systems include graph-based algorithms [42–44], embedding-based approaches [45–48] and hybrid methods [49–51].

The research presented within this paper aims at creating a proprietary continuing education knowledge graph, which will complement an existing occupation knowledge base, job vacancies databases, and a repository of jobseeker profiles. The knowledge graph will formalize knowledge on reskilling and upskilling options, enable the creation of knowledge-driven recommender systems, and support customers in quickly locating suitable career paths and educational offers. The introduced knowledge extraction and knowledge graph construction techniques are instrumental in populating the continuing education knowledge graph and improving the coverage of relevant areas in the existing knowledge base, particularly the sections on skills and degrees.

## 3. Method

Figure 1 provides an overview of the methods and components introduced within this paper. Section 3.1 describes the methods used to create an automatic knowledge graph construction pipeline that analyzes Web pages published by education providers to build and update a continuing education knowledge graph. Afterward, Section 3.2 introduces knowledge-driven recommender systems that leverage the created knowledge graph to suggest beneficial reskilling and upskilling paths, and continuing educations facilitating them.

The knowledge graph construction methods and the recommender systems require external knowledge for disambiguation, entity linking, and computing recommendations. The project's industry partner ( $\times 28$ ), therefore, contributed to the  $\times 28$  occupation knowledge base which formalizes domain knowledge on work-related areas such as skills, education, occupations, topics, and industries. The knowledge base comprises 42 database tables and also contains relations between instances, and links to other, publicly available schemas such as the Standard Occupational Classification (SOC) (<https://www.bls.gov/soc/>, accessed on 17 October 2022) and multiple industry directories. This background knowledge has been collected, formalized, and updated by  $\times 28$ 's domain experts, represents years of development, and is utilized together with domain-specific business logic and constraints

in the knowledge extraction processes and the developed recommender system, either directly or after transformation into a knowledge graph.

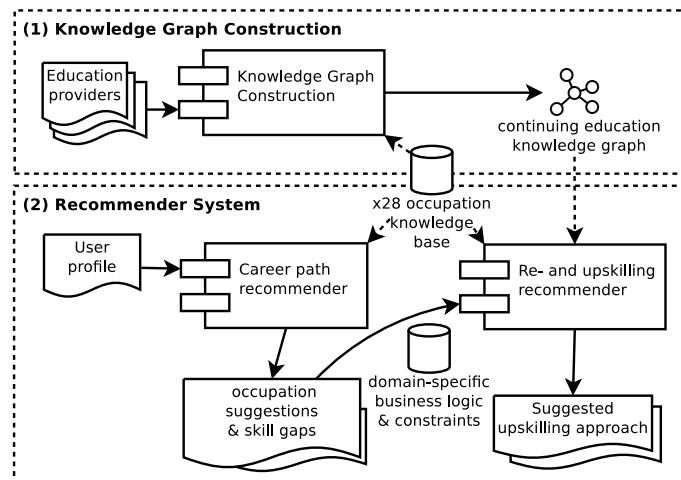


Figure 1. Use of knowledge graphs and ontologies for the suggestion of reskilling and upskilling paths.

Deploying knowledge graphs provides us with advantages in terms of interoperability (i.e., we are no longer bound to proprietary database schemas), facilitates data integration (e.g., via links to other knowledge graphs and datasets), and enables using the Semantic Web stack which provides tools and standards for storing, querying and reasoning upon the knowledge graph. These standards also facilitate interoperability with a number of powerful third-party tools. The current version of the recommender system, for example, only considers hierarchical relations between skills (Section 3.2.1). Due to the use of RDF-based knowledge graphs, more complex reasoning could be easily added by drawing upon reasoners such as DIG 2.0 [52] and Pellet (<https://github.com/stardog-union/pellet>, accessed on 17 October 2022). In addition, the presented methods for page segmentation (Section 3.1.1) and entity linking (Section 3.1.2) process knowledge graphs which makes them much more versatile and easier to adapt to new application domains.

### 3.1. Knowledge Graph Construction from the Web

Figure 2 illustrates the developed knowledge extraction and knowledge base population process, which currently operates on Web pages from 488 different education providers. Extracting structured knowledge from websites, and its integration into knowledge graphs, requires combining content extraction with context-aware knowledge extraction methods.

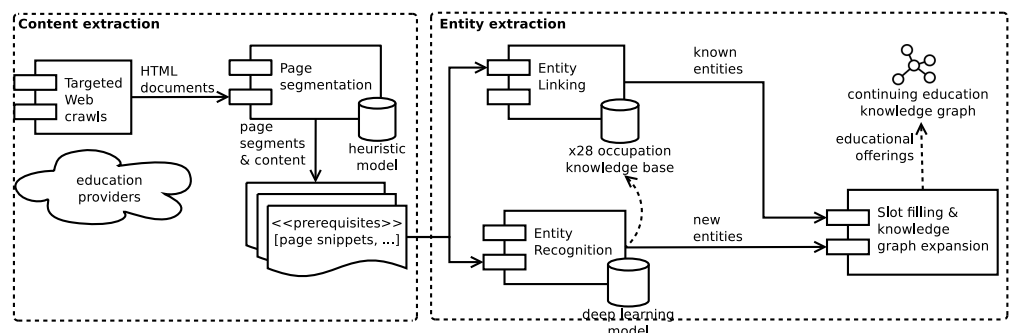



Figure 2. Overview of the automatic knowledge graph construction process.

The system’s knowledge graph population pipeline expands the continuing education knowledge graph by analyzing the Web pages  $doc_i$  of educational offerings  $i$ . A page segmentation and classification component extracts relevant page segments  $seg_i^{type} \in doc_i$  with  $type \in \{target\ group, prerequisite, learning\ objective, course\ content, certificates\ \&\ degree\}$  from

these pages. Afterward, entity linking identifies known entities  $e_{ij}^{known}$  such as skills and occupation within the segments and links them to the  $\times 28$  occupation knowledge base. We complement entity linking with entity recognition, which is capable of identifying entities  $e_{ij}^{new}$  that are not yet available in the knowledge base. The extracted new entities serve as candidate entities for extending the  $\times 28$  occupation knowledge base, and considerably improve the coverage of the entity extraction process.

Finally, the slot filling component fills for each educational offering  $i$  the slots outlined in Table 1 by contextualizing the extracted entities  $e_{ij} = e_{ij}^{known} \cup e_{ij}^{new}$  with the information on the page segment type  $seg_i^{type}$  from which they have been extracted. In the real-world example shown in Figure 3, the slot value with surface form “Programmierung” (cc:25002374 (identifier for the knowledge base concept “programming”)) extracted from the page segment “Lehrinhalte”, for example, fills the course content slot (cc:hasCourseContent), while the same entity extracted from the segment “Voraussetzungen” (so:programPrerequisites) would be considered a prerequisite for visiting the course.



**Figure 3.** Snippet of an annotated example course description taken from <https://www.sae.edu>, accessed on 17 October 2022. Light blue highlighting indicates identified entities, and the red border outlines the page segment to which they have been assigned.

**Table 1.** Target slot, valid entity types and cardinalities of course entities.

Slot	(Entity) Type (min, max)	Cardinality
title	title extracted from the page metadata	(1, 1)
school	school	(1, 1)
target group	degree, education, occupation, position, industry, topic	(0, *)
prerequisite	degree, education, occupation, position, skill, topic	(0, *)
learning objective	occupation, skill, topic	(1, *)
course content	skill, topic	(1, *)
certificates	degree, education	(0, *)

After the completion of the slot-filling process, the system integrates the information on each educational offering into the knowledge graph by generating the corresponding triples (e.g.,  $\langle \text{https://sae.edu/course01} \rangle$  cc:hasCourseContent cc:25002374.) that are then serialized in the RDF (Resource Description Framework) format.

The following sections discuss the applied methods in greater detail.

### 3.1.1. Page Segmentation

Most websites present educational offerings structured in sections such as course prerequisites and learning objectives. Extracting the content within these sections and correctly labeling it is key to ensuring that subsequent entity linking and slot-filling steps can correctly link and contextualize (e.g., distinguish between skills that are prerequisites versus learning outcomes) entities.

Page segmentation partitions the text in sections, identifies section titles and clusters, and finally assigns each section to a label defined in a task-specific classification schema (i.e., course title, target group, prerequisites, learning objectives, course content and degrees & certificates).

The text partitioning task splits the HTML documents into segments, whenever one of the following elements  $m = \{div, p, li, td, th, dt, dd, summary, legend, h1, h2, h3, h4, h5, h6\}$  encloses text that does not contain any other separator element  $m$ .

Afterwards, the algorithm tries to identify cluster titles based on the HTML title elements  $n = \{h1, h2, h3, h4, h5, h6\}$ . In practice, not all titles use a title element as specified in the HTML standard, in which case the following fallback heuristic may be applied for identifying non-standard title segments: (i) The segment’s text must not contain more than three words; (ii) At least one word of the candidate text must appear in a domain-specific list of commonly used title terms such as ‘prerequisite’, ‘content’, and ‘degree’.


