Can we GLUE it? Extending an Interactive Adhesive Selector with Knowledge Graphs

Ioannis Dasoulas[1,2,3], Simon Vandevelde[1,2,3], Jeroen Jordens[3], Duo Yang[1,2,3], Xuemin Duan[1,2,3], Abdellatif Bey-Temsamani[3], Joost Vennekens[4], and Anastasia Dimou[1,2]

[1] KU Leuven, Dept. of Computer Science, Belgium
firstname.lastname@kuleuven.be
[2] Leuven.AI - KU Leuven Institute for AI, Leuven, Belgium
[3] Flanders Make, Belgium
firstname.lastname@flandersmake.be
[4] Vrije Universiteit Brussel, Belgium
firstname.lastname@vub.be

Abstract. Adhesives are commonly used across industries for a wide variety of applications thanks to their high strength, durability, and design versatility. The selection of appropriate adhesives is a complex task that typically requires specialized knowledge and access to detailed adhesive and material specifications. Traditionally, adhesive selection is done by experts based on experience and engineering knowledge. The Circular Joining Tech (CJT) department within the Flanders Make industrial research institute specializes in adhesive bonding, working with manufacturing companies on adhesive applications, such as glass bonding, automotive, and building construction. To support its adhesive experts in selecting the most appropriate adhesives faster and cheaper, we develop the **Semantic Adhesive Selector**, an interactive adhesive selector system with semantic web technologies and symbolic reasoning. To achieve this, we defined a novel adhesive ontology, the **GLUE Ontology**, that models adhesives and their properties, and created an adhesive knowledge graph, the **GLUE KG**, by applying this ontology to documented adhesive data within CJT. The adhesive recommendations are generated by the IDP-Z3 reasoning engine that is integrated with the GLUE KG. Our framework automates knowledge base construction for adhesives, enabling reasoning scalability, while knowledge graphs are also used to improve interactivity and data exploration. We validate our approach and discuss its impact within and beyond CJT, its adoption by Mavom and Novatech, two industrial adhesive suppliers.

Keywords: Adhesive Selection · Semantic Web · Symbolic AI.

1 Introduction

In recent years, the manufacturing industry has seen a steady rise in the use of adhesives for bonding operations [10,17]. Thanks to their favorable characteristics, adhesives are used in a range of industrial sectors, e.g., automotive, electronics and construction. However, selecting the most appropriate adhesive is a challenging task, as there are many requirements to be considered. For example, for an agriculture vehicle, we may want to glue an aluminum component to a thermoplastic housing, with a strength requirement of 12MPa. It should be able to withstand temperatures up to 40 °C, and have good vibration dampening properties together with a good resistance to water/sunlight/humidity. This is a complex list of requirements, making it difficult to find a suitable adhesive.

The Circular Joining (CJT) Tech department of Flanders Make employs adhesive experts to offer adhesive selection consultancy services to industrial patterns. For each case, experts draft a list of candidate adhesives and choose the most appropriate ones after conducting experiments. While testing is essential, the selection of suitable adhesives for the testing phase is time-consuming and labor-intensive, given the large number of adhesives available [10,16]. Also, adhesive information is often updated with new properties or measurements. Hence, informed adhesive selection with a scalable data integration process is needed.

In prior work, we developed the *Adhesive Selector* for CJT, a tool which uses a logical reasoning engine together with a formal knowledge base on adhesive selection to perform decision-making. The Adhesive Selector is user-centric, supporting experts during the selection process. Experts enter their requirements through an interactive user interface, which guides them towards a list of suitable adhesives by continuously updating the remaining possibilities.

However, adhesives could only be introduced manually since the Adhesive Selector did not include a structured way or automated knowledge base update process. Adhesives could only be added by knowledge engineers, as adhesive experts are not familiar with the FO(·) formalism [14] used in the Adhesive Selector. Moreover, the lack of externally defined semantics and definitions for the data hinders knowledge base scalability, while its FO(·) knowledge base lacks interoperability with internal or external manufacturing sources.

In this work, we scale the Adhesive Selector into the **Semantic Adhesive Selector (SAS)**, incorporating automated data integration pipelines, semantic interoperability and improved interaction techniques within CJT. To achieve this, we leverage semantic web technologies: We formalize the relations between adhesives, substrates and materials, and more than 100 of their attributes, in the **GLUE Ontology**[1]. We use prominent semantic web practices to create scalable data integration pipelines that transform existing recorded attributes and experimental measurements to the **GLUE KG**, an adhesive knowledge graph (KG), enabling richer queries, scalable reasoning, and interoperability with external datasets. We incorporate a data quality assessment module and Large Language Models (LLMs) into our data integration pipeline to minimize the required input from experts when updates are made or new adhesives are added to our knowledge base. Lastly, we integrate the previous Adhesive Selector with the GLUE KG by adding a layer of SPARQL query templates. We provide a video demonstration that showcases the usage of our system[2].

[1] https://github.com/dtai-kg/GLUE-Ontology/

[2] https://www.youtube.com/watch?v=B2kTc1L51Fs

We validate our approach's maintainability, usefulness, and scaling potential. To assess maintainability and usefulness we conducted a user study within CJT. To demonstrate scaling potential for generating recommendations, we show how filtering with SPARQL queries reduces the required reasoning time. The results demonstrate that our semantic web-based approach is very well-perceived by experts that maintain and use the SAS, and its potential in scaling reasoning capabilities of knowledge-based decision support systems.

A two-track valorization strategy is pursued for the uptake of our system. Firstly, CJT uses SAS internally to consult industrial partners for adhesive selection use cases. Secondly, we target third-party adhesive suppliers to provide our system as a service. As of March 2026, two adhesive suppliers, Mavom[3] and Novatech[4], which consult industrial partners and supply them with products related to bonding, sealing, and maintenance, initiated pilot deployments of our SAS to evaluate its applicability in industrial consulting for adhesives' recommendation. Another two companies have show interest to do this in the future. Beyond the adhesives domain, we investigate the reuse of the developed functionalities for other manufacturing domains, such as industrial machine operations.

The remainder of the paper is structured as follows: Section 2 elaborates on the adhesive selection use case at CJT and Section 3 on our methodology for building the SAS with semantic web technologies. Section 4 demonstrates our validation, and Section 5 discusses its impact and uptake. We comment on related work in Section 6 and conclude our work in Section 7.

2 Interactive Adhesive Selection at CJT

At CJT, **adhesive selection** is primarily done manually: given a description of an adhesion use case, an expert derives the relevant constraints, and manually looks through internal spreadsheets and technical datasheets of commercial adhesives to come up with a shortlist of adhesives. Then, the expert performs tests after which a candidate is selected. Due to the manual work involved, this process is costly and time-consuming. **Interactive adhesive selection** aims to assist the experts in finding an initial adhesives shortlist that satisfies a set of requirements, thus, reducing the time and cost of adhesive selection.

