Quality Without Borders: A Modular Approach to Unified Knowledge Graph Assessment

Gabriele Tuozzo[1,*]

[1]Dipartimento di Informatica, Università degli Studi di Salerno, Fisciano (SA), ITALY

Abstract

Knowledge Graphs (KGs) have emerged as a critical infrastructure for data integration and semantic enrichment across diverse domains, from scientific research to enterprise applications. However, the quality assessment of KGs remains fragmented due to the coexistence of isolated evaluation paradigms, including KG-specific quality frameworks, the FAIR principles, and the 5-star open data scheme. This fragmentation limits metric reusability, and few efforts have been made till now to develop a unified framework for reusing quality measurements, making results comparable, and designing integrated and comprehensive quality assessment tools.

This three-year doctoral research proposes the design of a comprehensive Shared Framework that formally aligns existing quality assessment paradigms for KGs. The framework establishes systematic mappings between KG quality dimensions, FAIR principles, and the 5-star open data scheme, enabling the reuse and extension of existing quality assessment tools to efficiently evaluate FAIRness and openness without computational redundancy.

A preliminary mapping analysis between KG quality dimensions and FAIR principles reveals significant alignment opportunities while exposing critical gaps, particularly in Reusability aspects. Empirical evaluation of existing catalogs, such as the LOD Cloud, confirms widespread Findability issues, including broken links, empty datasets, and missing KGs. To address these challenges, the research proposes a modular, automated KG aggregator employing multiple discovery strategies—crawler-based indexing, search engine APIs, repository harvesting, and large language model guidance—to ensure comprehensive and timely coverage. The research contributes to establishing unified approaches for KG quality assessment and supports broader efforts toward FAIR, open, and high-quality Linked Data. The long-term vision includes developing an interactive "KG Weather Station" dashboard providing real-time, actionable insights on KG quality, FAIR compliance, and openness for both technical and non-technical stakeholders across the Semantic Web ecosystem.

Keywords

data quality, FAIR principles, 5-star open data, quality framework, framework alignment

1. Introduction

In recent decades, the widespread adoption of Semantic Web technologies has led to the creation and dissemination of an unprecedented number of datasets structured as Knowledge Graphs (KGs) [1, 2]. Currently, more than 10,000 datasets are available online, adhering to Linked Data (LD) principles [3]. These KGs exhibit substantial heterogeneity, not only in terms of domain coverage [4], but also in their generation and maintenance processes. Some are developed and curated internally by major technology companies such as Google, Microsoft, Apple, and Amazon, while others are collaboratively maintained by academic institutions and open communities, as exemplified by DBpedia [5] and Wikidata [6]. Their construction methods range from manual curation to semi-automated and fully automated generation pipelines. This diversity results in significant variation in data quality, encompassing highly curated, reliable sources as well as less refined, noisy datasets [7, 1]. Consequently, assessing and monitoring KG quality has become crucial. Quality assessment enables researchers, developers, and practitioners to identify datasets that best suit their specific requirements, supporting informed decisions regarding selection, integration, and reuse. It ensures that KGs are fit for purpose across quality dimensions such as completeness, accuracy, timeliness, and consistency [8]. This is particularly critical in high-impact domains, including Artificial Intelligence (AI) systems, digital assistants, recommender engines, and

ISWC 2025 Companion Volume, November 2–6, 2025, Nara, Japan

*Corresponding author.

gtuozzo@unisa.it (G. Tuozzo)

https://github.com/GabrieleT0 (G. Tuozzo)

0009-0004-5108-1995 (G. Tuozzo)

semantic search, where low-quality data can compromise system performance and trustworthiness. As KGs increasingly serve as foundational infrastructure for data integration, semantic enrichment, and intelligent querying across diverse sectors, ensuring their quality is imperative. High-quality KGs support reliability, reusability, and interoperability in both academic and real-world applications spanning government, industry, and healthcare. To address these needs, several frameworks have been proposed to enable systematic and structured quality evaluation of data [9, 10, 11, 12, 13, 14, 15, 16]. Building upon these foundational models, more specialized frameworks have emerged to address the unique structural and semantic characteristics of KGs. In particular, several KG-specific quality assessment frameworks have been developed [17, 18, 19], offering metrics tailored to the graph-based, interlinked nature of these resources.

Complementing these are more general-purpose frameworks that have gained wide acceptance across domains. Among these, the FAIR principles [20] have become a standard for promoting Findability, Accessibility, Interoperability, and Reusability of digital research outputs. Likewise, the 5-star open data scheme [21] introduced by Tim Berners-Lee provides a progressive model for publishing structured data on the Web, advocating for the use of URIs, RDF, and linking to external datasets to foster interoperability.

Despite the adoption of these three families of frameworks—KG-specific quality assessment models, the FAIR principles, and the 5-star schema—they are currently applied in isolation. There is no established mechanism to align or map metrics across them, making it difficult to reuse existing assessments, compare results across standards, or develop unified evaluation tools. This fragmentation poses a significant barrier to achieving consistent, scalable, and automated quality evaluation of LD.

The central goal of my three-year doctoral research is to design a Shared Framework that integrates these three prominent paradigms into a cohesive and extensible quality assessment system (towards replying to **RQ1**). In particular, the focus will be on integrating high-level, format-independent frameworks—such as FAIR and the 5-Star Open Data scheme—with more operational quality assessment frameworks specifically designed for Linked Data. The mapping process will be thoroughly documented to ensure compatibility with additional frameworks, while the implementation will follow a modular design, making it easy to adapt and extend to future needs. This Shared Framework will allow the reuse of quality assessment of existing tools, avoid redundant computations, and support consistent evaluation workflows across various domains and use cases.