Once title elements have been identified, they are used to determine the text cluster by merging the segments below the titles (Figure 4).

**Example HTML snippet**

```
<h2>Prerequisites</h2>
<div>Bachelor study</div>
<div>Programming skills</div>

<h2>Degree</h2>
<div>MSc in Computer Science</div>
```

**Text representation (segments and clusters)**




**Figure 4.** Text divided into segments based on the example HTML snippet. Violet fields indicate cluster titles, blue boxes the assigned content, and dashed lines cluster boundaries.

The final content classification step assigns clusters to the corresponding page segment  $seg_i^{type} \in doc_i$  with  $type \in \{course\ title, target\ group, prerequisites, learning\ objectives, course\ content, degrees\ \&\ certificates\}$  by comparing title terms with patterns recorded within the classification ontology. If the title terms do not match any known patterns, the cluster is classified as *unknown*.

3.1.2. Entity Linking

The system uses a graph-based entity linking method [53] that draws upon the project’s skill and education database. Based on the information needs of the developed knowledge-driven recommender system (Section 3.2), we customized the component to identify the following entity types used within the  $\times 28$  occupation knowledge base: (i) education, (ii) function (i.e., occupation and position), (iii) skill and (iv) topic.

Similar to other machine learning components, we differentiate between data preparation, training, and evaluation (Figure 5). The data preparation step facilitates the Protege Ontop (<https://github.com/ontop/ontop>, accessed on 17 October 2022) framework to transform the  $\times 28$  occupation knowledge base into a Linked Data representation. The resulting mapping has a significant impact on the further process since it defines the relations and data points the system can utilize for training the entity linking component.



**Figure 5.** Steps required to prepare a knowledge base for use in the entity linking (EL) component: data preparation transforms the data into a Linked Data repository, training mines the data repository to create a serializable EL profile, and evaluation uses the profile to annotate new and unknown documents.

The training step then draws upon the use of case-specific SPARQL queries to mine relevant entities and context information from the created Linked Data repository. For further processing, all concept labels mined from the knowledge base are either marked as qualifying names that fully identify entities (such as “hedge fund manager”), ambiguous names that are not sufficient for identifying an entity (e.g., “wolf of wall street”), or context information. Automatic graph mining extracts relations between entities which are then utilized by the graph-based disambiguation algorithm. The EL Profile Builder further applies multiple pre-processors and analyzers to the query results, which create artificial name variations such as plurals, possessive forms, and abbreviations to maximize the EL profile’s coverage. Afterward, analyzers determine the relevancy of the surface forms (queried and generated) within the profile, by classifying them into unambiguous surface forms, ambiguous surface forms, and context terms. The EL Profile Builder concludes training by serializing the model to a binary EL profile that contains all information required for the EL process.

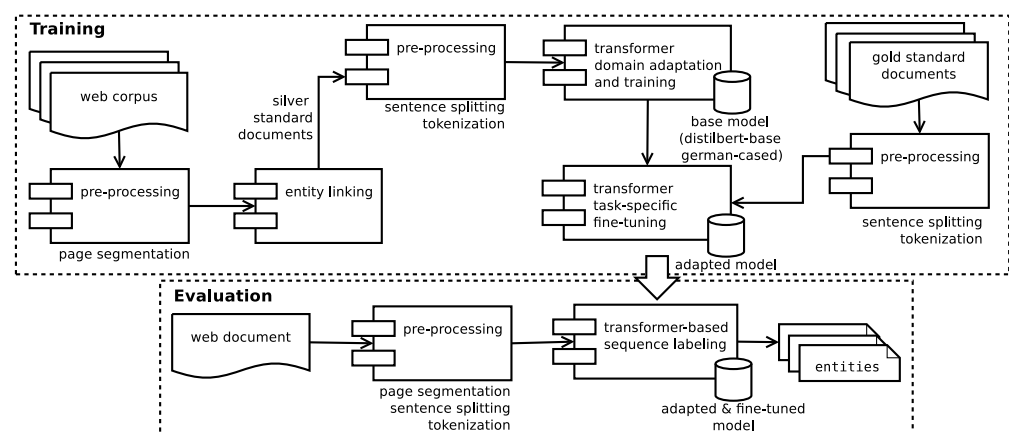
The EL Web Service, used in the evaluation step, draws upon a graph-based disambiguation algorithm [53] that operates on the serialized EL profiles. The service identifies entities within the provided page segments and links them to the corresponding nodes within the occupation knowledge base.

### 3.1.3. Entity recognition and Entity Classification

An evaluation of the entity linking component revealed that the deployed knowledge graph still misses approximately 45% of all relevant entities. We mitigate this issue by using entity recognition and entity classification as fallback methods for identifying entities that have not yet been included in the industry partner’s occupation knowledge base. This strategy does improve recall and provides candidate concepts for inclusion into the knowledge base, therefore, enhancing its coverage over time.

We model entity recognition and entity classification as a token classification problem, where a deep learning model takes a sequence of tokens (e.g., a sentence) and then provides labels for each of them. The entity classification component draws upon the *distilbert-base-german-cased* model provided by the popular transformers library [54].

A domain adaptation and initial training step use a domain corpus of 28,000 documents that has been enriched with silver standard annotations obtained from the previously described EL component. To improve the model’s capabilities of capturing new entities, we performed a fine-tuning step that draws upon the gold standard documents to generate the final model. A five-fold cross-evaluation procedure ensures that no training documents are used within the evaluation. Table 2 summarizes the parameters used for training and evaluation (all experiments have been performed on version 4.12.5 of the transformers library), and Figure 6 outlines the model structure and the training process.



**Figure 6.** Transformer model used for entity recognition and entity classification.



**Table 2.** Parameters used by the entity recognition and entity classification transformer model.

Parameter	Value
Solver (learning rate)	Adam ( $5 \times 10^{-5}$ )
Activation	Gaussian Error Linear Unit (GELU)
Base model	distilbert-base-german-cased
Attention dropout	0.1
Dimension	768
Dropout	0.1
Hidden layer dimensions	3072
Initializer range	0.02
Max position embeddings	512
N heads (N layers)	12 (6)
Qa dropout	0.1
Seq classification dropout	0.2

### 3.1.4. Knowledge Graph Expansion

The knowledge graph expansion constructs the continuing education knowledge graph and provides candidate concepts (e.g., new skills) for inclusion into the  $\times 28$  occupation knowledge base. It combines the entities obtained from the entity linking and entity recognition processes with context knowledge yielded by the page segmentation heuristic, to fill the slot values required for constructing the continuing education knowledge graph. New entities that have been discovered by the entity recognition and classification component and are, therefore, not yet in the occupation knowledge base are assigned temporary identifiers that can be used for linking them to the knowledge base at a later stage. Afterward, these new entities are forwarded to domain experts for review. If the experts decide to include an entity in the knowledge base, it is assigned a permanent identifier that is also propagated to the continuing education knowledge graph.

Table 1 provides an overview of the used slots, matching entity types, and the corresponding cardinalities. The component retrieves the title slot from the Web page metadata and fills all other slots by contextualizing the entities extracted in the previous steps.

Afterward, it draws upon the following RDF namespaces and the mapping outlined in Table 3 to serialize the extracted knowledge in an RDF graph:

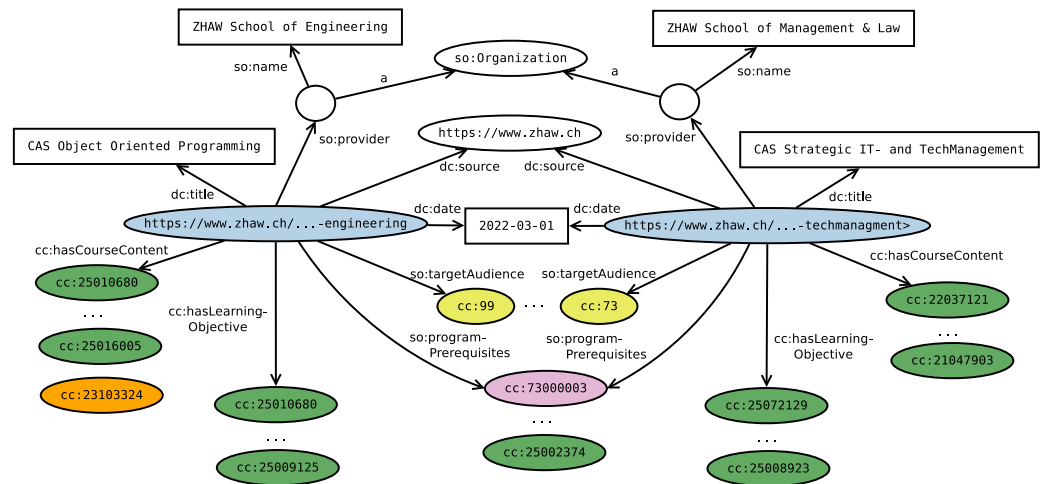
- cc: project-specific CareerCoach namespace that is used for custom vocabulary (e.g., course content, learning objectives) and for referring to entities within the industry partner's knowledge base
- dc: Dublin Core namespace used for the title, source, and date properties
- skos: Simple Knowledge Organization System namespace to indicate entities that haven't been assigned to a slot with the skos:related property
- so: Schema.org namespace to describe educational programs (e.g., credits and degrees awarded, program prerequisites, and target audiences) and the organizations offering these programs

**Table 3.** Vocabulary used for translating slot filling results to RDF statements.

Slot	RDF Property
title	dc:title
school	so:provider
target group	so:targetAudience
prerequisite	so:programPrerequisites
learning objectives	cc:hasLearningObjective
course content	cc:hasCourseContent
certificates	so:educationalCredentialAwarded

Figure 7 shows a small fraction of the created graph which comprises two courses offered by the Zurich University of Applied Sciences. The visualization replaces rdf:type

statements with color coding, and only shows a fraction of the assigned skills, industries, occupations, and degrees.



