



Intellectual Output 1

Modeling and Meta-modeling of Job Knowledge for Labor Market



This work is licensed under the Creative Commons Attribution 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/> or send a letter to Creative Commons, PO Box 1866, Mountain View, CA 94042, USA.

Disclaimer: The European Commission support for the production of this publication does not constitute an endorsement of the contents which reflects the views only of the authors, and the Commission cannot be held responsible for any use which may be made of the information contained therein.



Co-funded by
the European Union



Erasmus+
Enriching lives, opening minds.

Contents

1	Executive Summary	2
2	Introduction	3
2.1	Tasks Constituting IO1	4
2.1.1	O1/A1 - Ontology of the Job Knowledge	4
2.1.2	O1/A2 - Job Knowledge Hierarchical Architecture	4
2.1.3	O1/A3 - Application of the Job Knowledge Schema in Data Science as a Working Example	4
3	Exploring Existing Schemas	4
3.1	schema.org/JobPosting	5
3.2	hr-xml/JobHeader and /PositionHeader	7
4	Comparison of Schemas	8
4.1	Pending fields in schema.org	10
4.2	Drawbacks of Individual Schemas	10
5	Extending the Schema	11
5.1	Additional Attributes	11
5.2	Data Models for Competency Definitions	12
5.3	Semantic Enrichment of JKB Using ESCO	12
5.3.1	Occupations	13
5.3.2	Skills/Competences	13
5.3.3	Qualifications	14
6	Application in ‘Data Science’ – An Example	15
7	Conclusions and Future Work	18

1 Executive Summary

This IO covers the modeling and meta-modeling of the Job Knowledge (JK) as a theoretical and conceptual part of the DISKOW project. Accordingly, this report provides the required terminology for the project, describes the structure and ontology of the JK in detail, and demonstrates the hierarchical architecture of it. With an aim to provide a standard and extensible JK model which can be integrated in a wide variety of tools and software technologies, we have adopted an existing ontology for job knowledge from schema.org and extended it with additional properties following the best practices laid out with respect to linked open data. We showcase the application of the proposed JK model in the field of data science as an example. We thereby envision that the outcome of this IO will facilitate the adaptation of the JK model in a wide variety of sectors and case studies.

2 Introduction

In the last few years, substantial efforts have been invested by governments, international organizations and other institutes to improve our understanding of the dynamics of the labor market. These efforts have resulted in a series of applications, tools and sources that cover labor demands, supply or matching and that shed more light on skill gaps and mismatches. ESCO (European skills, competences, qualifications and occupations tool developed by the EC), Cedefop's skills panorama and EURES (EU's job portal) as well as European Qualifications Framework (EQF) are few examples of such efforts.

Traditionally, to get more insight into labor demands or supply, researchers and policymakers have relied on interviews, trade publications, surveys, and vacancies. While these traditional data sources have some clear advantages, they are also characterized by limitations that can be addressed by using web-based data instead. Web is a gold mine for job knowledge discovery. Linked open data, job announcements, social media, job search engines, forums, wikis, data streams and interlinked information are few examples of such valuable job-related sources on the net.

The main problem in this regard is not in the availability of data and how to retrieve them, but in how to clean, explore, visualize and interpret such a huge volume of various web data. With an aim to streamline this process and make such data suitable for further exploitation (e.g. consumed by specialised mobile apps) an open Job Knowledge Base (JKB) is proposed within the scope of the DISKOW project, that can be used by employers, employees, job seekers, labour market experts and policymakers. Such a JKB contains different types of information such as responsibilities and roles, required competencies (described using existing standards, such as ECF, EQF, etc.), that could be used to develop trainings and identify priorities, wage information, geographical and demographic trends, cultural issues, demands of the job markets in different domains, job announcement information and rates, job popularity and other useful statistics. The DISKOW project aims at creating a JKB prototype, based on existing open-source Business Intelligence platform in order to cover the most important factors in this regard such as required job knowledge for a specific job.

This IO covers the modeling and meta-modeling of the Job Knowledge (JK) as a theoretical and conceptual part of the DISKOW project. Accordingly, this report provides the required terminology for the project, describes the structure and ontology of the JK in detail, and demonstrates the hierarchical architecture of it. With an aim to provide a standard and extensible JK model which can be integrated in a wide variety of tools and software technologies, we have adopted an existing ontology for job knowledge from schema.org and extended it with additional properties following the best practices laid out with respect to linked open data. We showcase the application of the proposed JK model in the field of data science as an example. We thereby envision that the outcome of this IO will facilitate the

adaptation of the JK model in a wide variety of sectors and case studies.

2.1 Tasks Constituting IO1

This report addresses the three main tasks that constitute IO1, in accordance to the Description of Work (DoW) in the DISKOW Grant Agreement. These tasks are described below, and discussed in detail in the following sections of this report.

2.1.1 O1/A1 - Ontology of the Job Knowledge

In this task, we studied the general ontological architecture for job knowledge according to which we will develop the DISKOW JKB (job knowledge base). The outcomes of this task – a schema for job knowledge – will help in the identification of potentially useful data sets, and inform activities within IO2 - *Identification and Analysis of the European Labor Market Data* and in the development, customization activities of the knowledge base within IO4 - *Job Knowledge Base*.