An initial mapping—presented in Section 5—between the KG quality dimensions outlined by Zaveri et al. [18] and the FAIR principles, reveals both alignment and important gaps. Significant gaps emerge in Reusability, highlighting the need to introduce novel quality metrics to enable a comprehensive assessment of the FAIR principles on KGs (towards replying to **RQ2**). The mapping illustrated in Figure 1, has already been employed to assess the FAIRness of some datasets within the LOD Cloud related to the cultural heritage [22]. Preliminary results have shown that Findability represents one of the main challenges affecting these resources. Consequently, to address this issue this research also proposes the development of a modular, automated KG aggregator that employs a variety of discovery strategies—including crawler-based indexing, search engine APIs, repository harvesting, and the use of Large Language Models (LLMs) trained on scientific literature to detect KG projects associated with published research (towards replying to **RQ3**).

The long-term vision is to implement an interactive dashboard, envisioned as a "weather station" for KGs, capable of offering real-time, actionable insights into a dataset's quality, FAIR compliance, and 5-star status. This dashboard will be designed for both technical and non-technical users, supporting broader adoption and accountability in LD publishing.

The rest of this doctoral consortium is structured as follows. Section 2 outlines the significance and the impact expected from this research. Section 3 reviews the related literature. Section 4 introduces the research questions, discussing the associated challenges and the envisioned solutions. Section 5 reports preliminary results. Section 6 describes the planned evaluation methodologies for validating the research. Finally, Section 7 concludes the paper by outlining current limitations and future directions.

2. Importance

The design of a Shared Framework for KG quality assessment is expected to have substantial importance and long-lasting impact in the Semantic Web community and beyond. Its significance stems from its grounding in well-established and widely adopted frameworks, such as the KG-specific quality model by Zaveri et al. [18]. By building on this foundation, the framework enables a formal alignment between KG quality dimensions, the FAIR principles, and the 5-star open data scheme. A preliminary mapping—outlined in Section 5—demonstrates how existing quality dimensions can be directly reused to assess FAIR compliance, avoiding the need for redundant measurements. This seamless reuse of existing metrics ensures consistency while significantly reducing computational overhead.

This benefit extends to widely adopted quality assessment tools, such as KGHeartBeat [23] and LUZZU [24], which implement evaluations grounded in the theoretical framework proposed by Zaveri et al. The outputs of these tools can be repurposed through the proposed mapping to assess FAIRness and openness according to the 5-star model—without the need to recompute metrics or redesign evaluation pipelines. Even if metric computation is not particularly time-consuming, it still incurs unnecessary consumption of resources, including time and storage. Furthermore, the absence of alignment across frameworks leads to incompatible outputs, hindering triangulation of observations and limiting interoperability. Tools tailored to specific frameworks may also omit insights that are relevant under other perspectives. The adoption of a Shared Framework thus enhances efficiency, scalability, and sustainability in KG quality monitoring. Moreover, it facilitates a more nuanced and integrated understanding of quality dimensions, as opposed to isolated assessments using individual frameworks. Consider the evaluation of licensing: traditional quality frameworks typically assess whether a license is provided in a machine-readable or human-readable format; the FAIR principles focus solely on the presence of a license (specifically, principle R1.1); and the 5-star model requires the license to be open. Within the Shared Framework, a single measurement can simultaneously verify (i) the presence of a license (FAIR), (ii) its format (standard quality frameworks), and (iii) its openness (5-star model). This consolidated evaluation yields a more comprehensive understanding of the metric—an outcome impracticable through the isolated application of any individual framework.

The proposed framework holds considerable promise for a wide array of stakeholders. For researchers and developers, it supports standardized, cross-framework evaluations, enhancing comparability and replicability in empirical studies. For data consumers and analysts, it provides a consistent and transparent method to compare KGs, enabling the identification of high-quality, FAIR-compliant, and interoperable datasets tailored to specific use cases. For data providers and maintainers, it offers a structured path to monitor and improve compliance with quality, FAIRness, and openness standards—thus promoting best practices in LD publication.

Grounding the results of the preliminary alignment and issues observed from empirical assessment, a particularly important contribution of this research lies in its effort to address persistent challenges in KG Findability. In fact, current catalogs, such as the LOD Cloud, often contain broken links, outdated datasets, or incomplete coverage, impeding discovery and reuse. To tackle this, the research proposes a modular, automated KG aggregator that combines web crawlers, API access, repository harvesting, and LLM-guided strategies to dynamically detect, index, and update KG collections across domains.

In summary, this work aims to bridge fragmented quality paradigms through a unified, extensible, and operationally effective framework. By enabling interoperable, reusable, and comparable quality assessments, it has the potential to transform how KGs are monitored, published, and consumed—ultimately supporting a more trustworthy, FAIR, and usable Web of Data. To make the framework actionable and accessible, the research will culminate in the development of an interactive dashboard—conceptualized as a "weather station" for KGs. This tool will deliver real-time, visual insights into KG quality, FAIR compliance, and 5-star adherence, supporting both expert users (e.g., data engineers, knowledge scientists) and non-experts (e.g., domain researchers, decision-makers) in evaluating and selecting KGs.