**Figure 7.** Visualization of the extracted knowledge graph for two continuing education offers provided by the Zurich University of Applied Sciences. The visualization outlines a section of the assigned properties and replaces `rdf:type` properties with color coding. Blue nodes indicate education offerings, green nodes skills, orange nodes industries, yellow nodes occupations, and pink nodes degrees.

### 3.2. Knowledge-Driven Recommender System

The developed recommender system uses the  $\times 28$  occupation knowledge base and the created continuing education knowledge graph as knowledge sources. Furthermore, it allows the integration of additional filtering and ranking criteria, including real-time data on job vacancies to model the demand side of the job market.

#### 3.2.1. Background Knowledge

The recommender uses the following functions to query background knowledge from the occupation knowledge base and the continuing education graph:

- $S(target)$  queries the occupation knowledge base for all skills  $s_i \in S(target)$  required for the given target occupation; the query also considers hierarchical relations between skills (e.g., the skill “Java programming” will automatically imply “Programming”);
- $f(s_i)$  returns the number of occupations that require skill  $s_i$  from the occupation knowledge base; and
- $b(s_i, education)$  uses the continuing education graph to determine whether the given education provides skill  $s_i$ , considering hierarchical relations between skills specified in the occupation knowledge base.

The continuing knowledge graph provides additional information such as an education’s target audience and prerequisites that are not yet used by the recommender system.

#### 3.2.2. Business Logic and Constraints

Research conducted by experts from the World Economic Forum (WEF) [55,56] list the following criteria for evaluating migration paths between occupations:

1. job similarity, which considers work activities, necessary knowledge (e.g., completed educations), skills (i.e., cross-functional and specialized skills), abilities (e.g., physical and cognitive capabilities), and expertise (education, years of work experience in the job or job family)
2. similar job zones (i.e., expected level of education)
3. stable long-term prospects (i.e., demand for the occupation is not declining)
4. wage continuity or increase

In contrast to CareerCoach, the WEF's recommendations only propose reskilling paths but do not consider the problem of identifying suitable reskilling and upskilling opportunities that help in implementing the proposed changes.

Our recommender system draws upon a flexible knowledge-driven approach that allows considering use-case-specific ranking criteria and constraints in the provided suggestions. The WEF criteria have been particularly useful in designing the business logic that guides suggestions of suitable career paths:

1. prefer similar jobs over less similar ones, since they require lower reskilling or upskilling efforts
2. the suggested jobs should expect a similar level of education (i.e., do not suggest paths that would require significant additional education or would devalue past educations)
3. provides optional filters and ranking rules that consider a user's preferences regarding wage continuity or increase, expected long-term prospects, and geography (i.e., availability of suitable positions in a particular region)

### 3.2.3. Knowledge-Driven Occupation Recommendations

In our initial experiments, we provided the education recommender with the user's existing skills  $S_{user}$ , and the ones required by the target occupation  $S(target)$  to compute the corresponding skill gap  $S_{gap}(target)$  (i.e., the list of required skills that are missing in the user's profile).

$$S_{gap}(target) = S(target) \setminus S_{user} \quad (1)$$

Our experiments revealed the relative size of the skill gap  $S_{gap}^{rel}(target)$  as a good ranking criterion, since it provides an estimation of how related the two occupations are.

$$S_{gap}^{rel}(target) = \frac{|S_{gap}(target)|}{|S(target)|} \quad (2)$$

In addition, the recommender system supports filtering and re-ranking of the results based on the criteria outlined in Section 3.2.2. Filters enforcing similar job zones have been implemented, although the flexibility of the Swiss labor market limits their usefulness, since it allows different educational levels for many occupations. A programmer, for example, might have successfully completed an apprenticeship, a bachelor, a master, or even a doctoral degree. Consequently, it is often not feasible to apply suitable restrictions, if only a target occupation without additional context (e.g., from a job announcement) is provided.

### 3.2.4. Knowledge-Driven Continuing Education Recommendations

The most straightforward approach towards providing continuing education recommendations would be extending the method presented in the previous section to reskilling and upskilling recommendations by ranking education based on their capability to close the skill gap  $S_{gap}(target)$  between the user's skills and the ones required by the target occupation.

Nevertheless, initial experiments quickly revealed this strategy as ineffective, since it does not sufficiently distinguish between common cross-industry skills (e.g., project management), and skills that are specific to the target occupation (e.g., Java programming).

Therefore, a new approach, that computes skill weights by drawing upon the background knowledge available in the  $\times 28$  occupation knowledge base, has been selected.

As outlined in Equations (3) and (4), the proposed strategy first determines the frequency ( $F_{skill}^{min}(target)$ ) of the most specific skill for a certain target occupation and then computes a weight  $w(s_i, target) \in [0, 1]$ , which indicates the specificity of each skill  $s_i$  for that particular occupation. As a result, the most job-specific skills obtain a weight  $w(s_i, target)$  close to one, while less specific skills yield weights closer to zero.

$$F_{skill}^{min}(target) = \min_{s_i \in S(target)} f(s_i) \quad (3)$$

$$w(s_i, target) = F_{skill}^{min}(target) / f(s_i) \quad (4)$$

Further qualitative analysis of real-world use cases revealed that cross-industry skills are frequently essential for working in the target occupation. Consequently, it might be counterproductive to ignore them completely. We, therefore, score recommendations by combining the education's contribution towards closing the job-specific skill gap with its total coverage of skills:

$$score(S_{gap}(target), education) = \alpha \cdot \frac{\sum_{s_i \in S_{gap}(target)} b(s_i, education) \cdot w(s_i, target)}{\sum_{s_i \in S_{gap}(target)} w(s_i, target)} \quad (5)$$

$$+ (1 - \alpha) \cdot \frac{\sum_{s_i \in S_{gap}(target)} b(s_i, education)}{|S_{gap}(target)|}$$

$$b(s_i, education) = \begin{cases} 1 & \text{if } s_i \in S(education) \\ 0 & \text{otherwise.} \end{cases} \quad (6)$$

The factor  $b(s_i, education)$  indicates whether a skill  $s_i$  is provided by the suggested education. The first term of Equation (5) estimates the continuing education's relevance to the target occupation and is weighted with  $\alpha \in [0, 1]$ . The second term of the equation, in contrast, corresponds to the completeness and is weighted with the factor  $(1 - \alpha)$ . The computed score ( $score(S_{gap}(target), education)$ ) is then used for ranking continuous education options.

#### 4. Evaluation

The evaluation focuses on the following two objectives: (i) quantifying the performance of the introduced knowledge graph population component for the automatic population of the continuing education knowledge graph, and (ii) illustrating the capabilities of the recommender system that draws upon this graph.

We deploy both, evaluation datasets and expert assessments to evaluate the system's performance:

- The evaluation of the knowledge extraction and knowledge graph population components uses the CareerCoach 2022 gold standard which has been introduced at the 27th International Conference on Natural Language & Information Systems (NLDB 2022) [57] (Section 4.1);
- The career path recommender is evaluated based on a gold standard of expert recommendations (Section 4.2). Afterward, experts assess the usefulness of the provided continuing education recommendations (Section 4.3).

Tables 4 and 5 provide an overview of the components, metrics and evaluation objectives. The evaluation of the knowledge graph population approach assesses the methods used for identifying entities to be integrated into the knowledge graph (i.e., entity linking, entity classification, and entity recognition), and the contextualization of the extracted entities (i.e., page segment recognition and page segment classification) which is required for distinguishing between an education's prerequisites and outcomes. Benchmarking the slot-filling component compares the extracted slots to the gold standard data provided in the CareerCoach 2022 dataset, yielding precision (P), recall (R), and F1 measures for the overall slot-filling process.