2.1.2 O1/A2 - Job Knowledge Hierarchical Architecture

In this task, we explore the hierarchical architecture afforded by the DISKOW job knowledge schema, developed in alliance with activities within the IO4 - *Job Knowledge Base*. In this regard, competence representation standards will be employed in this task.

2.1.3 O1/A3 - Application of the Job Knowledge Schema in Data Science as a Working Example

In this task, we discuss how the proposed schema for job knowledge can be applied and customized to specific fields such as ‘Data Science’. In this regard, since the case study of this project is on Data Science career, IO2 - *Identification and Analysis of the European Labor Market Data* will also reflect on the availability of data sources to build and support the JKB for Data Science.

3 Exploring Existing Schemas

To identify and develop an easily extensible schema for job knowledge, we explored the existing standards. Two main schemas were shortlisted for comparison, based on their expressibility of job knowledge - the `schema.org/JobPosting` schema and the `hr-xml/JobHeaderAndPositionHeader` schema.

Schema.org is a collaborative, community activity with a mission to create, maintain, and promote schemas for structured data on the Internet, on web pages, in email messages, and beyond¹. The schema.org vocabulary can be used with many different encodings, including RDFa, Microdata and JSON-LD. These vocabularies cover entities, relationships between entities and actions, and can easily be extended through a well-documented extension model. Over 10 million sites use Schema.org to markup their web pages and email messages. Many applications from Google, Microsoft, Pinterest, Yandex and others already use these vocabularies to power rich, extensible experiences. On the other hand, the HR-XML Consortium is the only independent, non-profit, volunteer-led organization dedicated to the development and promotion of a standard suite of XML specifications to enable e-business and the automation of human resources-related data exchanges. Due to these reasons, we considered the `schema.org/JobPosting` and `hr-xml/JobHeaderAndPositionHeader` schemas as potentially viable options for structuring the job knowledge within the DISKOW JKB.

3.1 `schema.org/JobPosting`

The JobPosting schema is a listing that describes a job opening in a certain organization. A toy example using this schema for a job posting related to ‘Data Scientist’ is presented in the listing below. Figure 1 presents an excerpt of the `schema.org/JobPosting` schema.

```
<div>
  <h2>Data Scientist </h2>
  <p><strong>Location:</strong>
    Kirkland WA</p>
  <p><strong>Industry:</strong> Computer Software
  <br><strong>Occupational Category:</strong> 15–1132.00 Software Developers ,
  Application
  <br><strong>Hours:</strong> Full-time , 40 hours per week
  <br><strong>Salary:</strong> USD 100000
  </p>
  <p>
    <strong>Description:</strong> ABC Company Inc .
    seeks a full-time mid-level Data Scientist to analyze and
    develop in-house tools .
  </p>
  <p><strong>Responsibilities:</strong></p>
  <ul>
```

¹<https://schema.org/>

Property	Expected Type	Description
Properties from JobPosting		
applicantLocationRequirements	AdministrativeArea	The location(s) applicants can apply from. This is usually used for telecommuting jobs where the applicant does not need to be in a physical office. Note: This should not be used for citizenship or work visa requirements.
applicationContact	ContactPoint	Contact details for further information relevant to this job posting.
baseSalary	MonetaryAmount or Number or PriceSpecification	The base salary of the job or of an employee in an EmployeeRole.
datePosted	Date	Publication date of an online listing.
educationRequirements	EducationalOccupationalCredential or Text	Educational background needed for the position or Occupation.
employerOverview	Text	A description of the employer, career opportunities and work environment for this position.
employmentType	Text	Type of employment (e.g. full-time, part-time, contract, temporary, seasonal, internship).
employmentUnit	Organization	Indicates the department, unit and/or facility where the employee reports and/or in which the job is to be performed.
estimatedSalary	MonetaryAmount or MonetaryAmountDistribution or Number	An estimated salary for a job posting or occupation, based on a variety of variables including, but not limited to industry, job title, and location. Estimated salaries are often computed by outside organizations rather than the hiring organization, who may not have committed to the estimated value.
experienceRequirements	Text	Description of skills and experience needed for the position or Occupation.
hiringOrganization	Organization	Organization offering the job position.
incentiveCompensation	Text	Description of bonus and commission compensation aspects of the job. Supersedes incentives .