3. Related Work

This section reviews related work, beginning with studies that have explored the alignment of various data quality frameworks. It then examines research focused on assessing the FAIRness of KGs, followed by investigations into their Findability. The section concludes with a critical discussion of the current state of the art, highlighting gaps and positioning this research within the broader scientific context.

Quality Frameworks Alignment. The alignment of data quality frameworks has received limited attention in the literature, with only a few notable contributions addressing this challenge. One such effort is by Miller et al. [25], who present a comprehensive review of both generic and domain-specific quality frameworks and propose a standardized quality model that maps quality dimensions from diverse frameworks onto a shared vocabulary. This approach facilitates the development of interoperable tools and applications capable of evaluating data quality consistently across multiple existing frameworks.

Within the Semantic Web domain, the only significant attempt to align different quality paradigms is the work of Hasnain et al. [26], who explore the relationship between the 5-star open data scheme and the FAIR principles. Their study demonstrates that the FAIR principles can be viewed as an extension of the 5-star model, rather than vice versa. They establish several correspondences between the two, noting that in some cases, multiple FAIR principles map to the same star. For instance, principle F1 ("(Meta)data are assigned globally unique and persistent identifiers") aligns with the fourth star ("data items should have a URI and can be shared on the Web"). However, the mapping remains incomplete—principle A2 ("Metadata should be accessible even when the data is no longer available") is notably left unmapped. This highlights the need for further work to fully integrate and operationalize such frameworks.

FAIR principles and Knowledge Graphs. Since the introduction of the FAIR principles in 2016, numerous studies have sought to operationalize their assessment on KGs, primarily through the development of specialized tools. Table 1 compares relevant works assessing FAIRness in KGs and ontologies, reporting each contribution, reference, and year, covered principles, and support for evaluating KGs and ontologies. The table is sorted by publication year. The symbol "∼" denotes partial coverage of sub-principles or tools designed for generic rather than linked data types.

Among the most prominent tools are O'FAIR [27] and FOOPS![28], both tailored for ontologies. O'FAIR provides full FAIR coverage, while FOOPS! omits some principles (e.g., A1.2, I3). FAIR-Checker[29] leverages Semantic Web technologies to assess FAIRness of ontologies and web resources, with a strong focus on metadata quality, particularly Interoperability and Reusability, but lacks support for F3 and A2. Tools like FUJI [30] target heterogeneous data types, including KGs and ontologies, though they overlook the specificities of linked data.

In contrast, leveraging the Shared Framework together with existing quality assessment tools such as KGHeartBeat and LUZZU—which currently cover the widest range of quality metrics in the literature, though any other tool grounded in the KG quality framework can also be applied—enables comprehensive evaluation of the FAIR principles and supports the assessment of both ontologies and KGs.

Knowledge Graphs Findability. Over the years, numerous initiatives have sought to improve the Findability of datasets structured as KGs by collecting and cataloging them from across the Web. To support this effort, several crawlers have been developed, including LDspider [31], the tool proposed by Reihaneh et al. [32], and the more recent Squirrel [33]. LDspider provides a scalable, open-source solution for crawling RDF documents through URI dereferencing, with strong modularity but no support for SPARQL endpoint discovery. The tool proposed by Reihaneh et al. [32] enhances RDF retrieval by scoring HTML links based on their likelihood of leading to RDF content, improving efficiency but lacking endpoint discovery and adaptability. Squirrel was developed with the goal of overcoming the limitations of LDspider, positioning itself as a more efficient tool thanks to its native support for parallelized crawling and its ability to handle a broader range of serialization formats. It also maintains a modular architecture that enables easy extension to support additional formats. However, despite these advantages, the tool does not support the discovery of SPARQL endpoints. In contrast, SpEdD [34]—introduced by Yumusak et al.—addresses this gap by leveraging search engines to discover SPARQL endpoints and monitor their availability, offering broader endpoint coverage. Its main limitation

Table 1
Comparison of related works assessing the FAIR principles on KGs and ontologies. For each work, we report its contribution, reference with publication year, the covered FAIR principles, and whether the approach supports the assessment of KGs and ontologies. The symbol "∼" indicates partial coverage of sub-principles or that the approach targets generic data types rather than specifically addressing linked data. The table is sorted by publication year.

Contribution	Ref.	Year	F	A	I	R	Ontologies	KGs
Hasnain et al. + LUZZU, KGHeartBeat or similar	[26]	2018	✓	∼	✓	✓	✓	✓
F-UJI	[30]	2021	✓	✓	✓	✓	∼	∼
FOOPS!	[28]	2021	✓	∼	∼	✓	✓	
O'FAIR	[27]	2022	✓	✓	✓	✓	✓	
FAIR-Checker	[29]	2023	∼	∼	✓	✓	✓	✓
Shared Framework + LUZZU, KGHeartBeat or similar		2025	✓	✓	✓	✓	✓	✓

lies in its dependence on keyword extraction and a limited number of public search engine indexing. Additionally, no mechanism to discover RDF dump is implemented. However, a review of the code repositories of these four tools suggests that these tools are no longer actively maintained or supported.

In parallel, several catalogs have emerged to aggregate and index datasets structured as KGs across various domains. Among the most widely recognized are the LOD Cloud, DataHub, and LODAtlas. These serve as general-purpose catalogs, while others focus on specific domains—for example, AgroPortal for agronomic data, the Linguistic Linked Open Data (LLOD) cloud for linguistic resources, and BioPortal for biomedical ontologies. Additionally, several initiatives concentrate on national or regional-level aggregation. Notable examples include Europeana (Europe), DigitalNZ (New Zealand), the Digital Public Library of America, and Trove (Australia).