In this section, we go through our previous work in adhesive selection (Section 2.1) and the requirements we collected through experts to improve it (Section 2.2). We also list the data sources we use (Section 2.3) to build GLUE KG.

2.1 Preliminary Work: The Adhesive Selector

In prior work, we developed the **Adhesive Selector** [23,43], which relies on the IDP-Z3 reasoning engine [6] to drive its functionality. IDP-Z3 uses a formal description of a problem domain to perform various reasoning tasks, e.g.,

[3] `https://www.mavom.com`
[4] `https://novatech.eu`

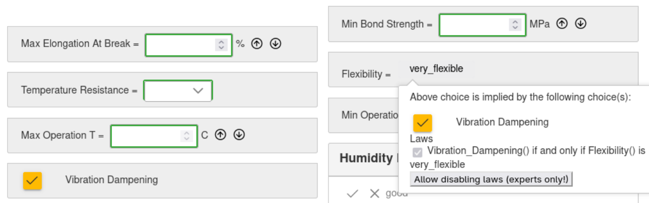

Fig. 1. Screenshot of some of the performance parameters in Adhesive Selector

finding possible solutions, deriving consequences, and generating explanations. Such a formal description for IDP-Z3 is frequently referred to as a knowledge base (KB), and is written in the FO(·) knowledge representation language [14] which extends First Order Logic. The KB for the Adhesive Selector was built in collaboration with CJT's experts through knowledge extraction workshops, in which we modeled adhesive domain knowledge that can be used to select adhesives. We refer to [43] for more information.

For user interaction, the Adhesive Selector uses the **Interactive Consultant** [5] interface. It is a generic interface for FO(·) KBs, acting as a front-end for IDP-Z3's inferences. In the Interactive Consultant, each symbol is represented using a *symbol tile* which can be used to set its value. For example, Figure 1 shows a part of the Adhesive Selector in which it indicated that the adhesive must be *Vibration Dampening*. When new information is added, the system automatically derives the consequences and shows the result. In the example, the Adhesive Selector derives that the adhesive must have a very high flexibility. As seen in the figure, we can also query the Interactive Consultant to explain why a value was derived. This explanation consists of the choices leading to it and the set of relevant formulas in the KB, leading to increased trust.

2.2 Expert Feedback and Requested Extensions

Through previous user studies [23], we found that the usability of the Adhesive Selector is promising: experts benefit from it, using it to affirm decisions and broaden their search. Yet, the it still had some crucial shortcomings: (i) the experts cannot add more adhesives as they are unfamiliar with the underlying FO(·) formalism; (ii) it can be slow as its reasoning does not scale well with an increasing number of adhesives; (iii) there is no interoperability between CJT data sources; and (iv) it lacks data exploration functionalities that can improve its interactivity. After receiving experts' feedback, we defined a set of requirements that our system's updated version should fulfill:

- **R1**: *User-friendly data updates.* The Adhesive Selector should provide a user-friendly way to update, add or remove information from its knowledge base.
- **R2**: *Automated updates.* The Adhesive Selector should automatically incorporate data updates to its knowledge base, without requiring knowledge engineers to assess quality of updated data.

Typical Properties*			
	TC-2002 A Resin	TC-2002 B Curative	Mixed
Appearance	Tan Paste	Gray Paste	Tan Paste
Viscosity, cP @ 25°C	600,000	325,000	500,000
Specific Gravity	1.71	1.24	1.67
Working Life, minutes @ 25°C	–	–	7-8

Fig. 2. Excerpt of Lord Cooltherm TC-2002 adhesive datasheet. Source: parker.com

- **R3**: *AI support for inferring missing information.* The Adhesive Selector should support experts in discovering missing information.
- **R4**: *Interactive visualizations.* The Adhesive Selector should support interactive visualizations of adhesives qualities and measurements for experts to further inspect and explore.
- **R5**: *Linking with external references and datasets.* The Adhesive Selector should link items with relevant external references, datasheets, and datasets.
- **R6**: *Performance improvement.* The Adhesive Selector should provide recommendations faster.

2.3 Data Sources for Interactive Adhesive Selection

We rely on four types of data sources to construct our knowledge base:

(i) ***Technical datasheets*** provided by adhesive manufacturers (e.g., Figure 2). These typically include quantitative specifications, e.g., application temperature and curing time. Also, they often contain qualitative notes about application guidelines for different material types and environmental conditions.

(ii) ***Knowledge derived from experts***. Their input is essential due to their experience and ability to translate datasheet guidelines to useful information, for instance regarding the resistance qualities of an adhesive.

(iii) ***Experimental data*** captured within CJT through laboratory testing of adhesives under controlled conditions. These tests provide empirical performance measurements, e.g., tensile strength in different temperature levels, validating or supplementing knowledge from technical datasheets and experts.

(iv) ***LLMs.*** Besides our above primary data sources, we also explore the usage of LLMs as data sources to infer missing information, validated by experts.

CJT maintains useful information derived from the first three data sources in internal spreadsheets, typically loosely structured. In the next section we explain how these data sources are used to create scalable knowledge acquisition pipelines to support adhesive selection.

3 The Semantic Adhesive Selector

We follow a modular and data-centric approach to build the **Semantic Adhesive Selector (SAS)** (Figure 3), enhancing adhesive selection using semantic web technologies, including OWL [30] for ontology formalization, RDF [9] for KG definition, and SPARQL [18] for KG querying, as well as prominent frameworks developed by the semantic web community which we describe in this section.

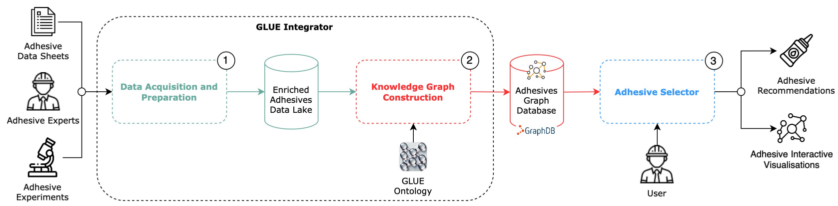

Fig. 3. Primary components of the Semantic Adhesive Selector.

The process begins with the knowledge acquisition from our primary data sources (Section 3.2). Acquired knowledge is stored in the Enriched Adhesives Data Lake, which includes diverse information regarding adhesives, materials, and substrates (Component ①). The next stage involves the construction (Component ②) of our adhesive KG, **GLUE KG** (see Section 3.3), using our **GLUE Ontology** (see Section 3.1), that formalizes the curated knowledge between adhesives, materials, and substrates. Components ① and ② comprise the **GLUE Integrator**, a framework tasked to automatically integrate available data to GLUE KG. GLUE Integrator is combined with the Adhesive Selector to form the SAS. The GLUE KG supports adhesive recommendation through reasoning, following the requirements that experts set (Component ③). It also provides interoperability with external resources, e.g., technical adhesive datasheets and datasets, as well as interactive graph visualizations. We provide a video demonstration that showcases the different SAS components and their usage[2].