**Table 4.** Evaluation components, metrics, and objectives for the knowledge graph population population method (Section 4.1).

Component	Evaluation Metrics	Objective
Page segment recognition	P, R, F1	evaluate content
Page segment classification	P, R, F1	extraction
Entity Recognition	P, R, F1	evaluate entity
Entity Classification	P, R, F1	extraction
Entity Linking	P, R, F1	
Slot filling	P, R, F1	evaluate overall slot filling process

The evaluation of the recommender system compares its ranking of career paths with a gold standard ranking provided by domain experts. Precision (P@3) and mean average ranking precision (MAP@3) for the top-three recommendations provide insights into the alignment between expert assessments and system recommendations. The same group of experts also evaluates the usefulness of the provided continuing education recommendations.

**Table 5.** Evaluation components, metrics, and objectives for the recommender system (Sections 4.2 and 4.3).

Component	Evaluation Metrics	Objective
career path recommender	P@3, MAP@3	evaluate feasibility and beneficial-ness of the suggested career paths
continuing education recommender	P <sup>b</sup> @3, P <sup>s</sup> @3	evaluate usefulness of the suggested educations

#### 4.1. Knowledge Graph Population

The evaluation of the knowledge graph population pipeline relies upon a domain-specific slot-filling gold standard that is described in Section 4.1.1. Afterwards, Section 4.1.2 discusses the evaluation of the content extraction subtasks, Section 4.1.3 the entity extraction subtasks, and Section 4.1.4 the slot filling task required for creating the knowledge graph (compare Figure 2).

##### 4.1.1. Gold Standard

Our experiments draw upon the publicly available CareerCoach 2022 gold standard dataset (<https://github.com/fhgr/careercoach2022>, accessed on 17 October 2022) [57] which comprises (i) a document partition for evaluating knowledge extraction and classification tasks (169 documents), and (ii) a second partition, which contains annotations for benchmarking entity extraction and slot filling (75 documents). In total, the dataset contains over 3800 annotations and 169 documents, which have been obtained from 89 different education providers.

##### 4.1.2. Content Extraction

Content extraction identifies and classifies text segments relevant to the slot-filling tasks. We distinguish between,

1. *T1: page segment recognition*—locates page segments within HTML pages *doc* and extracts the text string  $s_i \in doc$  from these segments.
2. *T2: page segment classification*—assigns each extracted text segment  $s_i$  to a class  $c_i \in \mathcal{C}$ . The page segment classification considers the classes 'target\_groups', 'prerequisites', 'learning\_objectives', 'course\_contents', and 'degrees & certificates'.

We evaluate the page segment recognition task (T1) by comparing the tokens in the extracted page segments  $t_i \in s_i$  with the tokens in the gold standard segments  $t_g \in s_g$ , computing precision ( $P$ ), recall ( $R$ ) and F1 measure as follows:

$$P = \frac{|t_g \cap t_i|}{t_i} \quad (7)$$

$$R = \frac{|t_g \cap t_i|}{t_g} \quad (8)$$

$$F1 = 2 \cdot \frac{P \cdot R}{P + R} \quad (9)$$

The evaluation of the page segment classification relies on the same metric, but requires the classes of the gold standard page segment  $t_g^c$  and of the extracted segment  $t_i^s$  to match.

#### 4.1.3. Entity Extraction

The entity extraction tasks aim at identifying mentions of entities of type  $t_i \in \mathcal{T}_i$  within the extracted page segments. The corpus contains annotations of the following entity types: 'skill', 'occupation', 'topic', 'position', 'school', 'industry', 'education', 'degree'.

1. *T3: entity recognition*—locates mentions  $m_i$  of entities within text segments.
2. *T4: entity classification*—assigns each mention  $m_i$  to the corresponding entity type  $t_i \in \mathcal{T}_i$ .
3. *T5: entity linking*—links mentions  $m_i$  to the appropriate entity  $e_i$  in the knowledge graph  $KG$ . Entities that are not yet available in the knowledge graph are handled as NIL entities (i.e., they are assigned a temporary identifier that is unique for all mentions which refer to the same entity).

The evaluation of the entity recognition (T3), entity classification (T4), and entity linking (T5) tasks use the precision ( $P$ ), recall ( $R$ ), and F1 measures:

$$P = \frac{|TP|}{|TP \cup FP|} \quad (10)$$

$$R = \frac{|TP|}{|TP \cup FN|} \quad (11)$$

$$F1 = 2 \cdot \frac{P \cdot R}{P + R} \quad (12)$$

We distinguish between two evaluation settings: strict and relaxed. In the strict setting, mentions  $m_i$  identified by the entity recognition component for task T3 are considered true positives ( $TP$ ) if they are identical to a gold standard mention  $m_g$ . The entity classification task (T4) also requires that both entities have been assigned to the same entity type  $t_i$ , and the linking task (T5) requires linking the mention to the correct knowledge base entity  $e_i$ .

The *relaxed setting* eases these conditions by also considering mentions that overlap a gold standard mention as correct.

Entities that do not appear in the gold standard are considered false positives ( $FP$ ), and false negatives ( $FN$ ) refer to gold standard entities which have been missed by the entity extraction task.

#### 4.1.4. Slot Filling

The slot-filling task (T6) combines all the tasks above. Page segment recognition (T1) identifies page segments. Afterward, the page segment classification (T2) assigns them to the corresponding segment cluster, and entity recognition (T3), entity classification (T4), and entity linking (T5) are performed. Finally, we contextualize the extracted entities  $e_i$  based on the classification of the page segment in which they have occurred and assign them to the corresponding slot.

#### 4.1.5. Experiments and Discussion

Table 6 summarizes the evaluation results for all six evaluation tasks. For the entity linking and slot-filling task, the evaluation also distinguishes between the strict and the relaxed setting.

**Table 6.** Slot filling and per component evaluation results.

Component	P	R	F1
T1: page segment recognition	0.82	0.84	0.83
T2: page segment classification	0.82	0.84	0.83
T3: entity recognition	0.82	0.66	0.73
T4: entity classification	0.78	0.63	0.70
T5: entity linking (strict)	0.67	0.80	0.73
T5: entity linking (relaxed)	0.67	0.82	0.74
T6: slot filling (strict)	0.48	0.60	0.54
T6: slot filling (relaxed)	0.50	0.62	0.55

A comparison of the page segment recognition (T1) and page segment classification (T2) performance reveals the same scores for both tasks. This confirms that the developed simple segment classification heuristic (Section 3.1.1) has been very effective and has classified all page segments correctly.

The evaluation also indicates that both entity recognition and entity classification have been optimized towards a higher precision, to spare domain experts, which need to confirm new entity types in the production process. Entity linking, in contrast, has been optimized towards a higher recall, as shown in the evaluation results.

The evaluation of the overall slot-filling process only considers correctly assigned slots (i.e., course, slot, and slot value are correct) as true positives. All other extracted values are considered false positives, and missing values as false negatives. The F1 score of the slot-filling process indicates that the system is not yet fully suitable for automated knowledge graph population, but rather enables a semi-automated process that significantly improves throughput when compared to the prior deployed manual approaches.

#### 4.1.6. Automatic Knowledge Graph Population

Running the presented system on 55-course descriptions retrieved from the gold standard's second partition extends the knowledge graph by 453 unique statements. Most of these statements (222) describe the course content, followed by target groups (90), learning objectives (61), course prerequisites (51), and certificates (29). In addition, 511 slot values have been marked as "related" since the system has not been able to unequivocally resolve their slot, due to shortcomings in the page segmentation process. This result indicates that improving the page partitioning process will be key to further enhancing the system's recall.

When applied to a corpus of 97,142 educations from 488 different education providers, the system yields a knowledge graph that comprises 73,969 nodes and 734,447 edges. This comprehensive knowledge graph provides the basis for the evaluation of the continuing education recommender in Section 4.3 and is also used in commercial settings as outlined in Section 5.

#### 4.2. Career Path Recommender

The following section presents the evaluation of the career path recommender, which suggests target occupations based on a user's current occupation, skills, and preferences.

One of our first insights, when designing the evaluation of the recommender system, has been the strong impact of user preferences on the evaluation outcome. Parameters such as salary constraints and expectations on the market's current and future demand for a particular occupation are highly user-specific and easy to integrate as subsequent filtering or re-ranking steps (Section 3.2.3).

To objectify the evaluation process we, therefore, only consider the user's current skills and assess career paths based on the required upskilling and reskilling efforts. In this setting, the recommender aims at minimizing the time jobseekers spend in retraining by suggesting occupations based on their similarity.

#### 4.2.1. Gold Standard

The gold standard dataset provides expert rankings of career paths for users with different occupations and varying educational backgrounds.

The annotation process involved three domain experts with at least 18 months of experience in the human resource domain who ranked suggested career paths based on information provided by career counseling services such as the official Swiss career counseling platform (<https://berufsberatung.ch>, accessed on 17 October 2022) and the following criteria:

1. similarity between the current occupation and the suggested target job, and
2. availability of shortened reskilling and upskilling programs for a given job pair.