Overall Discussion and Positioning of This Research. To the best of our knowledge, no existing work in the Semantic Web domain has attempted to integrate KG quality assessment frameworks with both the FAIR principles and the 5-star open data scheme. While Hasnain et al. [26] provide a valuable mapping between the 5-star scheme and the FAIR principles, their work does not incorporate any established KG quality frameworks. Moreover, their mapping remains partial, with key principles, e.g., A2 ("metadata should remain accessible even if the data are no longer available"), left unmapped.

The operationalization of FAIR for KGs remains an open challenge. Existing tools—such as O'FAIR, FOOPS!, and FAIR-Checker—focus primarily on ontologies or web resources and offer limited coverage of the full FAIR principles, with none addressing KG-level evaluation comprehensively.

Similarly, KG discoverability tools and catalogs remain fragmented and often outdated. Crawlers like LDspider [31], SpEdD [34], and Squirrel [33] are no longer actively maintained, limiting their effectiveness for continuous dataset discovery. Furthermore, no existing tool is capable of combining the discovery of RDF dumps together with the SPARQL endpoints.

Likewise, major cross-domain catalogs such as LOD Cloud, DataHub, and LODAtlas each exhibit shortcomings: LOD Cloud struggles with completeness and link decay, DataHub has ceased indexing new entries, and LODAtlas has not been updated since 2018.

This work addresses two key challenges in the KG ecosystem. First, it proposes a unified quality assessment framework by aligning the established model by Zaveri et al. [18] with the FAIR principles and the 5-star open data scheme. Second, to enhance KG discoverability, the study introduces a modular and extensible aggregator.

4. Research Questions

This section formulates the Research Questions (RQs), discusses the challenges associated with each of them, and outlines the strategies to address the RQs.

RQ1: To what extent could a quality framework tailored to KGs be formally aligned with broader paradigms such as the FAIR principles and the 5-star open data scheme?
A strategic choice of this research is to build upon existing quality frameworks, such as the widely adopted model proposed by Zaveri et al. [18], rather than reinventing the wheel.

Hypothesis: by aligning an established framework tailored to KGs with broader paradigms like the FAIR principles and the 5-star scheme, it enables the reuse of existing quality assessment tools to enable the integrated assessment of FAIRness and openness. It ensures methodological consistency while reducing duplication of effort.

The primary challenge lies in determining which quality dimensions are most appropriately linked to the various FAIR principles and 5-star criteria. This task is complicated by the fact that a single quality dimension may be relevant to multiple FAIR principles or stars.

To address this research question, a systematic analysis of theoretical frameworks and evaluation tools is planned in order to obtain a holistic view of the dimensions and metrics, and to propose an alignment assessed in terms of acceptance and correctness by domain experts, as well as clarity and usability by end users. A preliminary experiment was conducted using the Zaveri et al. framework, the alignment presented in Section 5, and was tested within the cultural heritage domain [22]. By operating at the dimension level, the alignment supports portability and adaptability with other quality frameworks, because the alignment starts at the dimensions level, which are generally more conceptual and transferable across different quality contexts, while metrics are often highly specific and implementation-dependent.

The proposed mapping will be thoroughly detailed, documenting the rationale behind the selection of each dimension and its corresponding metrics, as well as the methodology used for alignment. This approach aims to ensure the methodology is fully reproducible, enabling the alignment of additional frameworks by reusing the same procedures. Moreover, the implementation will be modular, allowing for easy adaptation and extension to accommodate future requirements.

RQ2: What new metrics are needed to achieve full compliance of individual frameworks with the entire Shared Framework?
Hypothesis: From the preliminary alignment proposed in Section 5, it emerged that the frameworks are not perfectly aligned. For example, principles such as R1.3 ("Data and metadata organized in standardized ways") and A2 ("Dataset registered in search engines") are not adequately covered by existing standard quality frameworks. Therefore, to support a comprehensive evaluation of FAIRness and 5-star open data, it is necessary to address these gaps by introducing new quality metrics.

For example, to assess whether metadata are defined using well-known vocabularies such as DCAT or VoID. Similarly, under the 5-star open data scheme, the requirement for datasets to be associated with an open license is explicit. However, most current quality frameworks only assess whether a license is present, without verifying whether the license itself meets open data standards. Overall, the development of these new metrics is essential to extend existing frameworks and enable a more complete and standards-aligned assessment of KG quality.

To address this research question, an in-depth analysis of the frameworks intended for alignment is required in order to identify and understand how to practically measure the targeted aspect, thus enabling its operationalization through the theorization of a novel quality metric.

RQ3: How can the Findability of Knowledge Graphs be improved?
Preliminary empirical results from applying the Shared Framework in the cultural heritage domain [22] highlight that Findability remains one of the most persistent and critical challenges within the KG ecosystem. In this context, the discoverability of KGs can be considered a necessary first step for any meaningful quality assessment. Locating KGs—especially in the absence of prior knowledge regarding their exact names or hosting platforms—is often a complex and non-trivial task.