Figure 1: An excerpt of the [schema.org/JobPosting](#) schema.

```

<li>Design and write specifications for tools for in-house customers</li>
<li>Build tools according to specifications</li>
</ul>

```

```

<p><strong>Educational requirements:</strong></p>
<ul>
  <li>Bachelor's Degree in Computer Science ,
    Information Systems or related fields of study.</li>
</ul>
<p><strong>Experience requirements:</strong></p>
<ul>
  <li>Minumum 3 years experience as a Data Science professional</li>
</ul>
<p><strong>Desired Skills:</strong></p>
<ul>
  <li>Web application development using Java/J2EE</li>
  <li>Web application development using Python or familiarity
    with dynamic programming languages</li>
</ul>
<p><strong>Qualifications:</strong></p>
<ul>
  <li>Ability to work in a team environment with members of
    varying skill levels.</li>
  <li>Highly motivated.</li>
  <li>Learns quickly.</li>
</ul>
<p><strong>Benefits:</strong></p>
<ul><li>ABC Corp provides top-tier employee compensation
  benefits and a relaxed , team-oriented work environment ,
  including: Medical , Life , Dental</li>
</ul>
  <p><strong>Incentives:</strong></p>
    <ul><li>Performance-based annual bonus plan ,
      project-completion bonuses</li>
    </ul>
  <p>If interested in this position , please email us your resume ,
  along with salary requirements and a cover letter to Jobs@abc.123.</p>
  <p>Date Posted: 2011-10-31</p>
</div>

```

3.2 hr-xml/JobHeader and /PositionHeader

HR-XML specifications have been completed for: Background Check, Benefits Enrollment, Competencies, Contact Method, DateTime Data Types, Effective Dating, Enrollment, Entity Iden-

tifiers, Job And Position Header, Payroll Benefit Contributions, PersonName, Postal Address, Provisional Envelope Specification, Resume, Staffing Exchange Protocol, Staffing Industry Data Exchange Standards (SIDES), Stock Plan Interface Specification, Time Expense Reporting, Time-Card Configuration, User Area, WorkSite and Environment.

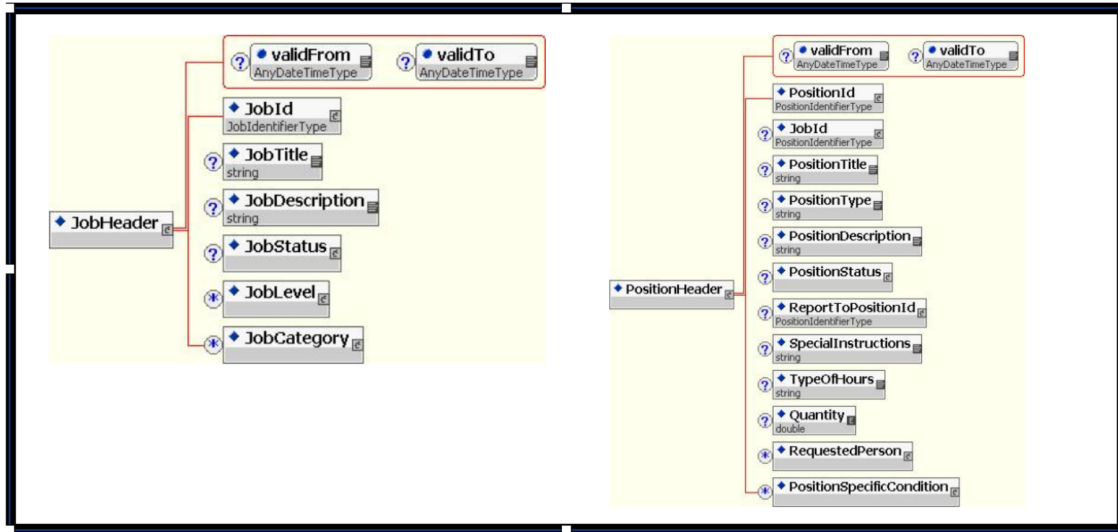


Figure 2: An illustration of the hr-xml/JobHeaderAndPositionHeader schema.

hr-xml/JobHeader and hr-xml/PositionHeader are two elements that can be used to describe a job and position. Both elements can be used together as they share the ‘*JobId*’ element. For convenience, we will refer to this shared schema as the hr-xml/JobHeaderAndPositionHeader schema. An illustration can be seen in Figure 2.

4 Comparison of Schemas

schema.org/JobPosting	hr-xml/JobHeaderAndPositionHeader
-	provides a JobLevel field which can be used to categorise jobs based on levels eg: “Developer”, “Senior Developer”, “Manager”, etc.
-	provides a PositionStatus field which shows status of position eg: filled, vacant, etc
-	provides a ReportToPositionId field which identifies the person/entity you will report to

Table 1 continued from previous page

schema.org/JobPosting	hr-xml/JobHeaderAndPositionHeader
-	provides a SpecialInstructions field eg: Report to 4th floor lobby
-	provides a Quantity field which allows to show the number of positions eg: 4 positions for this JobId
-	provides a PositionSpecificField that has information about any factors that could impact on the risk evaluation or work environment of the position
Provision for salary related information through baseSalary, estimatedSalary, incentiveCompensation fields	-
Provision for specific requirement information like experienceRequirement, educationRequirement, specialCommitment, qualifications, skills	No clear distinction but has only one field PositionSpecificCondition where all requirements can be described together
Provision for hiring entity information through hiringOrganization and industry fields	-
Provides jobBenefits field	-
Provides jobLocation field	-
provides relevantOccupation field to define occupation associated with jobPosting	provides position field which can be viewed as similar to it
provides responsibilities field	-
provides occupationalCategory field to categorise occupation but forces to use BLS O*NET-SOC taxonomy: http://www.onetcenter.org/taxonomy.html for occupationalCategory	provides jobCategory field to categorise job and jobCategory can be defined as we want
Pending fields educationalRequirements, estimatedSalary, experienceRequirements, occupationalCategory, qualifications, relevantOccupation, responsibilities, skills	No such pending or in-progress fields

Table 1: Comparison of two existing schemas.