The experts used pre-tests on two independent datasets to identify disagreements and improve the ranking guidelines accordingly. Entries in the first datasets were jointly ranked and discussed by the experts. Afterward, they individually ranked entries in the second dataset, compared their rankings, and discussed disagreements.

Finally, the gold standard has been created by randomizing and then ranking the recommender's top 15 job suggestions for the following six example profiles used in prior work by Inglin [58]:

- employees with no formal vocational education that work in occupations requiring little training (office assistant, production employee)
- employees working in a skilled craft or trade (painter, electrician)
- highly-skilled employees in occupations that require a university degree (junior business analyst, commercial computer scientist)

If multiple target occupations are equally well suited in terms of these criteria, the experts will award them with the same rank. For the commercial computer scientist, for example, all three experts considered the following target jobs as equally well-matched: *Applications Integrator, Business Intelligence Consultant, Data Architect, IT Consultant, IT Business Analyst, Requirements Engineer*.

The benchmarking dataset contains the experts' individual rankings and a consolidated one that has been derived by majority voting. The evaluation in the next section also reports the average expert agreement with the consolidated ranking to provide insights into the level of consensus between the experts.

#### 4.2.2. Evaluation Metrics and Results

Based on the provided rankings, we computed the precision ( $P@3$ ) and mean average ranking performance (MAP) between the top three expert and system suggestions as

$$P@3 = \frac{|\{\text{expert suggestions}\} \cap \{\text{system suggestions}\}|}{|\{\text{system suggestions}\}|} \quad (13)$$

$$MAP(3) = \frac{1}{3} \sum_{i=1}^3 \left( \frac{1}{i} \sum_{k=1}^i P@k \right) \quad (14)$$

with  $P@k$  indicating the precision of the first  $k$  elements returned by the recommender.

Table 7 summarizes the evaluation results as well as the expert agreement. A drill-down analysis revealed the following insights:

- The recommender works well if closely related target occupations exist. An office assistant, for example, shares many skills with management assistants, office managers, and commercial employees, which makes all three professions suitable career paths. The same is true for the painter, which yields a plasterer as an alternative. Again,



these two occupations are highly related and, therefore, share a considerable amount of skills.

- If direct reskilling paths are not available, the recommendations become much more difficult, which is also illustrated in a lower agreement between the domain experts (column “avg. experts” in Table 7). Consequently, suggesting career pathways for an electrician and production employee is a considerably harder task, which yields even significant disagreement among experts.
- A notable exception to these observations is the commercial computer specialist, for whom a lot of useful alternatives have been proposed. The low score for this use case has been caused by the different rankings produced by experts and the system. Nevertheless, the experts also considered the system’s three target occupations useful career suggestions.

**Table 7.** Mean average ranking performance, precision, and average expert agreement for the career path recommendations.

Prior Education	Occupation	System		avg. Experts	
		MAP(3)	P@3	MAP(3)	P@3
no formal education	office assistant	1.00	1.00	1.00	1.00
	production employee	0.28	0.33	0.87	0.78
craft or trade	electrician	0.28	0.33	0.83	0.67
	painter	0.61	0.33	1.00	1.00
university degree	junior business analyst	0.89	0.67	1.00	1.00
	commercial computer scientist	0.28	0.33	1.00	1.00

#### 4.3. Continuing Education Recommender

The evaluation of the continuing education recommender draws upon (i) the six example occupations used in the previous section, and (ii) the top three target occupations proposed by the domain experts for these occupations. The continuing education recommender then uses the knowledge graph created with the methods introduced in Section 3.1 to suggest suitable educations supporting the provided career paths.

##### 4.3.1. Limitations

Initially, we tried to rank the recommended educations, but analyzing the expert suggestions and subsequent discussions during pre-testing quickly revealed that the agreement between the experts has been too low for a joint ranking since

1. the perceived value of education options differed considerably between experts, which made it infeasible to provide a consolidated ranking;
2. some suggested career paths do not necessarily require any further education (e.g., the promotion from an office assistant to an office manager); and
3. a considerable number of career paths are not yet covered in the continuing education knowledge graph so no useful recommendations could be found. As outlined in Section 3.1, the knowledge graph population component draws upon the offerings of a curated list of education providers. Consequently, its recall is fairly well for formal education (e.g., studies and post-graduate courses), and popular continuing education topics (e.g., languages, computer skills, etc.). Apprenticeships, in contrast, are rarely covered, since they are typically offered by companies and trades rather than educational institutions. Future work will address this issue, by integrating knowledge from apprentice position directories.

We had to remove eight career paths not yet sufficiently covered by the continuing education knowledge graph, since they required education outside the scope of the crawled education providers: (i) paths involving apprenticeships such as the painter, electrician,

and the upskilling of an office assistant to a commercial employee; and (ii) the career path from junior business analyst to SAP business analyst.

#### 4.3.2. Evaluation Metrics and Results

We implemented an evaluation that considers these limitations by letting the domain experts judge a randomized set of the recommender's top 15 continuing education suggestions based on the following two criteria:

1. *benefit*, i.e., whether the suggested education facilitates working in the aspired target occupation (e.g., by providing required skills); and
2. *sufficiency* which requires the experts to judge whether candidates will be able to work in the target occupation once they complete the proposed education.

Even this simplified setting was still very challenging. The experts achieved for the first criteria (benefit) a Fleiss' Kappa of 0.42 which is considered a moderate agreement. The assessment of the education's sufficiency posed an even greater challenge and would have required a strategy for handling cases in which the education covered parts of the relevant skills but failed to consider others. Consequently, experts only achieved a Kappa of 0.19 (slight agreement) for this task.

The evaluation then computes the system's  $P@3$  score for both criteria with  $P^b@3$  indicating the education's beneficialness and  $P^s@3$  its sufficiency. Table 8 summarizes the evaluation results. Although most of the recommender's suggestions are beneficial as well as sufficient, the system does not yet correctly handle career paths that do not require additional formal education (e.g., from junior to more senior positions or from the office assistant to office manager). This has been particularly problematic for the career path to the IT business analyst, where none of the suggested educations has been judged beneficial by the domain experts ( $P^b@3 = 0$ ), although they improved the employee's fit to the target occupation.

Another serious problem is the system's failure to identify skills that are *required* rather than *desired* to work in a target occupation. Although the algorithm already considers a skill's uniqueness as a weighting criterion, this information is not yet sufficient to choose the right education for the career path between production employees and warehouse clerks. A drill-down analysis revealed that addressing this issue would require adding information on a skill's importance to the occupation knowledge base.

**Table 8.** Evaluation of the continuing education recommender.

Occupation Start	Target	$P^b@3$	System $P^s@3$
office assistant	assistant to the manager	0.33	1.00
	office manager	0.33	1.00
production employee	warehouse clerk	0.67	0.00
	logistician	0.67	0.33
	production specialist	1.00	0.67
junior business analyst	business analyst	0.67	1.00
	business analysis manager	0.67	1.00
commercial computer scientist	application integrator	0.33	1.00
	data architect	0.33	0.33
	IT business analyst	0.00	1.00

## 5. Discussion

Knowledge-driven recommender systems base their suggestions on formalized background knowledge rather than user behavior. They do not rely upon behavioral data and are, therefore, not affected by the "cold start" problem of collaborative approaches.

The research presented in this paper introduces

1. a knowledge graph construction method used for creating a real-time continuing education knowledge graph that summarizes knowledge extracted from education provider websites; and
2. a recommender system that draws upon an occupation knowledge base and the extracted knowledge on continuing education for suggesting career paths and education facilitating them.

Our evaluation of the knowledge graph construction method focused on five tasks performed by the knowledge extraction components (page segment recognition, page segment classification, entity recognition, entity classification, and entity linking), and the overall slot-filling task. The evaluation identified the content classification component, the recall of the entity classification task, and the disambiguation algorithm deployed for entity linking as major areas for improvement. Since slot filling relies upon the outcome of the preceding methods, addressing these shortcomings would also improve its overall performance.

The automatic knowledge graph construction component has been deployed to a corpus of 97,142 educations yielding a continuing education graph that comprises 73,969 nodes and 734,447 edges. The created knowledge graph has been applied to multiple industrial settings in which the reported error types do not significantly impact usefulness. External industrial stakeholders use the continuing education graph for applications such as enabling semantic search across education programs and data analytics. Stakeholder interviews which included experts and the senior management from three different companies that use the continuing education graph in their products revealed that stakeholders rated the data quality either as satisfactory (4 out of 6 points) or good (5 out of 6 points). The external stakeholders acknowledged the potential of the automatically created continuing education knowledge graph and see clear benefits from its current use within their businesses.