Although centralized and manually curated repositories such as the LOD Cloud aim to index and aggregate such datasets, they are not consistently maintained. As noted by Debattista et al. [35] and further confirmed by our preliminary results described in Section 5, many of the datasets listed in the LOD Cloud suffer from serious accessibility issues. These include broken links, inactive endpoints,

datasets with zero triples, unstructured content, or static data dumps lacking live querying functionality.

Hypothesis: the Findability of KGs can be significantly improved through the design of an automated, modular, and continuously updated aggregation infrastructure that minimizes reliance on manual submissions and fragmented discovery tools.

To address this research question, the proposed research will investigate multiple complementary discovery strategies: crawler-based aggregation over widely used repositories such as Zenodo and GitHub; the use of Google Search APIs to identify new or relocated SPARQL endpoints; integration with existing catalogs through public APIs (e.g., KGHeartBeat [23]); and the exploitation of LLMs trained on scientific literature to detect KG projects associated with published research, even when such projects are not indexed in traditional catalogs.

The ultimate goal is to develop a unified, continually updated catalog of KGs, offering a consistent and user-friendly interface. This catalog will serve as a centralized entry point for KG discovery, eliminating the need to navigate multiple, fragmented, and heterogeneous repositories, thereby significantly enhancing the overall Findability of KGs across domains.

5. Preliminary Results

Initial progress toward addressing **RQ1** has focused on the alignment of KG quality frameworks with the FAIR principles. A preliminary mapping [22], illustrates how high-level quality dimensions from the widely adopted framework by Zaveri et al. [18] can be systematically aligned with individual FAIR principles. The proposed alignment is visualized using a Sankey diagram (Figure 1), highlighting the conceptual flow between KG quality dimensions and FAIR criteria. This alignment maintains an abstraction at the level of quality dimensions, rather than metrics. This strategic design choice ensures compatibility with other quality models and supports extensibility, thus making the Shared Framework usable even with frameworks other than the one defined by Zaveri et al.

Observing the Figure 1, it is possible to see how the quality dimension Availability is shown to contribute to multiple FAIR principles, including both Findability and Accessibility, demonstrating how single quality aspects may support multiple FAIR objectives. From this high-level abstraction, the alignment then extends toward the metric level, enabling the operationalization of FAIR assessments on concrete KGs. Current mappings reveal that certain FAIR principles—particularly R1.3 ("Data and metadata organized in standardized ways") and A2 ("dataset registered in search engines")—remain uncovered, exposing gaps in existing quality frameworks. These findings directly relate to **RQ2**, highlighting the need for new quality metrics that can capture metadata-specific aspects not addressed by current approaches.

In parallel, progress toward **RQ3** has been initiated through the development of CHeCLOUD [22], the Cultural Heritage Cloud, conceived as a potential LOD sub-cloud. The same work provides evidence of several structural issues affecting the current LOD Cloud. Notably, many listed KGs link to non-functional endpoints, have been relocated without updates, contain zero triples, or are not structured as valid KGs. Moreover, the catalog favors static data dumps over live access methods, such as SPARQL endpoints, limiting both reusability and real-time querying capabilities [35, 36]. Critically, several noteworthy datasets—such as ATLAS and ArCo—are entirely absent from the LOD Cloud, further demonstrating its lack of completeness and the need for a more reliable, dynamic, and comprehensive aggregation mechanism for KGs. While CHeCLOUD currently targets a specific domain, it serves as a proof of concept for the broader objective of this research: to develop a generalized, automated aggregator that supports Findability for KGs across all domains. The goal is to create a unified, modular platform that addresses the fragmented landscape of KG catalogs and supports more effective discovery, monitoring, and reuse of semantic data resources.

Together, these preliminary efforts represent foundational steps toward building a Shared Framework for unified KG quality assessment and discoverability, addressing key challenges across the Semantic Web ecosystem.