4.1 Pending fields in schema.org

According to `schema.org` the pending extension is a staging area for work-in-progress terms which have yet to be accepted into the core vocabulary. Pending terms are subject to change and should be used with caution.

This means that, although `schema.org` has a larger vocabulary describing job posting, if there are any changes in the pending fields before they cease being tagged as ‘pending’, these changes would need to be applied to our project as well.

At this point, one can reflect on whether we need a larger vocabulary to define job postings. Choosing a larger vocabulary may give flexibility to describe the listing the way one wants. However, it requires us to accommodate ‘pending vocabulary’ which would require one to maintain and update.

4.2 Drawbacks of Individual Schemas

We analyzed the potential drawbacks of using each of the schemas. These are discussed below. With respect to the `schema.org/JobPosting` schema:

1. There are ‘pending’ fields which requires maintaining and update if necessary.
2. There is a provision to categorise occupation but forces to use BLS O*NET-SOC taxonomy: <http://www.onetcenter.org/taxonomy.html> for `occupationalCategory`.
3. Missing vocabulary to provide number of positions available for a job.

With respect to the `hr-xml/JobHeaderAndPostionHeader` schema:

1. Missing vocabulary to define hiring organisation or entity.
2. Missing vocabulary to provide job location information.
3. Missing vocabulary to provide job salary related information.
4. Missing vocabulary to provide specific required competence like that of education, experience, etc., however, all required competence can be dumped into element - ‘PositionSpecificCondition’ provided by hr-xml.

Bing supports `schema:JobPosting` markup, Google’s Job listing markup advises `schema.org` annotation, and job postings wishing to have a higher ranking in search results would follow Bing and Google’s guidelines. All these aspects imply that there are larger benefits of adopting `schema.org`. Based on a thorough analysis of existing schemas for job knowledge, and the overall benefits of using a `schema.org` listing, we decided to adopt `schema.org/JobPosting` as the core schema for the JKB within the scope of the DISKOW project. To overcome the limitations of missing attributes, we introduce them as additional attributes stemming from the DISKOW project. These two attributes are described in the following section.

5 Extending the Schema

5.1 Additional Attributes

Our investigation revealed that the JobPosting schema did not capture the desirable attributes to describe job knowledge in its entirety. Figure 3 shows some of the existing attributes within the `schema.org/JobPosting` schema.

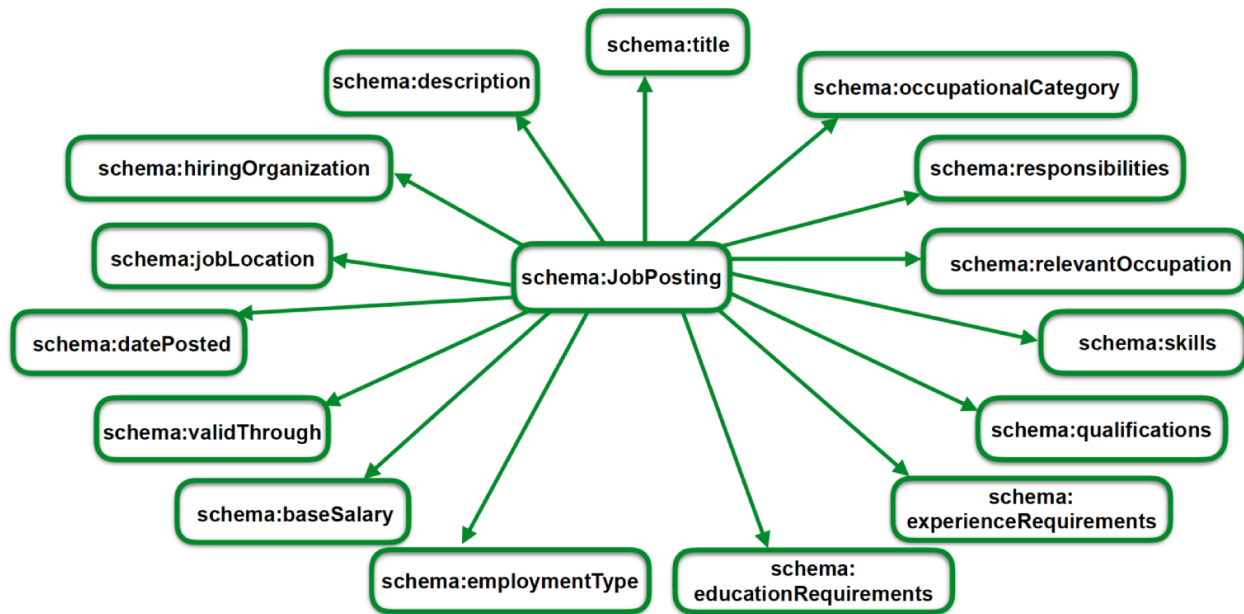


Figure 3: Some existing attributes within the `schema.org/JobPosting` schema.