3.1 The GLUE Ontology

To formally represent adhesives, their relationships with substrates, and their attributes, we developed the GLUE Ontology. The development is based on an in-depth analysis of CJT data sources as well as collaboration with experts for representing the relationships between concepts within the field of adhesives, following the Linked Open Terms (LOT) [34] methodology. The GLUE Ontology builds on the prominent high-level ontologies Schema.org[5] and PMD Ontology [3] to then define more specific adhesion-related entities and relationships. The ontology is available at GitHub[1], and published at `http://w3id.org/glue`.

Figure 4 showcases an overview of the primary ontology entities and relationships. Adhesives (`glue:Adhesive`) belong to an adhesive family (`glue:AdhesiveFamily`). Adhesive families are characterized by adhesion quality measures (`glue:AdhesionQuality`) depending on the material types (`glue:MaterialFamily`) of the substrates (`glue: Substrate`) that are adhered and by the applications (`glue:AdhesionApplication`) they are typically used in, across different types of manufacturing industries (`glue:Industry`).

The GLUE Ontology represents the attributes of adhesives, adhesive families, and materials documented by CJT throughout years of research and experiments, to support a comprehensive view of adhesives and substrates involved. We

[5] `https://schema.org`

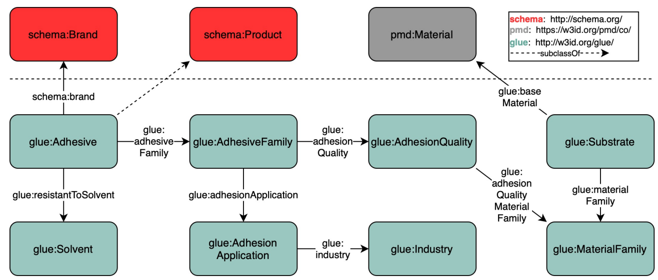

Fig. 4. Overview of GLUE Ontology main classes and properties.

define 34 data properties describing adhesives (e.g., maximum storage temperature, maximum pot life), 27 data properties describing adhesive families (e.g., water resistance, maximum detachment temperature), and 39 data properties describing substrates (e.g., maximum thermal conductivity, transparency).

The development of the GLUE Ontology is based on the LOT industrial-oriented methodology [34], which includes four major stages: Requirements Specification, Implementation, Publication, and Maintenance. We defined the ontology requirements in collaboration with CJT experts, based on its scope, i.e., the recommendation of adhesives for different application, environmental or material conditions. We collaborated with experts to document metadata for all properties, such as their definition, domain and range, and unit of measure. The ontology was designed in OWL [30] using Protegé [33]. To evaluate the ontology, we used SPARQL [18] queries and the OOPS! [35] validator. Ontology and LOT documents are maintained through GitHub[1], facilitating community suggestions.

3.2 GLUE Integrator: Data Acquisition and Preparation

An overview of our automated data acquisition and preparation pipeline is shown in Figure 5 (Component 1 of Figure 3). The process begins with knowledge acquisition, performed with experts. Experts contributed their insights as well as structured data extracted from technical datasheets and adhesion experiments. Together, these sources formed the initial input into the Adhesives Data Lake.

To build the initial data lake, we systematically communicated with experts to define all data fields within CJT data sources, their relationships, expected value types and ranges. This process ultimately enables data attributes to be mapped directly to our GLUE Ontology. We also used experts' feedback to infer missing values (e.g., the adhesive family of an adhesive) and resolve semantic inconsistencies (e.g., the labeling of the same adhesive with different labels).

To reduce required input from experts and ongoing reliance on knowledge engineers for routine data curation, we introduced a modular enrichment and data validation framework. This is essential due to the large amount of adhesives, substrates, and corresponding attributes. Two key modules enable this:

(i) ***Data Enrichment using LLMs***. To address missing categorical data, we integrated LLMs into our pipeline. LLMs are tasked with proposing values

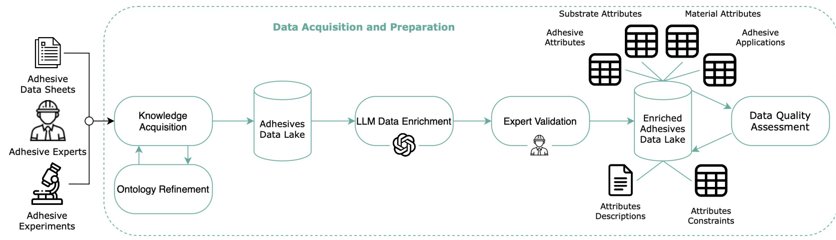

Fig. 5. Overview of data acquisition and preparation pipeline.

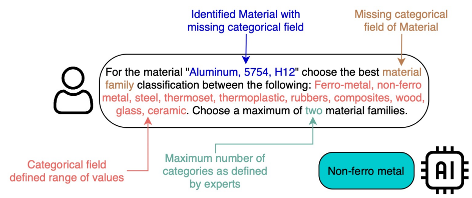

Fig. 6. High-level illustration of missing information inference using LLMs.

for incomplete entries by selecting from a predefined range of allowed values per attribute. To achieve that, we provide the LLM with the following: a) an entity that has missing categorical information; b) the name of the missing categorical information field; c) the range of possible values, as defined by experts; d) the maximum number of values the LLM should infer. This method allows the automated generation of prompts for missing categorical data. We comment on its results in the following section. All LLM-suggested values are subsequently subjected to expert validation before incorporation into the Enriched Adhesives Data Lake. Figure 6 is a high-level example of our automated LLM interactions to infer the missing material family of the *Aluminum, 5754, H12* material.

(ii) ***Data Quality Assessment***. After every update of the Enriched Adhesives Data Lake, a Data Quality Assessment module ensures data consistency. Updates are evaluated based on the defined value types and expected ranges of value for each attribute. Automatic repairs are performed for misalignments that can be fixed, such as the removal of text comments from data fields that expect numeric values (e.g., "100°C" is corrected to "100" using regular expressions). Misalignments that cannot be automatically repaired are flagged for expert review, such as the input of a value outside the expected value range for a corresponding field (e.g., "High temperature" is flagged for review if a numeric value is expected at this place). This module allows experts to maintain and update data autonomously, without requiring intervention from knowledge engineers, only by updating their internal spreadsheets.