One reason for the positive stakeholder assessment lies within the distribution of continuing education offerings across the analyzed websites. Large education providers not only tend to correctly implement web standards (which improves the accuracy of the introduced content extraction components) but also contribute most of the relevant education. The CareerCoach 2022 gold standard used in Section 4.1, in contrast, has been designed with source variety in mind and, therefore, underestimated the system's performance in a real-world setting.

The evaluation of the recommender system outlined in Section 3.2 also leverages the created continuing education knowledge graph in conjunction with the  $\times 28$  occupation knowledge base. For use cases that require higher levels of precision and recall, the system could be deployed as part of a semi-automatic knowledge graph-building process that increases efficiency and effectiveness by providing domain experts with suggestions for integrating new concepts and relations into the knowledge graph.

The evaluation of the recommender system aimed at (i) obtaining information on the system's performance for career path and education suggestions, (ii) outlining the method's potential, and (iii) collecting information on how weaknesses of the continuing education knowledge graph impact the system's performance.

Literature indicates, that the sophistication of the employees' current occupation considerably impacts their willingness to participate in reskilling and upskilling activities. Highly-skilled workers are considerably more likely to partake in continuing education than employees that work in occupations that require little or no formal vocational education and a lot of routine tasks [59]. Our experiments aimed at considering these different employee segments by benchmarking system performance for occupations with different educational requirements.

The career path recommender suggests career paths based on the industry partner's occupation knowledge base. An in-depth analysis of its performance revealed that it worked particularly well if closely related target occupations exist. In cases where such closely related occupations have not been available, the quality of its recommendations deteriorated, but so does the agreement between experts, which provided the gold standard for ranking the recommender.

Evaluation results for the education recommender, in contrast, have been strongly affected by the occupation's coverage in the continuing education ontology. In addition, the recommendation task has been significantly more difficult, which is also reflected in a low expert agreement (moderate agreement for an education's beneficialness and only slight agreement for its sufficiency). The evaluation yielded the following key insights:

- suggesting education is a challenging task and even experts struggle with providing consistent recommendations. Future work will mitigate this issue by developing strategies for edge cases such as education that only covers parts of the relevant skills.
- one of the system's biggest strengths, the availability of real-time information on online courses and educational offerings that have been directly obtained from the provider's websites, also became its major weakness, since education that has not been covered in the input sources are not considered. In Switzerland, crafts, and trades, for example, are taught through apprenticeships. Consequently, the coverage of continuing education for crafts and trades has been insufficient within the continuing education ontology forcing us to remove a total of eight career paths from the evaluation. In addition, the career path to SAP business analyst had to be discarded, since no suitable education had been available in the knowledge graph.
- the system does not yet consider the efforts required for completing further education. Consequently, it preferred more comprehensive education over quicker ones. Edge cases demonstrating this problem have been career paths where on-the-job experience could have been sufficient for advancing to a more prestigious occupation (e.g., from office assistant to office manager). Although all the system's recommendations have been suitable and would have been beneficial towards a possible promotion, domain experts did not see a requirement for further education.

## 6. Outlook and Conclusions

This paper introduced a knowledge graph construction system that combines knowledge extraction methods towards slot filling to extract continuing education offerings from websites. We then defined six evaluation tasks for benchmarking slot filling and knowledge extraction on the CareerCoach 2022 gold standard, which provided detailed information on the component's performance.

Applying the system to a corpus of Swiss education sites yielded a knowledge graph that comprises 73,969 nodes and 734,447 edges covering educational offerings such as academic programs, continuing education programs, courses, seminars, and online courses.

Both, the created continuing education knowledge graph and the industry partner's occupation knowledge base, act as knowledge sources for a recommender system that suggests career paths and the corresponding continuing education offerings to its users. The system avoids the cold-start problem by drawing upon these knowledge sources and supports filtering and re-ranking criteria such as similarity of job zones (i.e., education levels), the job's long-term prospects in terms of expected demand, and salary restrictions, which can be added via user preferences.

An evaluation used six well-defined use cases to compare the system's mean average ranking performance and precision for career path recommendations to an expert ranking. Afterward, experts rated the system's continuing education recommendations for the suggested career paths based on their beneficialness and sufficiency.

Future work will focus on improving upon the shortcomings identified in the system's evaluation, particularly the knowledge graph's coverage, by integrating companies that offer apprenticeships, providing means for career paths for which further education is beneficial rather than required, and considering the education's cost (particularly the required time and efforts) in the recommendations. We also plan to improve slot-filling performance by enhancing page segmentation, and further fine-tuning the entity classification and linking components.

**Author Contributions:** Conceptualization, A.W.; methodology, A.W., A.F., R.W., A.v.S., and P.K.; software, R.W., A.F., A.v.S., and P.K.; validation, R.W., A.F., and P.K.; formal analysis, A.W.; resources, R.W. and A.v.S.; data curation, R.W. and A.v.S.; writing—original draft preparation, A.W., R.W., A.F., A.v.S., and P.K.; writing—review and editing, A.W., R.W., A.F., A.v.S., and P.K.; visualization, A.W., R.W., P.K., and A.F.; supervision, A.W.; project administration, A.W.; funding acquisition, A.W. All authors have read and agreed to the published version of the manuscript.

**Funding:** The research presented in this paper has been conducted within the CareerCoach project ([www.fhgr.ch/CareerCoach](http://www.fhgr.ch/CareerCoach), accessed on 17 October 2022) which is funded by Innosuisse under grant number 48713.1 IP-ICT.

**Institutional Review Board Statement:** Not applicable.

**Informed Consent Statement:** Not applicable.

**Data Availability Statement:** The CareerCoach 2022 dataset has been created and utilized within this study. The dataset is publicly available at <https://github.com/fhgr/careercoach2022>, accessed on 17 October 2022.

**Acknowledgments:** The authors would like to thank Cornel Müller and Matthias Hewelt for their support in acquiring and performing the CareerCoach project.

**Conflicts of Interest:** The authors declare no conflict of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript, or in the decision to publish the results.

## References

- Daiber, J.; Jakob, M.; Hokamp, C.; Mendes, P.N. Improving Efficiency and Accuracy in Multilingual Entity Extraction. In Proceedings of the 9th International Conference on Semantic Systems, I-SEMANTICS '13, Graz, Austria, 4–6 September 2013; ACM: New York, NY, USA, 2013; pp. 121–124. [\[CrossRef\]](#)
- Weichselbraun, A.; Kuntschik, P.; Brasoveanu, A.M. Name Variants for Improving Entity Discovery and Linking. In Proceedings of the Second Conference on Language, Data and Knowledge (LDK 2019), Leipzig, Germany, 20–23 May 2019; OpenAccess Series in Informatics: Leipzig, Germany, 2019; Volume 70, pp. 14:1–14:15. [\[CrossRef\]](#)
- Hogan, A.; Blomqvist, E.; Cochez, M.; d'Amato, C.; de Melo, G.; Gutierrez, C.; Gayo, J.E.L.; Kirrane, S.; Neumaier, S.; Polleres, A.; et al. Knowledge Graphs. *ACM Comput. Surv.* **2022**, *54*, 1–37. [\[CrossRef\]](#)
- World Economic Forum-Centre for the New Economy and Society. *The Future of Jobs Report 2020*; Technical report; World Economic Forum-Centre for the New Economy and Society: Davos, Switzerland, 2020.
- Noy, N.; Gao, Y.; Jain, A.; Narayanan, A.; Patterson, A.; Taylor, J. Industry-scale Knowledge Graphs: Lessons and Challenges. *Commun. ACM* **2019**, *62*, 36–43. [\[CrossRef\]](#)
- Bizer, C.; Lehmann, J.; Kobilarov, G.; Auer, S.; Becker, C.; Cyganiak, R.; Hellmann, S. DBpedia—A crystallization point for the Web of Data. *J. Web Semant. Sci. Serv. Agents World Wide Web* **2009**, *7*, 154–165. [\[CrossRef\]](#)
- Vrandečić, D.; Krötzsch, M. Wikidata: A Free Collaborative Knowledgebase. *Commun. ACM* **2014**, *57*, 78–85. [\[CrossRef\]](#)
- Lin, X.; Chen, L. Canonicalization of Open Knowledge Bases with Side Information from the Source Text. In Proceedings of the 2019 IEEE 35th International Conference on Data Engineering (ICDE), Macao, China, 8–11 April 2019; pp. 950–961. [\[CrossRef\]](#)
- Dubey, M.; Banerjee, D.; Abdelkawi, A.; Lehmann, J. LC-QuAD 2.0: A Large Dataset for Complex Question Answering over Wikidata and DBpedia. In Proceedings of the The Semantic Web—ISWC 2019, Auckland, New Zealand, 26–30 October 2019; Ghidini, C., Hartig, O., Maleshkova, M., Svátek, V., Cruz, I., Hogan, A., Song, J., Lefrançois, M., Gandon, F., Eds.; Lecture Notes in Computer Science; Springer International Publishing: Cham, Switzerland, 2019; pp. 69–78. [\[CrossRef\]](#)
- Usbeck, R.; Ngomo, A.C.N.; Haarmann, B.; Krithara, A.; Röder, M.; Napolitano, G. 7th Open Challenge on Question Answering over Linked Data (QALD-7). In Proceedings of the Semantic Web Challenges, Portoroz, Slovenia, 28 May–1 June 2017; Communications in Computer and Information Science; Dragoni, M., Solanki, M., Blomqvist, E., Eds.; Springer International Publishing: Cham, Switzerland, 2017; pp. 59–69. [\[CrossRef\]](#)
- Elsahar, H.; Vougiouklis, P.; Remaci, A.; Gravier, C.; Hare, J.; Laforest, F.; Simperl, E. T-REx: A Large Scale Alignment of Natural Language with Knowledge Base Triples. In Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018), Miyazaki, Japan, 7–12 May 2018; European Language Resources Association (ELRA): Miyazaki, Japan, 2018.
- Glass, M.; Gliozzo, A. A Dataset for Web-Scale Knowledge Base Population. In Proceedings of the The Semantic Web, Heraklion, Greece, 3–7 June 2018; Lecture Notes in Computer Science; Gangemi, A., Navigli, R., Vidal, M.E., Hitzler, P., Troncy, R., Hollink, L., Tordai, A., Alam, M., Eds.; Springer International Publishing: Cham, Switzerland, 2018; pp. 256–271. [\[CrossRef\]](#)