For stakeholders who benefit from the DISKOW job knowledge base, there may be attributes beyond those described in the `schema.org/JobPosting` schema that are of interest. Due to this, we attempted to exhaustively identify attributes which could be essential. This led to the identification of two attributes: (i) a job category attribute that can be useful to classify jobs, and (ii) an attribute to describe the number of positions available with respect to a given job posting. At this point, it is useful to mention that every property in `schema.org` has an `additionalType` attribute which refers to a URL for an extended property. This URL can refer to `hr-xml jobCategory` (which is openly defined and allows us to exploit the hierarchical ESCO categorization) and `hr-xml positionHeader/RequestedPerson` (which will refer to the number of positions available).

5.2 Data Models for Competency Definitions

We compared the benefits of adopting either the *IEEE 1484.20.1* standard for reusable competency definitions or relying on *ESCO* (European Skills, Competences, Qualifications and Occupations), the European multilingual classification of Skills, Competences, Qualifications and Occupations². The IEEE Standard for Learning Technology-Data Model for Reusable Competency Definitions (*IEEE 1484.20.1*), defines a data model for describing, referencing, and sharing competency definitions, primarily in the context of online and distributed learning. This Standard provides a way to represent the key characteristics of a competency formally and independently of its use in any particular context. This enables interoperability among learning systems that deal with competency information by providing a means for them to refer to common definitions with common meanings. On the other hand, ESCO works as a dictionary, describing, identifying and classifying professional occupations, skills, and qualifications relevant for the EU labour market and education and training. Those concepts and the relationships between them can be understood by electronic systems, which allows different online platforms to use ESCO for services like matching job-seekers to jobs on the basis of their skills, suggesting trainings to people who want to reskill or upskill etc. ESCO is therefore very suitable for describing job knowledge within the scope of the DISKOW project, and we thereby choose to enrich the job knowledge within DISKOW by leveraging ESCO.

5.3 Semantic Enrichment of JKB Using ESCO

The aim of ESCO is to support job mobility across Europe and therefore a more integrated and efficient labour market, by offering a “common language” on occupations and skills that can be used by different stakeholders on employment and education and training topics. As a project that is focusing on the European Labor Market, it is of paramount importance to build on existing standards such as ESCO.

ESCO provides descriptions of 2942 occupations and 13.485 skills linked to these occupations, translated into 27 languages (all official EU languages plus Icelandic, Norwegian and Arabic). Linking the job knowledge available within the DISKOW JKB with ESCO descriptions can semantically enrich available job knowledge, and further enhance end-user scenarios by facilitating deeper, more complex analysis visualizations.

Semantically enriched job postings with the help of ESCO concepts and descriptions can help stakeholders better understand:

- what knowledge and skills are usually required when working in a specific occupation;
- what knowledge, skills and competences are obtained as a result of a specific qualification;
- what qualifications are demanded or often requested by employers from those searching for work in a specific occupation.

²<https://ec.europa.eu/esco/portal>

Next, we describe the affordances using the so called three hierarchical pillars of ESCO³.

5.3.1 Occupations

Each occupation corresponds to an occupational profile. The profiles contain an explanation of the occupation in the form of description, scope note and definition. Furthermore, they list the knowledge, skills and competences that experts considered relevant terminology for this occupation on a European scale. The following listing presents the levels of occupations.

- 0 Armed forces occupations
- 1 Managers
- 2 Professionals
- 3 Technicians and associate professionals
- 4 Clerical support workers
- 5 Service and sales workers
- 6 Skilled agricultural, forestry and fishery workers
- 7 Craft and related trades workers
- 8 Plant and machine operators and assemblers
- 9 Elementary occupations

In ESCO, each occupation is mapped to exactly one ISCO-08 code. ISCO-08 can therefore be used as a hierarchical structure for the occupations pillar. ISCO-08 provides the top four levels for the occupations pillar. ESCO occupations are located at level 5 and lower.

5.3.2 Skills/Competences

This pillar distinguishes between i) skill/competence concepts and ii) knowledge concepts by indicating the skill type. There is however no actual distinction between skills and competences. Each of these concepts comes with one preferred term and any number of non-preferred terms and hidden terms in each of the ESCO languages. It also includes an explanation of the concept in the form of description, scope note and definition. The skills pillar of ESCO does not contain a full hierarchical structure but is structured in four different manners:

- Through their relationship with occupations, i.e. by using occupational profiles as an entry point;
- In the part of the transversal knowledge, skills and competences through a skills hierarchy;
- Through relationships indicating how knowledge, skills and competences are relevant to other knowledge, skills and competences (in particular in cases of skill contextualisation);
- Through functional collections that allow to select subsets of the skills pillar.

The following listing presents the hierarchical structure of the skills/competences.