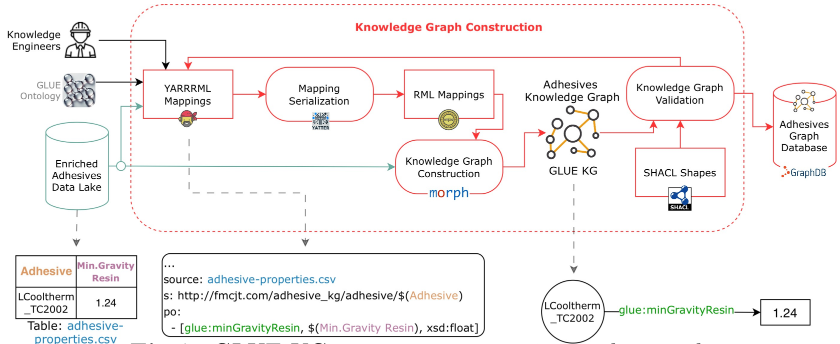

Fig. 7. GLUE KG construction overview and example.

3.3 GLUE Integrator: Adhesive Knowledge Graph Construction

We leverage prominent semantic web practices to streamline the construction of an adhesive KG, i.e. the GLUE KG, every time the data lake is updated. An overview of the process in showcased in Figure 7 (Component 2 of Figure 3).

The GLUE KG is constructed by applying the GLUE Ontology to the enriched CJT data lake. We use the RDF Mapping Language (RML) [15,22] to map data from the data lake to RDF triples, enabling declarative and standardized data transformations. We leverage the YARRRML [20] serialization using Yatter [21] to generate RML mapping rules in a user-friendly way. Morph-KGC [2,11] then transforms the original data to the GLUE KG, using the defined RML mappings. Once the KG is constructed, it is validated using SHACL [25]. We designed the shapes by first automatically extracting them through the GLUE Ontology using the Astrea [8] framework, and then manually reviewing and refining them. If errors are identified via SHACL shapes validation, the YARRRML mappings are refined accordingly. All shapes and coverage evaluation reports are included in our public repository. Figure 7 illustrates an example of KG construction by using the mappings defined by knowledge engineers in collaboration with adhesive experts. KG triples are stored in a GraphDB[6] instance.

Automated Data Updates. The GLUE Integrator enables the fully automated update of the GLUE KG whenever the adhesive experts update an adhesive attribute or add a new adhesive in the data lake, without requiring the interaction with semantic technologies or mapping languages. The GLUE Integrator applies the declarative mappings on the updated data lake to construct and publish a new version of the GLUE KG in GraphDB as a new named graph. In case of flagged errors, adhesive experts can again act independently and fix the corresponding data lake field, guided by the error message. Alternatively, they can contact the knowledge engineers that maintain the system to collaboratively solve issues. The intervention of knowledge engineers is only required if new adhesive/substrate attributes are added to the data lake; in which case they have to refine the ontology in collaboration with the experts.

[6] `https://graphdb.ontotext.com`

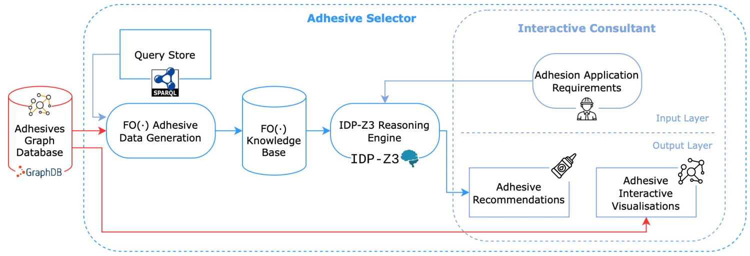

Fig. 8. Overview of the integration of the GLUE KG with the Adhesive Selector.

3.4 KG integration in the Adhesive Selector

We integrate the Adhesive Selector with the GLUE KG, using a layer of template SPARQL queries as seen in Figure 8 (Component 3 of Figure 3). Experts select their application requirements through the Interactive Consultant (see Section 2.1). The corresponding data are fetched using SPARQL queries, and are then formatted as part of the FO(·) knowledge base. We define "template queries" that can be used to automatically generate queries for specific data, based on expert inputs. We distinguish between two such templates, namely one for the types (i.e., the set of adhesives or adhesive families), and one for the properties (i.e., the strength of each adhesive or adhesive family). The generated FO(·) knowledge base is used by the IDP-Z3 reasoning engine to recommend adhesives, as described in Section 2.1.

To provide interactive data exploration, we link adhesives, adhesive families, and substrates to interactive GLUE KG visualizations which we generate with GraphDB. These enable the exploration of recorded data instances, attributes, and relationships without manually searching for them. An example is presented in Figure 9 for the adhesive family Methyl Methacrylate (MMA).

4 Validation

We validate the SAS via user-evaluation. Note that the accuracy and usefulness of the recommendation process itself has already been validated in our previous work [43]. In this paper, this reasoning process remains unchanged. Instead, we focus on assessing the newly introduced aspects, namely the semantic web-based data update autonomy introduced, the LLM-based data enrichment module, the interactive visualizations, and the interoperability mechanisms through GLUE KG. We also demonstrate an example of how GLUE KG is used to fetch data, and show how KG-based filtering accelerates recommendations.

User evaluation. To assess the effectiveness of the SAS, we conducted a formative user study with three CJT experts that have used the Adhesive Selector in the past. Since experts both update the SAS data lake and use the SAS for decision support, we evaluate the usefulness both for features that aim to improve system

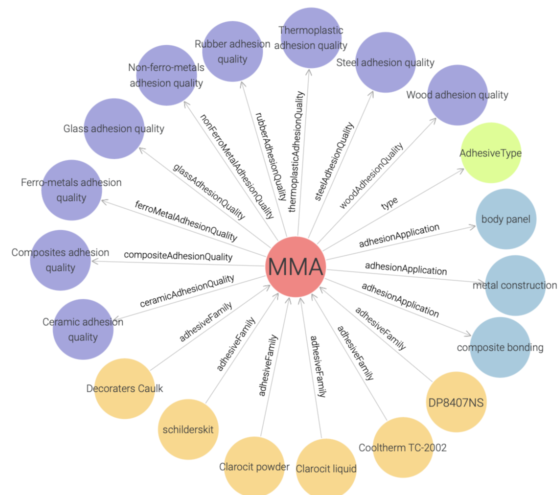

Fig. 9. Example of interactive visualization for the Methyl Methacrylate (MMA) adhesive family, integrated from GraphDB for data exploration by SAS users.

maintenance and user experience. Experts were asked to evaluate the usefulness of each feature on a scale of $1 - 5$ and share their feedback:

- ***Data update autonomy*** (R1, R2). Experts found that autonomously updating the data of the SAS is very useful for maintaining and extending the system. Experts can achieve this now through the internal spreadsheets they are already familiar with, without having to communicate with knowledge engineers or manually update the FO($\cdot$) knowledge base. The built-in Data Quality Assessment guides experts to correct errors they make on update attempts. Experts rated the usefulness of this feature at 5/5.
- ***LLM-Based data enrichment*** (R3). Experts were shown the results of LLM-based inference of missing categorical information and were asked to evaluate the model's response. In this evaluation, the GPT-4o[7] model was used for inferring missing information for three categorical substrate fields, namely material family, transparency, and magnetic type, and two adhesive family categorical fields, namely relative chemical resistance and adhesion quality with non-ferro metals. The experts assessed that 81.19% (116/143) of LLM generated answers were entirely correct, showing strong performance for material classification (family: 100%, transparency: 87.5%, magnetic type: 100%), but lower accuracy for specific adhesives (chemical resistance: 68.8%, non-ferro metals adhesion: 50%). In other words, while the LLM reliably categorizes general material properties, it struggles with highly specific characteristics of adhesives without access to deeper domain knowledge. Still, the experts found this feature moderately useful, though

[7] `https://openai.com/index/hello-gpt-4o/`, using default settings

Listing 1.1. SPARQL query to retrieve adhesives meeting specific requirements

```
PREFIX glue: <http://w3id.org/glue/>
SELECT ?adh ?fam ?strength ?waterq ?q1 ?q2 WHERE {{
    ?adh a glue:Adhesive.
    ?adh glue:adhesiveFamily ?fam.
    OPTIONAL {?fam glue:waterResistance ?waterq}
    OPTIONAL {?fam glue:maxShearStrength ?strength}
    OPTIONAL { ?subs1q glue:adhesionQualityMaterialFamily
        <http://w3id.org/fm_adhesive_kg/material_family/non_ferro_metal>.
        ?fam glue:nonFerroMetalAdhesionQuality ?subs1q.
        BIND(?subs1q AS ?q1source)}
    OPTIONAL {?subs2q glue:adhesionQualityMaterialFamily
        <http://w3id.org/fm_adhesive_kg/material_family/thermoplastic>.
        ?fam glue:thermoplasticAdhesionQuality ?subs2q.
        BIND(?subs2q AS ?q2source)}
    OPTIONAL {?q1source glue:adhesionQualityEstimation ?q1.}
    OPTIONAL {?q2source glue:adhesionQualityEstimation ?q2.}
    FILTER(!BOUND(?strength) ||  ?strength > 12)
    FILTER(!BOUND(?waterq) || ?waterq = "good")
    FILTER(!BOUND(?q1) || ?q1 = "good" || ?q1 = "treatment")
    FILTER(!BOUND(?q2) || ?q2 = "good" || ?q2 = "treatment")}}
```

human validation is still necessary during the process. This feature was rated at 3.33/5.

- ***Interactive visualizations*** (R4). Experts found that interactive GLUE KG visualizations are promising for data exploration and discovery of connections. However, they commented that data exploration may be hindered for node neighborhoods that are too crowded, hence rating it at 3.5/5.
- ***Linking with external references and datasets*** (R5). Experts found that linking with references, e.g., the adhesive technical datasheets and the potential interoperability with external datasets from the material science domain, is an important feature for data exploration. It was rated at 4.67/5.

Use case example. To demonstrate adhesive selection using the GLUE KG, we consider a use case with the following three requirements:

1. Glue **Aluminum** (non-ferro metal) to **POM** (thermoplastic).
2. Minimum bonding strength of **12MPa**.
3. **Good** water resistance.

Listing 1.1 shows the SPARQL query that corresponds to our adhesion use case. Using FILTERs, we filter out the adhesives that do not meet the hard requirements specified. This remaining list is formulated as $\mathrm{FO}(\cdot)$ to enable reasoning with the IDP-Z3 engine for adhesives recommendation, using its knowledge base, as discussed in Section 2.1.

Speeding up the Adhesive Selector. In the Adhesive Selector's previous version, reasoning was performed over the entirety of the $\mathrm{FO}(\cdot)$ knowledge base to generate adhesive recommendations, becoming slower with the increasing number of adhesives. To alleviate this, we use SPARQL queries to filter out adhesives not relevant for a use case. This allows the SAS to exclude adhesives that do not meet a hard requirement, thus reducing reasoning time (R6).

Table 1. Adhesive recommendation reasoning time using the Adhesive Selector and the Semantic Adhesive Selector, over 10 common adhesive selection use cases. The minimum and maximum times correspond to the same use cases for both methods.

Reasoning Method	Reasoning Time (s)		
	Min.	Max.	Avg.
IDP-Z3 [6] (Adhesive Selector)	0.88	2.11	1.39
IDP-Z3 w/ KG Filtering (SAS)	**0.81**(-8.03%)	**1.43**(-32.28%)	**1.13**(-16.89%)

Table 2. Full experiment results comparing the recommendation reasoning time with and without SAS filter. In total, our KG contained 54 adhesives.

Experiment	Avg. time w/o filter	Avg. time. w/ filter	Avg. speed increase	Avg. stdev increase	nb of adhesives remaining
1	1.64 ± 0.84	1.09 ± 0.48	33.36%	43.24%	49
2	1.17 ± 0.63	1.12 ± 0.65	4.69%	-3.22%	49
3	1.31 ± 0.77	1.17 ± 0.67	10.45%	13.02%	32
4	1.44 ± 0.93	1.11 ± 0.52	23.25%	43.35%	21
5	0.88 ± 0.33	0.81 ± 0.39	8.03%	-20.64%	22
6	1.48 ± 0.88	1.32 ± 0.77	10.69%	12.76%	32
7	0.99 ± 0.43	0.75 ± 0.27	24.41%	38.69%	21
8	2.11 ± 2.19	1.43 ± 0.82	32.28%	62.54%	47
9	1.25 ± 0.74	1.35 ± 0.86	-7.67&	16.02%	31
10	1.67 ± 1.94	1.18 ± 0.82	29.47%	57.65%	33
Avg	1.39 ± 0.97	1.13 ± 0.63	16.89%	35.41%	33.7

To quantify the speed increase, we defined 10 common adhesive use cases and their corresponding requirements: cases 1 and 2 defined four requirements, whereas the other cases defined six. We measured the required time to generate recommendations with and without using the GLUE KG to filter out adhesives that do not meet hard parameters, using the same hardware. To account for randomness in the solving speed, each experiment was repeated four times, each time with a random order for the requirements. We have left out the SPARQL querying and data conversion, as its total time remains consistently small and did not significantly affect performance (averaging 40ms). The total experiment KG consisted of 3939 triples in total.

Table 1 contains the summarized results, showing that the KG-based filtering achieves an average recommendation speed increase of 16.89%. The full results of the experiments are shown in Table 2, including the standard deviations, and the remaining number of adhesives after filtering. On average, filtering reduced the number of adhesives by 38% in our experiments, which explains the significant efficiency gain for the recommendation. Interestingly, for experiment 9 the overall reasoning time actually *increased*, showing that reducing the number of adhesives does not always result in speed gain, a phenomenon stemming from the way IDP-Z3 works internally.