13. Mesquita, F.; Cannaviccio, M.; Schmidek, J.; Mirza, P.; Barbosa, D. KnowledgeNet: A Benchmark Dataset for Knowledge Base Population. In Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), Hong Kong, China, 3–7 November 2019; Association for Computational Linguistics: Hong Kong, China, 2019; pp. 749–758. [\[CrossRef\]](#)
14. Lin, X.; Li, H.; Xin, H.; Li, Z.; Chen, L. KBPearl: A knowledge base population system supported by joint entity and relation linking. *Proc. Vldb Endow.* **2020**, *13*, 1035–1049. [\[CrossRef\]](#)
15. Li, J.; Sun, A.; Han, J.; Li, C. A Survey on Deep Learning for Named Entity Recognition. In *IEEE Transactions on Knowledge & Data Engineering*; IEEE Computer Society: Los Alamitos, CA, USA, 2020; pp. 50–70. [\[CrossRef\]](#)
16. Fu, J.; Huang, X.; Liu, P. SpanNER: Named Entity Re-/Recognition as Span Prediction. *arXiv* **2021**, arXiv:2106.00641.
17. Yosef, M.A.; Hoffart, J.; Bordino, I.; Spaniol, M.; Weikum, G. AIDA: An Online Tool for Accurate Disambiguation of Named Entities in Text and Tables. *PVLDB* **2011**, *4*, 1450–1453. [\[CrossRef\]](#)
18. Guo, Y.; Che, W.; Liu, T.; Li, S. A Graph-based Method for Entity Linking. In Proceedings of the 5th International Joint Conference on Natural Language Processing, Chiang Mai, Thailand, 9–11 November 2011; Asian Federation of Natural Language Processing: Chiang Mai, Thailand, 2011; pp. 1010–1018.
19. Moro, A.; Raganato, A.; Navigli, R. Entity Linking meets Word Sense Disambiguation: A Unified Approach. *Trans. Assoc. Comput. Linguist.* **2014**, *2*, 231–244. [\[CrossRef\]](#)
20. Usbeck, R.; Ngonga Ngomo, A.C.; Röder, M.; Gerber, D.; Coelho, S.; Auer, S.; Both, A. AGDISTIS - Graph-Based Disambiguation of Named Entities Using Linked Data. In Proceedings of the Semantic Web—ISWC 2014, Riva del Garda, Italy, 19–23 October 2014; Lecture Notes in Computer Science; Springer International Publishing: Cham, Switzerland, 2014; Volume 8796, pp. 457–471.
21. Moussallem, D.; Usbeck, R.; Röder, M.; Ngomo, A.C.N. MAG: A Multilingual, Knowledge-base Agnostic and Deterministic Entity Linking Approach. In Proceedings of the Knowledge Capture Conference on—K-CAP 2017, Austin, TX, USA, 4–6 December 2017; pp. 1–8. [\[CrossRef\]](#)
22. Cao, Y.; Hou, L.; Li, J.; Liu, Z. Neural Collective Entity Linking. *arXiv* **2018**, arXiv:1811.08603.
23. Wu, J.; Zhang, R.; Mao, Y.; Guo, H.; Soflaei, M.; Huai, J. Dynamic Graph Convolutional Networks for Entity Linking. In Proceedings of the Web Conference, WWW '20, Taipei, Taiwan, 20–24 April 2020; Association for Computing Machinery: Taipei, Taiwan, 2020; pp. 1149–1159. [\[CrossRef\]](#)
24. Ding, W.; Chaudhri, V.K.; Chittar, N.; Konakanchi, K. JEL: Applying End-to-End Neural Entity Linking in JPMorgan Chase. *Proc. Aaai Conf. Artif. Intell.* **2021**, *35*, 15301–15308. [\[CrossRef\]](#)
25. Lim, S.; Kwon, S.; Lee, S.; Choi, J. UNIST SAIL System for TAC 2017 Cold Start Slot Filling. In Proceedings of the Text Analysis Conference TAC 2017, Gaithersburg, MD, USA, 13–14 November 2017.
26. Roth, B.; Barth, T.; Wiegand, M.; Singh, M.; Klakow, D. Effective Slot Filling Based on Shallow Distant Supervision Methods. *arXiv* **2014**, arXiv:1401.1158.
27. Wang, H.; Zhang, F.; Wang, J.; Zhao, M.; Li, W.; Xie, X.; Guo, M. Exploring High-Order User Preference on the Knowledge Graph for Recommender Systems. *ACM Trans. Inf. Syst.* **2019**, *37*, 32:1–32:26. [\[CrossRef\]](#)
28. Ritze, D.; Lehmborg, O.; Oulabi, Y.; Bizer, C. Profiling the Potential of Web Tables for Augmenting Cross-domain Knowledge Bases. In Proceedings of the 25th International Conference on World Wide Web, WWW '16, Montréal, QC, Canada, 11–15 April 2016; International World Wide Web Conferences Steering Committee: Montréal, QC, Canada, 2016; pp. 251–261. [\[CrossRef\]](#)
29. Siddique, A.; Jamour, F.; Hristidis, V. Linguistically-Enriched and Context-Aware Zero-shot Slot Filling. In Proceedings of the Web Conference, WWW '21, 2021, Virtual Event / Ljubljana, Slovenia, 19–23 April 2021; Association for Computing Machinery: New York, NY, USA, 2021; pp. 3279–3290. [\[CrossRef\]](#)
30. Coucke, A.; Saade, A.; Ball, A.; Bluche, T.; Caulier, A.; Leroy, D.; Doumouro, C.; Gisselbrecht, T.; Caltagirone, F.; Lavril, T.; et al. Snips Voice Platform: An embedded Spoken Language Understanding system for private-by-design voice interfaces. *arXiv* **2018**, arXiv:1805.10190.
31. Liu, X.; Eshghi, A.; Swietojanski, P.; Rieser, V. Benchmarking natural language understanding services for building conversational agents. In *Increasing Naturalness and Flexibility in Spoken Dialogue Interaction*; Marchi, E., Siniscalchi, S.M., Cumani, S., Salerno, V.M., Li, H., Eds; Springer: Singapore, 2021. [\[CrossRef\]](#)
32. Zang, X.; Rastogi, A.; Sunkara, S.; Gupta, R.; Zhang, J.; Chen, J. MultiWOZ 2.2 : A Dialogue Dataset with Additional Annotation Corrections and State Tracking Baselines. In Proceedings of the 2nd Workshop on Natural Language Processing for Conversational AI, Online, 9 July 2020; Association for Computational Linguistics: Florence, Italy, 2020; pp. 109–117. [\[CrossRef\]](#)
33. Rastogi, A.; Zang, X.; Sunkara, S.; Gupta, R.; Khaitan, P. Towards Scalable Multi-domain Conversational Agents: The Schema-Guided Dialogue Dataset. In Proceedings of the AAAI Conference on Artificial Intelligence, New York, NY, USA, 7–12 February 2020; Volume 34, pp. 8689–8696.
34. Liu, Z.; Winata, G.I.; Xu, P.; Fung, P. Coach: A Coarse-to-Fine Approach for Cross-domain Slot Filling. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, Online, 5–10 July 2020; Association for Computational Linguistics: Florence, Italy, 2020; pp. 19–25. [\[CrossRef\]](#)
35. Shah, D.; Gupta, R.; Fayazi, A.; Hakkani-Tur, D. Robust Zero-Shot Cross-Domain Slot Filling with Example Values. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, Florence, Italy, 28 July–2 August 2019; Association for Computational Linguistics: Florence, Italy, 2019; pp. 5484–5490. [\[CrossRef\]](#)