³<https://ec.europa.eu/esco/portal/occupation>

- Description
- Definition
- Scope notes
- Alternative label
- Skill type
- Skill reusability level
- Broader skills/competences
- Narrower skills/competences
- Essential skills and competences
- Essential Knowledge
- Optional skills and competences
- Optional Knowledge
- Essential skill/competence of
- Optional skill/competence of
- Regulatory aspect
- Status
- Replaces
- Replaced by
- Concept URI

5.3.3 Qualifications

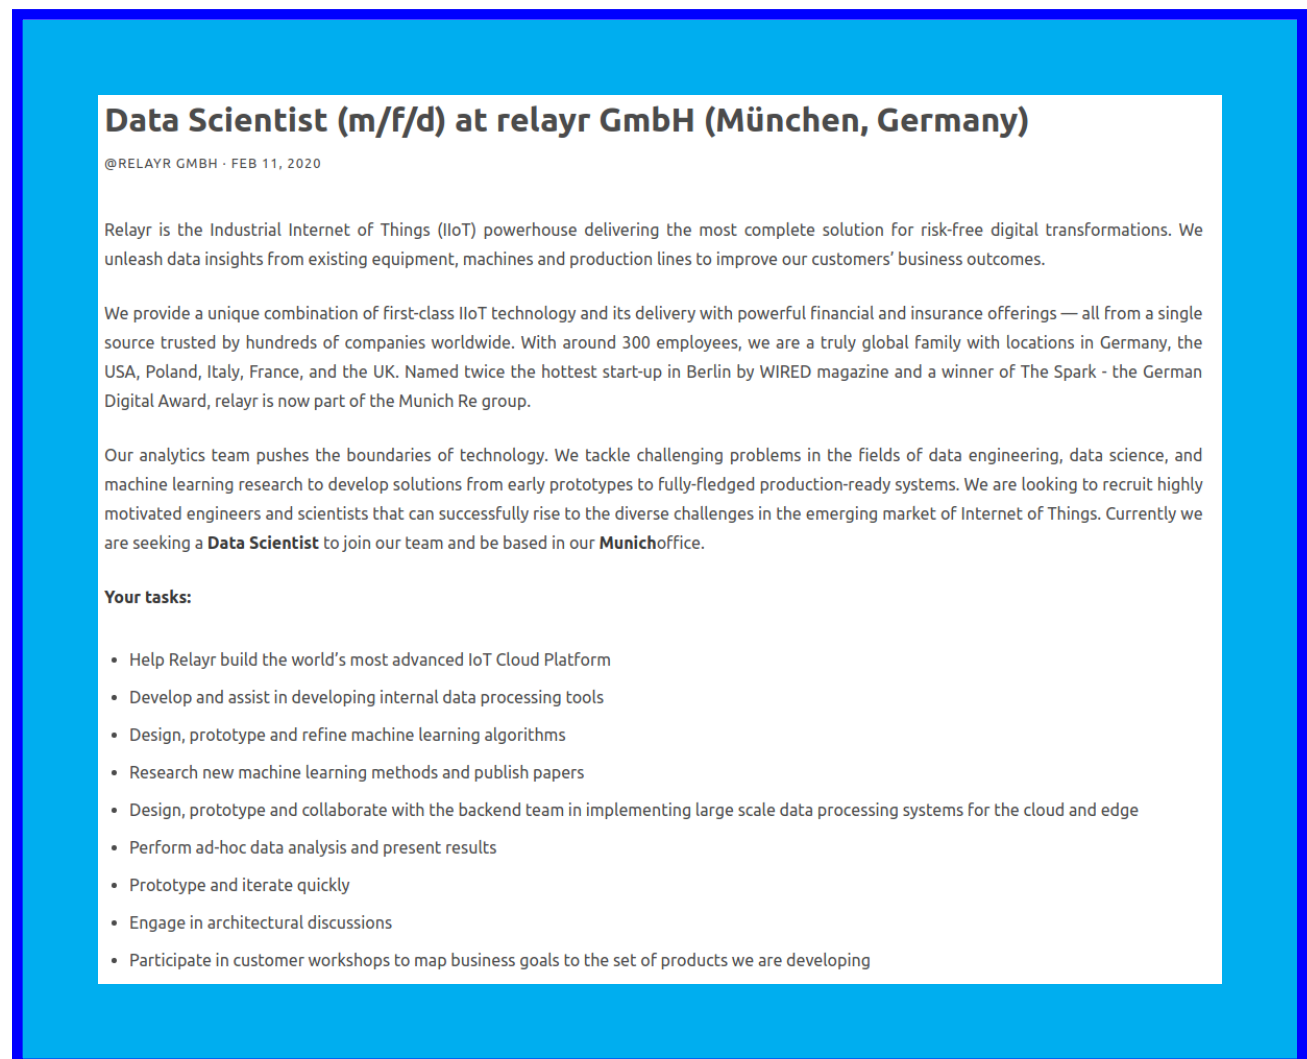
These are the formal outcome of an assessment and validation process which is obtained when a competent body determines that an individual has achieved learning outcomes to given standards. Qualifications displayed in ESCO come from databases of national qualifications that are owned and managed by the European Member States.