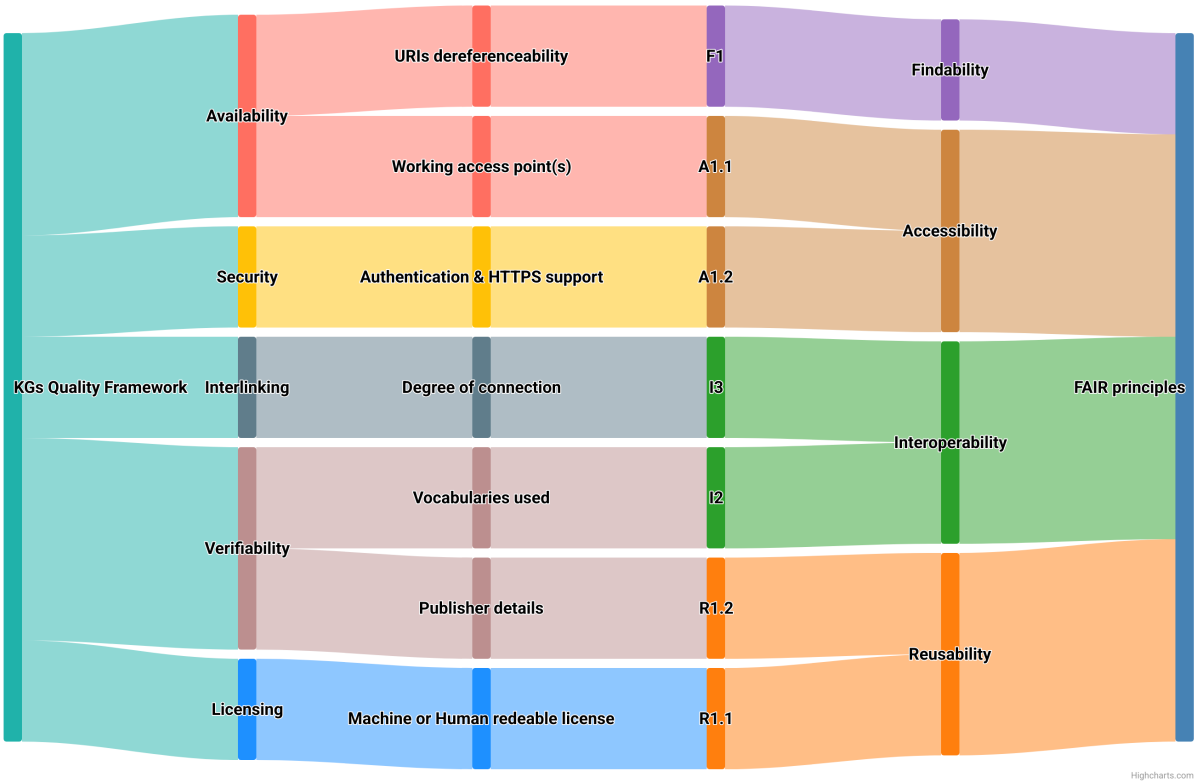

Figure 1: Mapping between KGs quality framework and FAIR principles proposed in [22]. On the left side of the diagram, the KG quality framework employed is the one proposed by Zaveri et al. [18]

6. Planned Evaluation

To ensure the theoretical robustness and practical relevance of the Shared Framework, a multi-phase evaluation strategy will be employed—combining expert feedback, tool benchmarking, and user-centered testing. This approach aims to validate both conceptual alignment with expert expectations and real-world applicability for diverse Semantic Web stakeholders.

- **Expert Validation via Delphi Panel**: To ensure conceptual robustness, a Delphi panel study will be conducted involving a selected group of domain experts with recognized experience and scholarly contributions in the areas of KG quality, data governance, and FAIR principles. Through a structured, iterative process, panelists will be invited to review, comment on, and refine the design and components of the proposed framework. The goal of this study is to achieve consensus on the framework's design logic, identify overlooked dimensions, and validate the alignment between quality dimensions, FAIR principles, and the 5-star open data model. Therefore, this method will be employed to validate the effectiveness in addressing **RQ1** and **RQ2**.
- **Comparative Operational Assessment**: The operational effectiveness of the Shared Framework will be evaluated through a comparative study involving existing quality assessment tools, including KGHeartBeat [23], LUZZU [24] using archived results available in literature due to its deprecation, O'FAIR [27], FOOPS! [28], and FAIR-Checker [29]. The evaluation will be performed across a diverse collection of KGs, spanning multiple domains (e.g., cultural heritage, life sciences, government data), and will assess metrics such as coverage, accuracy, consistency, and precision of quality reporting across different frameworks. This analysis will help determine the added value of the proposed framework in bridging multiple assessment paradigms and validate, from an empirical point of view, the effectiveness in addressing **RQ1** and **RQ2**.
- **Empirical evaluation of the catalog completeness**: To assess whether the objective defined for **RQ3** has been achieved, an analysis will be conducted on the number of KGs and ontologies indexed in the proposed catalog. This evaluation will verify the extent to which the catalog is

as comprehensive as possible—that is, whether it successfully aggregates all resources already indexed in well-known catalogs such as DataHub and the LOD Cloud. Additionally, the assessment will include measuring the number of new datasets that have been indexed and are not listed in any existing catalog.

- **Usability Testing of the "Weather Station" Dashboard**: To assess the utility and accessibility of the framework's human-facing components, a third evaluation stream will focus on the interactive dashboard—conceptualized as a "weather station" for KGs. Usability testing sessions will be conducted with both technical users (e.g., data engineers, researchers) and non-technical users (e.g., librarians, policymakers, and domain experts), measuring usability dimensions such as learnability, efficiency, user satisfaction, and clarity of visualizations.

7. Limitations and Future Directions

The proposed Shared Framework introduces a promising approach to unifying KG quality assessment. This work contributes to the broader vision of a more interoperable, FAIR-compliant, and trustworthy Web of Data by offering a foundation for harmonized assessment strategies and actionable insights for diverse stakeholders. However, at this current stage, the preliminary mapping between the Zaveri et al. quality dimensions and the FAIR principles is not yet complete. As illustrated in Figure 1, certain principles, most notably R1.3 and A2, do not currently align with any existing quality dimensions. Furthermore, the framework has not yet been validated across KGs from diverse domains. Future research will address these limitations. First, a systematic validation of the Shared Framework will be carried out across KGs spanning various domains to assess its robustness and adaptability. Concurrently, the framework will be extended to cover all remaining FAIR principles by introducing new quality metrics where necessary. Following this, the construction of a KGs aggregator will facilitate automated KG discovery to solve the Findability. Finally, the research will culminate in the development of an interactive "KG Weather Station" dashboard, designed to deliver real-time insights into KGs quality for a broad audience, including both technical and non-technical users.

As future limitations expected at the conclusion of this research project, the scalability of the envisioned KG Weather Station. When extended to assess thousands of KGs, it will encounter infrastructure challenges related to storage, computation, bandwidth, and long-term maintenance. Ensuring sustainability will require integration with scalable architectures (e.g., cloud-native services) and the establishment of external collaborations. Currently, the framework operates largely independently of broader data governance initiatives, such as those led by W3C and ISO. Aligning with these efforts and engaging with their communities will be essential to foster adoption, enhance interoperability, and ensure long-term impact.