36. Bapna, A.; Tur, G.; Hakkani-Tur, D.; Heck, L. Towards zero-shot frame semantic parsing for domain scaling. In Proceedings of the Interspeech 2017, Stockholm, Sweden, 20–24 August 2017; pp. 2476–2480. [[CrossRef](#)]
37. Cerezo-Costas, H.; Martín-Vicente, M. Relation Extraction for Knowledge Base Completion: A Supervised Approach. In Proceedings of the Semantic Web Challenges, Heraklion, Greece, 3–7 June 2018; Communications in Computer and Information Science; Buscaldi, D., Gangemi, A., Reforgiato Recupero, D., Eds.; Springer International Publishing: Cham, Switzerland, 2018; pp. 52–66. [[CrossRef](#)]
38. Chabchoub, M.; Gagnon, M.; Zouaq, A. Collective Disambiguation and Semantic Annotation for Entity Linking and Typing. In Proceedings of the Semantic Web Challenges, Heraklion, Greece, 29 May–2 June 2016; Communications in Computer and Information Science; Sack, H., Dietze, S., Tordai, A., Lange, C., Eds.; Springer International Publishing: Cham, Switzerland, 2016; pp. 33–47. [[CrossRef](#)]
39. Sakor, A.; Onando Mulang, I.; Singh, K.; Shekarpour, S.; Esther Vidal, M.; Lehmann, J.; Auer, S. Old is Gold: Linguistic Driven Approach for Entity and Relation Linking of Short Text. In Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers), Minneapolis, MN, USA, 2–7 June 2019; Association for Computational Linguistics: Minneapolis, MN, USA, 2019; pp. 2336–2346. [[CrossRef](#)]
40. Zhang, S.; Yao, L.; Sun, A.; Tay, Y. Deep Learning Based Recommender System: A Survey and New Perspectives. *ACM Comput. Surv.* **2019**, *52*, 5. [[CrossRef](#)]
41. Adomavicius, G.; Tuzhilin, A. Toward the Next Generation of Recommender Systems: A Survey of the State-of-the-Art and Possible Extensions. *IEEE Trans. Knowl. Data Eng.* **2005**, *17*, 734–749. [[CrossRef](#)]
42. Hu, B.; Shi, C.; Zhao, W.X.; Yu, P.S. Leveraging Meta-path based Context for Top- N Recommendation with A Neural Co-Attention Model. In Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD'18, London, UK, 19–23 August 2018; Association for Computing Machinery: London, UK, 2018; pp. 1531–1540. [[CrossRef](#)]
43. Yu, X.; Ren, X.; Sun, Y.; Gu, Q.; Sturt, B.; Khandelwal, U.; Norick, B.; Han, J. Personalized entity recommendation: A heterogeneous information network approach. In Proceedings of the 7th ACM International Conference on Web Search and Data Mining, WSDM '14, New York, NY, USA, 24–28 February 2014; Association for Computing Machinery: New York, NY, USA, 2014; pp. 283–292. [[CrossRef](#)]
44. Zhao, H.; Yao, Q.; Li, J.; Song, Y.; Lee, D.L. Meta-Graph Based Recommendation Fusion over Heterogeneous Information Networks. In Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '17, Halifax, Canada, 13–17 August 2017; Association for Computing Machinery: Halifax, NS, Canada, 2017; pp. 635–644. [[CrossRef](#)]
45. Huang, J.; Zhao, W.X.; Dou, H.; Wen, J.R.; Chang, E.Y. Improving Sequential Recommendation with Knowledge-Enhanced Memory Networks. In Proceedings of the The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval, SIGIR '18, Ann Arbor, MI, USA, 8–12 July 2018; Association for Computing Machinery: Ann Arbor, MI, USA, 2018; pp. 505–514. [[CrossRef](#)]
46. Wang, H.; Zhang, F.; Zhao, M.; Li, W.; Xie, X.; Guo, M. Multi-Task Feature Learning for Knowledge Graph Enhanced Recommendation. In Proceedings of the The World Wide Web Conference, WWW '19, San Francisco, CA, USA, 13–17 May 2019; Association for Computing Machinery: San Francisco, CA, USA, 2019; pp. 2000–2010. [[CrossRef](#)]
47. Wang, H.; Zhang, F.; Xie, X.; Guo, M. DKN: Deep Knowledge-Aware Network for News Recommendation. In Proceedings of the 2018 World Wide Web Conference, WWW '18, Lyon, France, 23–27 April 2018; International World Wide Web Conferences Steering Committee: Lyon, France, 2018; pp. 1835–1844. [[CrossRef](#)]
48. Zhang, F.; Yuan, N.J.; Lian, D.; Xie, X.; Ma, W.Y. Collaborative Knowledge Base Embedding for Recommender Systems. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '16, San Francisco, CA, USA, 13–17 August 2016; Association for Computing Machinery: San Francisco, CA, USA, 2016; pp. 353–362. [[CrossRef](#)]
49. Sun, Z.; Yang, J.; Zhang, J.; Bozzon, A.; Huang, L.K.; Xu, C. Recurrent knowledge graph embedding for effective recommendation. In Proceedings of the 12th ACM Conference on Recommender Systems, RecSys '18, Vancouver, Canada, 2–7 October 2018; Association for Computing Machinery: Vancouver, BC, Canada, 2018; pp. 297–305. [[CrossRef](#)]
50. Wang, H.; Zhang, F.; Wang, J.; Zhao, M.; Li, W.; Xie, X.; Guo, M. RippleNet: Propagating User Preferences on the Knowledge Graph for Recommender Systems. In Proceedings of the 27th ACM International Conference on Information and Knowledge Management, CIKM '18, Torino, Italy, 22–26 October 2018; Association for Computing Machinery: Torino, Italy, 2018; pp. 417–426. [[CrossRef](#)]
51. Wang, H.; Zhao, M.; Xie, X.; Li, W.; Guo, M. Knowledge Graph Convolutional Networks for Recommender Systems. In Proceedings of the The World Wide Web Conference, WWW '19, San Francisco, CA, USA, 13–17 May 2019; Association for Computing Machinery: San Francisco, CA, USA, 2019; pp. 3307–3313. [[CrossRef](#)]
52. Turhan, A.Y.; Bechhofer, S.; Kaplunova, A.; Liebig, T.; Luther, M.; Möller, R.; Noppens, O.; Patel-Schneider, P.; Suntisrivaraporn, B.; Weithöner, T. DIG 2.0—Towards a Flexible Interface for Description Logic Reasoners. In Proceedings of the CEUR Workshop on OWL: Experiences and Directions, Athens, GA, USA, 10–11 November 2006; RWTH Aachen University: Aachen, Germany, 2006.
53. Weichselbraun, A.; Kuntschik, P.; Brasoveanu, A.M.P. Mining and Leveraging Background Knowledge for Improving Named Entity Linking. In Proceedings of the 8th International Conference on Web Intelligence, Mining and Semantics, WIMS 2018, Novi Sad, Serbia, 25–27 June 2018. [[CrossRef](#)]

54. Wolf, T.; Debut, L.; Sanh, V.; Chaumond, J.; Delangue, C.; Moi, A.; Cistac, P.; Rault, T.; Louf, R.; Funtowicz, M.; et al. Transformers: State-of-the-art natural language processing. In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, Online, 16–20 November 2020; Association for Computational Linguistics: Florence, Italy, 2020; pp. 38–45.
55. World Economic Forum—Centre for the New Economy and Society. *Towards a Reskilling Revolution*; Technical Report; World Economic Forum-Centre for the New Economy and Society: Davos, Switzerland, 2018.
56. World Economic Forum. *Towards a Reskilling Revolution: Industry-Led Action for the Future of Work*; Technical Report; World Economic Forum: Cologny, Switzerland, 2019.
57. Weichselbraun, A.; Waldvogel, R.; Fraefel, A.; van Schie, A.; Kuntschik, P. Slot Filling for Extracting Reskilling and Upskilling Options from the Web. In Proceedings of the 27th International Conference on Natural Language & Information Systems, Valencia, Spain, 15–17 June 2022.
58. Inglin, M. Re- Und Upskilling-Empfehlung: Kriterien Für Die Automatische Auswahl von Re- Und Upskilling-Angeboten. Bachelor Thesis, University of Applied Sciences of the Grisons, Chur, Switzerland, 2022.
59. Heß, P.; Janssen, S.; Leber, U. *Digitalisierung Und Berufliche Weiterbildung: Beschäftigte, Deren Tätigkeiten Durch Technologien Ersetzbar Sind, Bilden Sich Seltener Weiter*; Institut Für Arbeitsmarkt- Und Berufsforschung (IAB) 16/2019: Nürnberg, Germany, 2019.