6 Application in ‘Data Science’ – An Example

The core objective of DISKOW is to mine job postings freely available on the Web and social media, and populate the JKB in accordance to the job knowledge schema. We can consider an example use case scenario of an individual on one hand, who is interested in applying for ‘Data Science’ jobs who uses the DISKOW JKB. On the other hand, within the scope of IO3 - *Job Knowledge Analysis Engine and Visualization Application Programming Interfaces (APIs)*, we will develop methods to extract job-related information from Web and social media data to populate the JKB and aid and such an end-user in finding appropriate job postings.

Figure 4 and 5 present an example job posting from `datasciencejobs.de`, a website that hosts job postings related to the field of Data Science in Germany. As can be observed in the figures, the job posting describes a Data Scientist role that is available at a company called Relayr in Munich, Germany. Data extraction methods developed within the scope of IO3 will extract relevant content from the job posting, and classification models will help map the relevant content to corresponding fields of the DISKOW JKB schema.

For example, in Figure 4 the tasks that are expected from the Data Scientist can be mapped onto the JKB schema attribute of `responsibilities`; the location of Munich, Germany can be mapped to the attribute `jobLocation`, etc. Similarly, as shown in Figure 5 the skills expected from the Data Scientist can be mapped onto the `skills` attribute within the JKB schema; the minimum experience expected can be mapped onto the `experienceRequirements` attribute, and so on.

The image is a screenshot of a job posting for a Data Scientist at relayr GmbH. The posting is set against a bright blue background. At the top, the title 'Data Scientist (m/f/d) at relayr GmbH (München, Germany)' is displayed in bold black text. Below the title, the text '@RELAYR GMBH · FEB 11, 2020' is shown in a smaller font. The main body of the posting consists of three paragraphs. The first paragraph describes relayr as an IIoT powerhouse. The second paragraph details the company's global presence and recent awards. The third paragraph describes the analytics team and the specific role being sought. Below the paragraphs, the section 'Your tasks:' is followed by a bulleted list of responsibilities, including building the IoT Cloud Platform, developing internal tools, designing machine learning algorithms, and participating in customer workshops. The entire job posting is enclosed within a white rectangular area with a thin black border, which is itself centered on the blue background.

Data Scientist (m/f/d) at relayr GmbH (München, Germany)

@RELAYR GMBH · FEB 11, 2020

Relayr is the Industrial Internet of Things (IIoT) powerhouse delivering the most complete solution for risk-free digital transformations. We unleash data insights from existing equipment, machines and production lines to improve our customers' business outcomes.

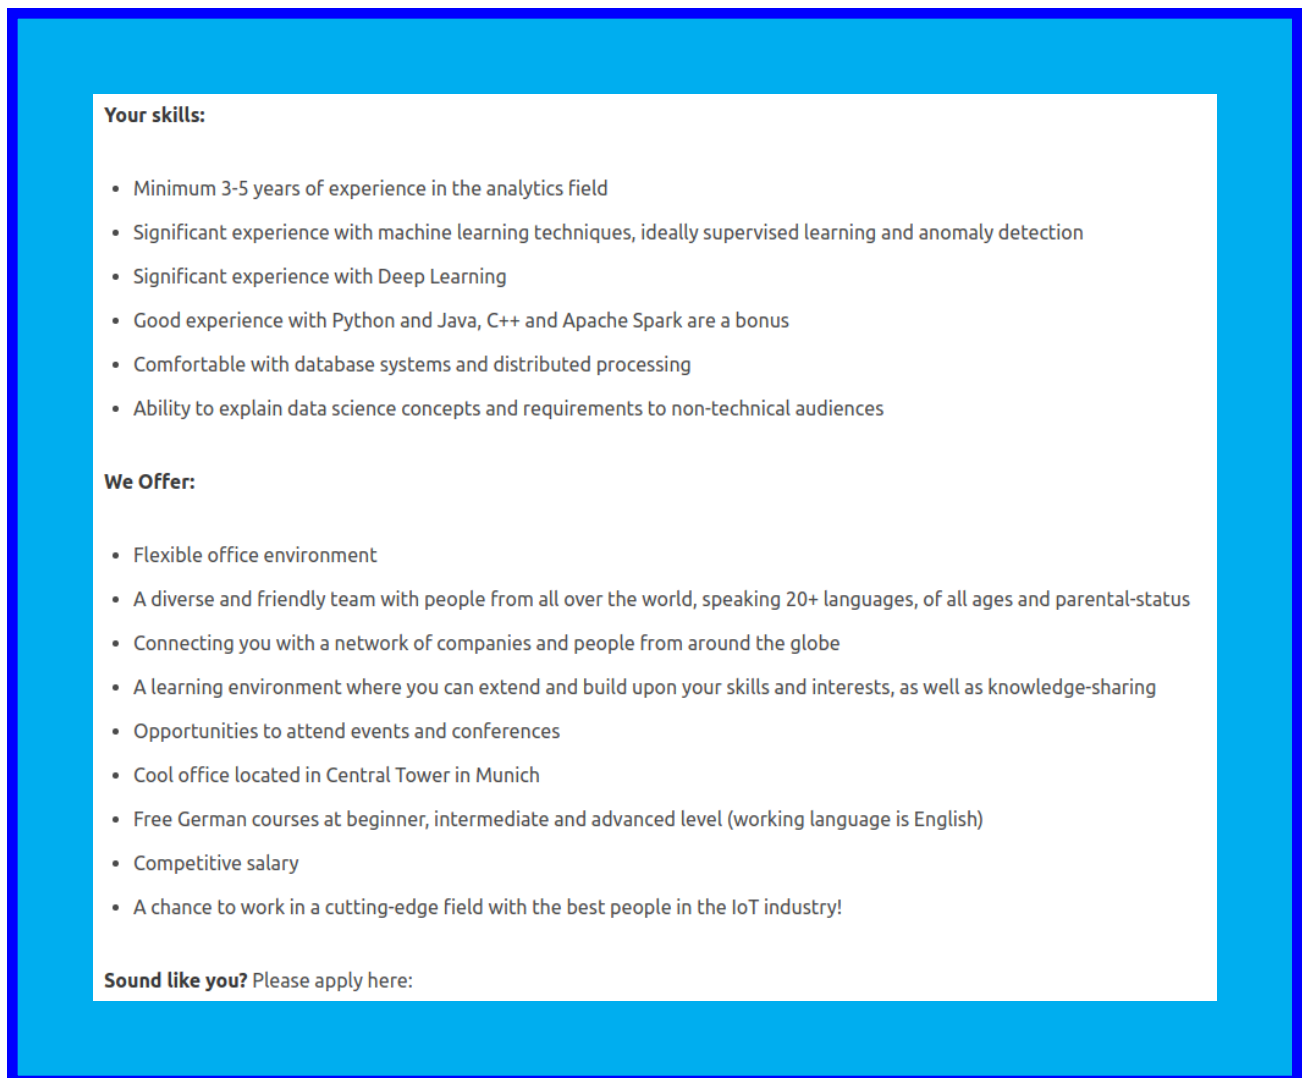
We provide a unique combination of first-class IIoT technology and its delivery with powerful financial and insurance offerings — all from a single source trusted by hundreds of companies worldwide. With around 300 employees, we are a truly global family with locations in Germany, the USA, Poland, Italy, France, and the UK. Named twice the hottest start-up in Berlin by WIRED magazine and a winner of The Spark - the German Digital Award, relayr is now part of the Munich Re group.