Acknowledgments

This work is funded by the European Alliance NEOLAiA (Project 101124794: "NEOLAiA – Transforming Regions for an Inclusive Europe"). I would like to express my sincere gratitude to my supervisors, Prof. Vittorio Scarano and Prof. Maria Angela Pellegrino, for their invaluable guidance and support.

Declaration on Generative AI

During the preparation of this work, the author used chatGPT in order to: Grammar and spelling check.

References

[1] D. Kontokostas, P. Westphal, S. Auer, S. Hellmann, J. Lehmann, R. Cornelissen, Databugger: a test-driven framework for debugging the web of data, in: 23rd International World Wide Web Conference, WWW, ACM, 2014, pp. 115–118. doi:10.1145/2567948.2577017.

[2] A. Langer, V. Siegert, C. Göpfert, M. Gaedke, SemQuire - Assessing the Data Quality of Linked Open Data Sources Based on DQV, in: Current Trends in Web Engineering - ICWE International Workshops, MATWEP, EnWot, KD-WEB, WEOD, TourismKG, volume 11153, Springer, 2018, pp. 163–175. doi:10.1007/978-3-030-03056-8_14.

[3] G. Sejdiu, A. Rula, J. Lehmann, H. Jabeen, A Scalable Framework for Quality Assessment of RDF Datasets, in: The Semantic Web – ISWC, Springer International Publishing, 2019, pp. 261–276. doi:10.1007/978-3-030-30796-7_17.

[4] M. Schmachtenberg, C. Bizer, H. Paulheim, Adoption of the linked data best practices in different topical domains, in: ISWC, Springer, 2014, pp. 245–260. doi:10.1007/978-3-319-11964-9_16.

[5] J. Lehmann, R. Isele, M. Jakob, A. Jentzsch, D. Kontokostas, P. N. Mendes, S. Hellmann, M. Morsey, P. Van Kleef, S. Auer, et al., DBpedia–a large-scale, multilingual knowledge base extracted from Wikipedia, Semantic web 6 (2015) 167–195. doi:10.3233/SW-140134.

[6] D. Vrandečić, M. Krötzsch, Wikidata: a free collaborative knowledge base, Communications 57 (2014) 78–85. doi:10.1145/2629489.

[7] T. Käfer, A. Abdelrahman, J. Umbrich, P. O'Byrne, Observing Linked Data Dynamics, in: The Semantic Web: Semantics and Big Data, Springer Berlin Heidelberg, 2013, pp. 213–227. doi:10.1007/978-3-642-38288-8_15.

[8] J. Juran, Juran's Quality Control Handbook, Mcgraw-Hill, 1988.

[9] C. Cichy, S. Rass, An overview of data quality frameworks, IEEE Access 7 (2019) 24634–24648. doi:10.1109/ACCESS.2019.2899751.

[10] Y. W. Lee, D. M. Strong, B. K. Kahn, R. Y. Wang, Aimq: a methodology for information quality assessment, Information Management 40 (2002) 133–146. doi:10.1016/S0378-7206(02)00043-5.

[11] C. Batini, C. Cappiello, C. Francalanci, A. Maurino, Methodologies for data quality assessment and improvement, ACM Comput. Surv. 41 (2009). doi:10.1145/1541880.1541883.

[12] L. L. Pipino, Y. W. Lee, R. Y. Wang, Data quality assessment, Commun. ACM 45 (2002) 211–218. doi:10.1145/505248.506010.

[13] M. del Pilar Angeles, F. García-Ugalde, A data quality practical approach, International Journal on Advances in Software 2 (2009). URL: https://personales.upv.es/thinkmind/dl/journals/soft/soft_v2_n23_2009/soft_v2_n23_2009_7.pdf.

[14] B. Carlo, B. Daniele, C. Federico, G. Simone, A data quality methodology for heterogeneous data, International Journal of Database Management Systems 3 (2011) 60–79. doi:10.5121/IJDMS.2011.3105.

[15] C. Cappiello, P. Ficiaro, B. Pernici, Hiqm: A methodology for information quality monitoring, measurement, and improvement, in: J. F. Roddick, V. R. Benjamins, S. Si-said Cherfi, R. Chiang, C. Claramunt, R. A. Elmasri, F. Grandi, H. Han, M. Hepp, M. D. Lytras, V. B. Mišić, G. Poels, I.-Y. Song, J. Trujillo, C. Vangenot (Eds.), Advances in Conceptual Modeling - Theory and Practice, Springer Berlin Heidelberg, Berlin, Heidelberg, 2006, pp. 339–351. doi:10.1007/11908883_41.

[16] R. Y. Wang, A product perspective on total data quality management, Commun. ACM 41 (1998) 58–65. doi:10.1145/269012.269022.

[17] B. Xue, L. Zou, Knowledge graph quality management: A comprehensive survey, IEEE Transactions on Knowledge and Data Engineering 35 (2023) 4969–4988. doi:10.1109/TKDE.2022.3150080.

[18] A. Zaveri, A. Rula, A. Maurino, R. Pietrobon, J. Lehmann, S. Auer, Quality assessment for linked data:a survey: A systematic literature review and conceptual framework, Semantic Web 7 (2015) 63–93. doi:10.3233/SW-150175.