5 Impact and Uptake

The user study confirmed the usefulness of the newly developed features for the Semantic Adhesive Selector (SAS). In particular, our automated knowledge base update approach transforms how adhesives' knowledge is applied within CJT, as it eliminates the need for direct involvement of knowledge engineers. This increased autonomy alleviates a key bottleneck in scaling the available data lake, enhancing SAS and thereby enabling its potential to be established as an enterprise system. Furthermore, the introduction of visual data exploration and interoperability with external references improves its transparency, while semantic search reduces the time required to provide recommendations.

Uptake. A dual track valorization strategy for the SAS is pursued in line with CJT's and Flanders Make's broader commitment to promote knowledge transfer to industrial practice. On the one hand, CJT is internally uptaking SAS aiming at expanding its application from internal use towards broader industrial adoption. CJT experts use the system to provide adhesive selection support in industrial projects in which the department is involved, enabling the collection of feedback to refine the system's usability, ontology coverage, and data integration pipeline.

On the other hand, we also focus on offering the SAS as a service to industrial partners, including adhesive manufacturers, suppliers, and consultancies. Close to twenty companies have been identified as potential adopters. In March 2026, two major adhesive supplying companies, Mavom[3] and Novatech[4], initiated pilot deployments of SAS to evaluate its applicability in industrial consulting for adhesives' recommendation. Both companies develop solutions and provide technical support related to adhesive bonding on large-scale industrial projects.

The knowledge acquisition process has started with both partners. Through these pilot deployments, both companies provide proprietary product datasets, expert decision rules, and industrial application cases. From a technical perspective, this deployment involves adjusting SAS's components for each, to capture company-specific data attributes and requirements (e.g., extending our base GLUE Ontology or customizing parts of the GLUE Integrator), while reusing the foundational components of our system (e.g., primary reasoning rules, ontology classes and properties, and KG construction pipeline). Thanks to these two collaborations, SAS has entered the testing phase in applications related to materials gluing on walls, safety rails, and steel plates on boats, marking its first steps toward broader industrial adoption, with another two companies having declared interest to adopt it. This process also acts as a feedback loop for improving and scaling our base SAS version. Lessons learned and best practices from our industrial integrations inform subsequent iterations of SAS, enabling targeted improvements in its functionality, robustness, and usability. These insights guide also its adoption to diverse industrial domains.

Extensibility in other manufacturing domains. During discussions and demonstrations with manufacturing experts, it became clear that the architectural approach we used leveraging semantic web technologies holds promise beyond adhesives. Similar systems could support selection tasks in related areas, such as selecting manufacturing components or configuring production processes. The core methodology, semantic modeling, linked data integration, LLM-assisted enrichment with human in the loop, and controlled reasoning after filtering offers a generalizable decision support framework. This positions our approach as a potential foundation for broader industrial applications, where structured knowledge is central for recommendation systems. We are actively working towards a similar approach for decision-support in industrial machines operation.

6 Related Work

Past works on adhesive selection have focused on decision guides or expert systems, without leveraging semantic web technologies. In parallel, semantic web technologies have been increasingly adopted in manufacturing for decision support, with the application of adhesive selection remaining largely unexplored.

Adhesive Selection. Existing software support for adhesive selection falls in website guides [10] or expert systems [26,31,32,38,41]. The former are typically limited in their data and still require manual work for decision-making. The latter are typically based on decision trees or selection tables, using pre-defined rules, but they are limited in their interpretability, usability, and user-friendliness [23].

Semantic Web for Decision Support in Manufacturing. Previous works showcased the advantages of applying semantic web technologies in the manufacturing domain to support decision-making and facilitate automation [37,44]. **ManuHub** [4] and **Lu and Xu** [27] propose semantic web-based frameworks to generate manufacturing service plans using semantic similarity. The **ExeKG framework** [24,45] enables operators to construct and visualize ML manufacturing workflows based on ontologies. **Mabkhot et al.** [28] supports decision-making for manufacturing process selection by using an ontology to apply automatic reasoning and similarity retrieval. **Psarommatis and Kiritsis** [36] automate decision-making and planning when defects are detected in manufacturing processes. **Uddin et al.** [42] assist in real-time decision-making to optimize performance indicators of manufacturing systems, leveraging ontology-based knowledge acquisition. **Singh et al.** [40] provide an ontology-based decision support system for supporting decision-making in supply chain management.

Manufacturing and Materials Ontologies. Different ontologies have been proposed to describe manufacturing knowledge and knowledge in the field of material sciences. At a broader, more general level, these ontologies include **QUDT** [1] for representing quantities, units, and dimensions; **DOLCE+DnS Ultralite** [39], an upper-level ontology, for modeling entities, processes, and roles; and **Procedural Knowledge Ontology** [7], which can be tailored to the manufacturing domain. **EMMO** (Elementary Multiperspective Material Ontology) [13] is an upper-level ontology that represents concepts in physics, chemistry, and materials science. **ChEBI** (Chemical Entities of Biological Interest) [12,29] provides a structured ontology and dictionary for representing molecular entities, covering atoms, molecules, ions, radicals, and related chemical structures. **PMD Ontology** [3] builds on these foundations (EMMO and ChEBI), covering more domain-neutral concepts. **eClassOwl** [19] offers a standardized representation of products and services to facilitate interoperability in electronic business, but with less emphasis on manufacturing applications. However, existing ontologies do not provide a comprehensive representation for adhesives, especially for the selection process that considers environmental conditions, substrate attributes, adhesive resistance and other application-specific requirements.

7 Conclusion

In this work, we described the SAS, our solution to extend the Adhesive Selector using semantic web technologies for the CJT. We created GLUE Ontology for modeling adhesive knowledge and built a scalable data integration pipeline that transforms curated data to an adhesive KG, GLUE KG, using standard semantic web practices. We experimented with LLM-based data enrichment and data quality assessment methods to reduce the dependence of CJT experts to knowledge engineers and enable experts to maintain and use the SAS independently. The GLUE KG enables interoperability, interactive visualizations, and complex querying to reduce reasoning time for recommendations. Our newly developed features were well received by the experts and industry contacts, showcasing how semantic web technologies can be integrated to improve existing systems and leading to formal industrial pilots with two major adhesive suppliers. In the future, we plan to focus on SAS's adoption in industry, directly incorporating and evaluating it on more use cases, gathering feedback to further improve and extend its capabilities, and facilitating collaboration with experts by adopting modern human-in-the-loop ontology engineering automated approaches.

Supplemental Material Statement: The GLUE Ontology is openly available at `https://github.com/dtai-kg/GLUE-Ontology` and published at `http://w3id.org/glue`. The implementation of the Semantic Adhesive Selector and curated data include confidential work and cannot be shared. A video demonstration is provided at `https://www.youtube.com/watch?v=B2kTc1L51Fs`.