Our analytics team pushes the boundaries of technology. We tackle challenging problems in the fields of data engineering, data science, and machine learning research to develop solutions from early prototypes to fully-fledged production-ready systems. We are looking to recruit highly motivated engineers and scientists that can successfully rise to the diverse challenges in the emerging market of Internet of Things. Currently we are seeking a **Data Scientist** to join our team and be based in our **Munich** office.

Your tasks:

- Help Relayr build the world's most advanced IoT Cloud Platform
- Develop and assist in developing internal data processing tools
- Design, prototype and refine machine learning algorithms
- Research new machine learning methods and publish papers
- Design, prototype and collaborate with the backend team in implementing large scale data processing systems for the cloud and edge
- Perform ad-hoc data analysis and present results
- Prototype and iterate quickly
- Engage in architectural discussions
- Participate in customer workshops to map business goals to the set of products we are developing

Figure 4: Example of a 'Data Science' posting originating from Data Science Jobs.de (1/2).



Your skills:

- Minimum 3-5 years of experience in the analytics field
- Significant experience with machine learning techniques, ideally supervised learning and anomaly detection
- Significant experience with Deep Learning
- Good experience with Python and Java, C++ and Apache Spark are a bonus
- Comfortable with database systems and distributed processing
- Ability to explain data science concepts and requirements to non-technical audiences

We Offer:

- Flexible office environment
- A diverse and friendly team with people from all over the world, speaking 20+ languages, of all ages and parental-status
- Connecting you with a network of companies and people from around the globe
- A learning environment where you can extend and build upon your skills and interests, as well as knowledge-sharing
- Opportunities to attend events and conferences
- Cool office located in Central Tower in Munich
- Free German courses at beginner, intermediate and advanced level (working language is English)
- Competitive salary
- A chance to work in a cutting-edge field with the best people in the IoT industry!

Sound like you? Please apply here:

Figure 5: Example of a ‘Data Science’ posting originating from Data Science Jobs.de (2/2).

7 Conclusions and Future Work

In this report, we presented the progress made and the goals achieved within the scope of the DISKOW Intellectual Output 1 - *Modeling and Meta-Modeling of Job Knowledge for Labor Market*. With an aim to provide a standard and extensible JK model which can inform and structure the data within the DISKOW job knowledge base, we have adopted an existing ontology for job knowledge from `schema.org` and extended it with additional properties from `hr-xml` following the best practices laid out with respect to linked open data.

The schema thus lays the foundation for data extraction methods that will be developed through the course of the rest of the DISKOW project within IO3; inform the services and structures of the JKB within IO4; and finally facilitate a wide-range of application scenarios catering to a variety of stakeholders. In the future, this report will be updated to reflect changes in the `schema.org/JobPosting` schema if and when pending attributes are integrated.