[19] H. Chen, G. Cao, J. Chen, J. Ding, A practical framework for evaluating the quality of knowledge graph, in: X. Zhu, B. Qin, X. Zhu, M. Liu, L. Qian (Eds.), Knowledge Graph and Semantic Computing:

Knowledge Computing and Language Understanding, Springer Singapore, Singapore, 2019, pp. 111–122. doi:10.1007/978-981-15-1956-7_10.

[20] M. D. Wilkinson, M. Dumontier, I. J. Aalbersberg, G. Appleton, M. Axton, A. Baak, N. Blomberg, J.-W. Boiten, L. B. da Silva Santos, P. E. Bourne, et al., The fair guiding principles for scientific data management and stewardship, Scientific data 3 (2016) 1–9. doi:10.1038/sdata.2016.18.

[21] 5-star Open Data, Conformant licenses, 2018. URL: https://5stardata.info/en/, [Online. Last access 2025, May].

[22] A. Lieto, M. A. Pellegrino, G. Tuozzo, The fairness of checloud, the cultural heritege linked open data cloud, Semantic web (""). URL: https://www.semantic-web-journal.net/content/fairness-checloud-cultural-heritage-linked-open-data-cloud.

[23] M. A. Pellegrino, A. Rula, G. Tuozzo, Kgheartbeat: An open source tool for periodically evaluating the quality of knowledge graphs, in: International Semantic Web Conference, Springer, 2024, pp. 40–58. doi:10.1007/978-3-031-77847-6_3.

[24] J. Debattista, S. Auer, C. Lange, Luzzu-A Methodology and Framework for Linked Data Quality Assessment, Journal of Data and Information Quality 8 (2016). doi:10.1145/2992786.

[25] R. Miller, S. H. M. Chan, H. Whelan, J. Gregório, A comparison of data quality frameworks: A review, Big Data and Cognitive Computing 9 (2025). doi:10.3390/bdcc9040093.

[26] A. Hasnain, D. Rebholz-Schuhmann, Assessing fair data principles against the 5-star open data principles, in: The Semantic Web: ESWC 2018 Satellite Events: ESWC 2018 Satellite Events, Heraklion, Crete, Greece, June 3-7, 2018, Revised Selected Papers, Springer-Verlag, Berlin, Heidelberg, 2018, p. 469–477. doi:10.1007/978-3-319-98192-5_60.

[27] E. Amdouni, S. Bouazzouni, C. Jonquet, O'faire: Ontology fairness evaluator in the agroportal semantic resource repository, in: P. Groth, A. Rula, J. Schneider, I. Tiddi, E. Simperl, P. Alexopoulos, R. Hoekstra, M. Alam, A. Dimou, M. Tamper (Eds.), The Semantic Web: ESWC 2022 Satellite Events, Springer International Publishing, 2022, pp. 89–94. doi:10.1007/978-3-031-11609-4_17.

[28] D. Garijo, O. Corcho, M. Poveda-Villalón, Foops!: An ontology pitfall scanner for the fair principles., in: ISWC (Posters/Demos/Industry), 2021. URL: https://ceur-ws.org/Vol-2980/paper321.pdf.

[29] A. Gaignard, T. Rosnet, F. De Lamotte, V. Lefort, M.-D. Devignes, Fair-checker: supporting digital resource findability and reuse with knowledge graphs and semantic web standards, Journal of Biomedical Semantics 14 (2023) 7. doi:10.1186/s13326-023-00289-5.

[30] A. Devaraju, R. Huber, An automated solution for measuring the progress toward fair research data, Patterns 2 (2021) 100370. URL: https://www.sciencedirect.com/science/article/pii/S2666389921002324. doi:https://doi.org/10.1016/j.patter.2021.100370.

[31] R. Isele, J. Umbrich, C. Bizer, A. Harth, Ldspider: An open-source crawling framework for the web of linked data, in: International Workshop on the Semantic Web, 2010. URL: https://api.semanticscholar.org/CorpusID:6290825.

[32] R. Emamdadi, M. Kahani, F. Zarrinkalam, A focused linked data crawler based on html link analysis, in: 2014 4th International Conference on Computer and Knowledge Engineering (ICCKE), 2014, pp. 74–79. doi:10.1109/ICCKE.2014.6993406.

[33] M. Röder, G. de Souza Jr, A.-C. N. Ngomo, Squirrel – crawling rdf knowledge graphs on the web, in: The Semantic Web – ISWC 2020, Springer International Publishing, Cham, 2020, pp. 34–47. doi:10.1007/978-3-030-62466-8_3.

[34] H. K. A. K. P.-Y. V. Semih YUMUSAK, Erdogan DOGDU, Spend: Linked data sparql endpoints discovery using search engines, IEICE TRANSACTIONS on Information E100-D (2017) 758–767. doi:10.1587/transinf.2016DAP0025.

[35] J. Debattista, J. Attard, R. Brennan, D. O'Sullivan, Is the lod cloud at risk of becoming a museum for datasets? looking ahead towards a fully collaborative and sustainable lod cloud, in: The World Wide Web Conference, 2019, pp. 850–858. doi:10.1145/3308560.33170.

[36] T. Gabriele, Navigating the LOD Subclouds: Assessing Linked Open Data Quality by Domain, in: Companion Proceedings of the Web Conference, ACM, 2025. doi:10.1145/3701716.3717569.