Acknowledgments. This research was partially supported by Flanders Make (REXPEK project), the strategic research centre for the manufacturing industry, and the Flemish Government under the "Onderzoeksprogramma Artificiële Intelligentie (AI) Vlaanderen" program. Illustrations were designed using resources from `flaticon.com`.

Declaration of use of Generative AI

During the preparation of this work, the authors used generative AI to help with grammar and to perform spelling check. The authors carefully reviewed and edited the content as needed and take full responsibility for the publication's content.

References

1. FAIRsharing.org: QUDT; Quantities, Units, Dimensions and Types (2025). https://doi.org/10.25504/FAIRsharing.d3pqw7, last Edited: Tuesday, March 25th 2025
2. Arenas-Guerrero, J., Chaves-Fraga, D., Toledo, J., Pérez, M.S., Corcho, O.: Morph-KGC: Scalable Knowledge Graph Materialization with Mapping Partitions. Semantic Web (2022). https://doi.org/10.3233/SW-223135
3. Bayerlein, B., Schilling, M., Birkholz, H., Jung, M., Waitelonis, J., Mädler, L., Sack, H.: Pmd core ontology: Achieving semantic interoperability in materials science. Materials Design **237**, 112603 (Jan 2024). https://doi.org/10.1016/j.matdes.2023.112603
4. Cai, M., Zhang, W.Y., Zhang, K.: Manuhub: A semantic web system for ontology-based service management in distributed manufacturing environments. IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans **41**(3), 574–582 (2011). https://doi.org/10.1109/TSMCA.2010.2076395
5. Carbonnelle, P., Aerts, B., Deryck, M., Vennekens, J., Denecker, M.: An interactive consultant. In: Proceedings of the 31st Benelux Conference on Artificial Intelligence (BNAIC 2019) and the 28th Belgian Dutch Conference on Machine Learning (Benelearn 2019), Brussels, Belgium, November 6-8, 2019. CEUR Workshop Proceedings, vol. 2491. CEUR-WS.org (2019), http://ceur-ws.org/Vol-2491/demo45.pdf
6. Carbonnelle, P., Vandevelde, S., Vennekens, J., Denecker, M.: IDP-Z3: A reasoning engine for FO(.) (2022). https://doi.org/10.48550/arXiv.2202.00343
7. Carriero, V.A., Scrocca, M., Baroni, I., Azzini, A., Celino, I.: Procedural knowledge ontology (pko) (2025), https://arxiv.org/abs/2503.20634
8. Cimmino, A., Fernández-Izquierdo, A., García-Castro, R.: Astrea: Automatic generation of shacl shapes from ontologies. In: The Semantic Web: 17th International Conference, ESWC 2020, Heraklion, Crete, Greece, May 31–June 4, 2020, Proceedings. p. 497–513. Springer-Verlag, Berlin, Heidelberg (2020). https://doi.org/10.1007/978-3-030-49461-2_29, https://doi.org/10.1007/978-3-030-49461-2_29
9. Cyganiak, R., Wood, D., Lanthaler, M.: RDF 1.1 Concepts and Abstract Syntax. Recommendation, World Wide Web Consortium (W3C) (Feb 2014), http://www.w3.org/TR/rdf11-concepts/
10. da Silva, L.F.M., Öchsner, A., Adams, R.D.: Handbook of Adhesion Technology, vol. 1. Springer (2018)
11. Dasoulas, I., Chaves-Fraga, D., Garijo, D., Dimou, A.: Declarative RDF construction from in-memory data structures with RML. In: Proceedings of the 4th International Workshop on Knowledge Graph Construction co-located with 20th Extended Semantic Web Conference ESWC 2023. vol. 1613, p. 0073 (2023)
12. Degtyarenko, K., de Matos, P., Ennis, M., Hastings, J., Zbinden, M., McNaught, A., Alcantara, R., Darsow, M., Guedj, M., Ashburner, M.: Chebi: a database and ontology for chemical entities of biological interest. Nucleic Acids Research **36**, D344–D350 (Dec 2007). https://doi.org/10.1093/nar/gkm791
13. Del Nostro, P., Friis, J., Ghedini, E., Goldbeck, G., Holtz, O., Roscioni, O., Zaccarini, F., Toti, D.: Elementary multiperspective material ontology: Leveraging perspectives via a showcase of emmo-based domain and application ontologies. In: Proceedings of the 16th International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management. p. 135–142. SCITEPRESS - Science and Technology Publications (2024). https://doi.org/10.5220/0012910200003838

14. Denecker, M.: Extending classical logic with inductive definitions. In: Lloyd, J., Dahl, V., Furbach, U., Kerber, M., Lau, K.K., Palamidessi, C., Pereira, L.M., Sagiv, Y., Stuckey, P.J. (eds.) Computational Logic — CL 2000. pp. 703–717. Springer Berlin Heidelberg, Berlin, Heidelberg (2000)
15. Dimou, A., Vander Sande, M., Colpaert, P., Verborgh, R., Mannens, E., Van de Walle, R.: RML: A Generic Language for Integrated RDF Mappings of Heterogeneous Data. In: Bizer, C., Heath, T., Auer, S., Berners-Lee, T. (eds.) Proceedings of the 7th Workshop on Linked Data on the Web. CEUR Workshop Proceedings, vol. 1184. CEUR (2014)
16. Ewen J.C., K.: 1 - Key issues in selecting the right adhesive. In: Dillard, D.A. (ed.) Advances in Structural Adhesive Bonding, pp. 3–19. Woodhead Publishing in Materials, Woodhead Publishing (2010). https://doi.org/10.1533/9781845698058.1.3, https://www.sciencedirect.com/science/article/pii/B9781845694357500018
17. Habenicht, G.: Applied adhesive bonding: a practical guide for flawless results. John Wiley & Sons (2008)
18. Harris, S., Seaborne, A.: SPARQL 1.1 Query Language. Recommendation, World Wide Web Consortium (W3C) (Mar 2013), https://www.w3.org/TR/sparql11-query/
19. Hepp, M.: GoodRelations: An Ontology for Describing Products and Services Offers on the Web, p. 329–346. Springer Berlin Heidelberg (2008). https://doi.org/10.1007/978-3-540-87696-0_29
20. Heyvaert, P., De Meester, B., Dimou, A., Verborgh, R.: Declarative Rules for Linked Data Generation at your Fingertips! In: The Semantic Web: ESWC 2018 Satellite Events: ESWC 2018 Satellite Events, Heraklion, Crete, Greece, June 3-7, 2018, Revised Selected Papers 15 (2018)
21. Iglesias-Molina, A., Chaves-Fraga, D., Dasoulas, I., Dimou, A.: Human-Friendly RDF Graph Construction: Which One Do You Chose? In: Garrigós, I., Murillo Rodríguez, J.M., Wimmer, M. (eds.) Web Engineering. pp. 262–277. Springer Nature Switzerland, Cham (2023). https://doi.org/10.1007/978-3-031-34444-2_19
22. Iglesias-Molina, A., Van Assche, D., Arenas-Guerrero, J., De Meester, B., Debruyne, C., Jozashoori, S., Maria, P., Michel, F., Chaves-Fraga, D., Dimou, A.: The RML Ontology: A Community-Driven Modular Redesign After a Decade of Experience in Mapping Heterogeneous Data to RDF. In: International Semantic Web Conference. pp. 152–175. Springer (2023). https://doi.org/10.1007/978-3-031-47243-5_9
23. Jordens, J., Vandevelde, S., Van Doninck, B., Witters, M., Vennekens, J.: Adhesive selection via an interactive, user-friendly system based on symbolic AI. Procedia CIRP, Elsevier (2022)
24. Klironomos, A., Zhou, B., Tan, Z., Zheng, Z., Mohamed, G.E., Paulheim, H., Kharlamov, E.: Exekglib: knowledge graphs-empowered machine learning analytics. In: ESWC. pp. 123–127. Springer (2023)
25. Knublauch, H., Kontokostas, D.: Shapes Constraint Language (SHACL). Recommendation, World Wide Web Consortium (W3C) (2017), https://www.w3.org/TR/shacl/
26. Lammel, C., Dilger, K.: Software for a RULE-BASED ADHESIVE-SELECTION SYSTEM. Adhesives & Sealants Industry **9**(5), 42–43 (2002)
27. Lu, Y., Xu, X.: A semantic web-based framework for service composition in a cloud manufacturing environment. Journal of Manufacturing Systems **42**, 69–81 (2017). https://doi.org/https://doi.org/10.1016/j.jmsy.2016.11.004, https://www.sciencedirect.com/science/article/pii/S0278612516300838

28. M. Mabkhot, M., Al-Samhan, A., Hidri, L.: An Ontology-Enabled Case-Based Reasoning Decision Support System for Manufacturing Process Selection. Advances in Materials Science and Engineering **2019**, 1–18 (08 2019). https://doi.org/10.1155/2019/2505183
29. de Matos, P., Alcántara, R., Dekker, A., Ennis, M., Hastings, J., Haug, K., Spiteri, I., Turner, S., Steinbeck, C.: Chemical entities of biological interest: an update. Nucleic Acids Research **38**, D249–D254 (Oct 2009). https://doi.org/10.1093/nar/gkp886
30. McGuinness, D.L., Van Harmelen, F., et al.: Owl web ontology language overview. W3C Recommendation (2004)
31. Meyler, K.L., Brescia, J.A.: Design of a computer expert system for adhesive selection using artificial intelligence techniques. Tech. rep., U.S. army armament research, development and engineering center (1993)
32. Moseley, LG., Cartwright, M.: The development of an expert system for operational use in the selection of industrial adhesives. Engineering Applications of Artificial Intelligence **5**(4), 319–328 (1992)
33. Musen, M.A.: The Protégé Project: A Look Back and a Look Forward. AI Matters **1**(4), 4–12 (Jun 2015). https://doi.org/10.1145/2757001.2757003
34. Poveda-Villalón, M., Fernández-Izquierdo, A., Fernández-López, M., García-Castro, R.: LOT: An industrial oriented ontology engineering framework. Engineering Applications of Artificial Intelligence **111**, 104755 (May 2022). https://doi.org/10.1016/j.engappai.2022.104755
35. Poveda-Villalón, M., Gómez-Pérez, A., Suárez-Figueroa, M.C.: OOPS! (OntOlogy Pitfall Scanner!): An on-line tool for ontology evaluation. International Journal on Semantic Web and Information Systems (IJSWIS) **10**(2), 7–34 (2014). https://doi.org/10.4018/ijswis.2014040102
36. Psarommatis, F., Kiritsis, D.: A hybrid Decision Support System for automating decision making in the event of defects in the era of Zero Defect Manufacturing. Journal of Industrial Information Integration **26**, 100263 (2022). https://doi.org/https://doi.org/10.1016/j.jii.2021.100263, https://www.sciencedirect.com/science/article/pii/S2452414X21000613
37. Ramos, L.: Semantic Web for manufacturing, trends and open issues: Toward a state of the art. Computers & Industrial Engineering **90**, 444–460 (2015). https://doi.org/https://doi.org/10.1016/j.cie.2015.10.013, https://www.sciencedirect.com/science/article/pii/S0360835215004064
38. RB, Allen., HH, Vanderveldt.: Expert-System Selects Adhesives for Composite-Material Joints. Adhesives Age **38**(7), 16–19 (1995)
39. Scherp, A., Franz, T., Saathoff, C., Staab, S.: F–a model of events based on the foundational ontology dolce+dns ultralight. In: Proceedings of the fifth international conference on Knowledge capture. p. 137–144. K-CAP '09, ACM (Sep 2009). https://doi.org/10.1145/1597735.1597760, http://dx.doi.org/10.1145/1597735.1597760
40. Singh, S., Ghosh, S., Jayaram, J., Tiwari, M.: Enhancing supply chain resilience using ontology-based decision support system. International Journal of Computer Integrated Manufacturing **32**, 1–16 (04 2019). https://doi.org/10.1080/0951192X.2019.1599443
41. Su, YY., Srihari, K., Adriance, J.: A knowledge update mechanism for an adhesive advisor. Computers & industrial engineering **25**(1-4), 111–114 (1993)
42. Uddin, K., Puttonen, J., Scholze, S., Dvoryanchikova, A., Martinez Lastra, J.L.: Ontology-Dased context-Sensitive computing for FMS optimization. Assembly Automation **32**, 163–174 (04 2012). https://doi.org/10.1108/01445151211212316

43. Vandevelde, S., Vennekens, J., Jordens, J., Van Doninck, B., Witters, M.: Knowledge-Based Support for Adhesive Selection: Will it Stick? Theory and Practice of Logic Programming pp. 1–21 (2024). https://doi.org/10.1017/S1471068424000024
44. Yahya, M., Breslin, J.G., Ali, M.I.: Semantic Web and Knowledge Graphs for Industry 4.0. Applied Sciences **11**(11) (2021). https://doi.org/10.3390/app11115110, https://www.mdpi.com/2076-3417/11/11/5110
45. Zheng, Z., Zhou, B., Zhou, D., Zheng, X., Cheng, G., Soylu, A., Kharlamov, E.: Executable knowledge graphs for machine learning: a Bosch case of welding monitoring. In: International Semantic Web Conference. pp. 791–809. Springer (2022)