

# A Dual Knowledge Graph-Driven Computation of SDG Indicators

Wissal Benjira  
De Vinci Higher Education, De Vinci  
Research Center & LASTIG, IGN  
Paris, France  
wissal.benjira@devinci.fr

Nicolas Travers  
De Vinci Higher Education  
De Vinci Research Center  
Paris, France  
nicolas.travers@devinci.fr

Faten Atigui  
CEDRIC, CNAM  
Paris, France  
faten.atigui@lecnam.net

Bénédicte Bucher  
LASTIG, IGN  
Paris, France  
benedicte.bucher@ign.fr

Malika Grim-Yefsah  
LASTIG, IGN  
Paris, France  
malika.grim-yefsah@ign.fr

## ABSTRACT

Open data is a key resource for extracting indicators that support decision making in urban planning, environmental monitoring, and sustainable development. While indicators follow standardized definitions, open data remains fragmented, schema less, and inconsistent, making integration and computation challenging. This paper presents SDG-KG, a Knowledge Graph (KG) framework for the integration, mapping, and validation of indicators using multiple datasets. We illustrate our approach with the Sustainable Development Goals (SDGs), for which our framework structures textual definitions, guides data retrieval, and formalizes indicator computation. Using graph based modeling, we build both an SDG Graph and a Metadata Graph, resolve conflicts between sources, and generate query driven execution workflows. Experiments on six countries and two SDG indicators demonstrate the effectiveness, and interpretability of the proposed framework.

## CCS CONCEPTS

• **Information systems** → **Graph-based database models**; • **Computing methodologies** → **Semantic networks**; • **Software and its engineering** → *Automatic programming*.

## KEYWORDS

Knowledge Graph, Open Data, Data integration, SDG

### ACM Reference Format:

, Wissal Benjira, Nicolas Travers, Faten Atigui, Bénédicte Bucher, and Malika Grim-Yefsah. 2018. A Dual Knowledge Graph-Driven Computation of SDG Indicators. In *Proceedings of (IEEE Big Data'25)*. ACM, New York, NY, USA, 10 pages. <https://doi.org/XXXXXXXX.XXXXXXX>

## 1 INTRODUCTION

Tracking progress is important for understanding how societies evolve. It helps measure economic growth, social development, and

environmental impact. Indicators play a central role in this process. They provide concrete evidence to evaluate policies, identify gaps, and guide decision making. In urban planning, they reveal population growth and infrastructure needs. In environmental monitoring, they track pollution and resource consumption. For sustainable development, they measure progress towards global goals. Reliable and accurate indicators are essential for effective action and sustainable planning [17, 21].

Open data is a powerful resource for building these indicators. It is accessible, transparent, and available across many domains. Governments, organizations, and international agencies publish datasets on health, education, environment, and economic development [9, 26]. These datasets provide insights at global, national, and local levels. They enable policymakers to monitor progress and make informed decisions [3].

However, open data is often fragmented and inconsistent. It is stored in different formats like CSV, GeoJSON, SQL dumps, and APIs. It varies in granularity, spatial coverage, and update frequency. These differences make it hard to combine and analyze data effectively [1, 20]. Traditional methods struggle to harmonize these diverse datasets. There is a need to integrate and compute indicators across multiple sources.

In the context of sustainable development, the United Nations defined 232 indicators to monitor progress across 17 *Sustainable Development Goals* [13] (SDG). Many indicators, especially *Tier II* and *Tier III*, lack data or standard methods. Several national and international bodies encourage the use of open data to improve their coverage and computability. This requires finding relevant datasets, aligning schemas, and formalizing indicator logic.

A promising solution to this challenge is the use of Knowledge Graphs (KGs), which provides a flexible and semantically rich data model for integrating diverse sources while preserving relationships between entities [4, 8]. KGs have been widely adopted for data integration and semantic interoperability across disparate sources [14]. Additionally, Large Language Models (LLMs) have demonstrated strong capabilities in schema matching and entity linking, making them valuable for aligning heterogeneous datasets [16].

To understand how it works, consider the query shown in Figure 1. It focuses on SDG Indicator 11.2.1, which measures the share of population with access to low capacity public transport. The needed data exists in open platforms, but with different schemas and variable names. For example, *geometry* and *ArRGeopoint* describe the same thing but are not linked. Values like *bus* must be

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

*IEEE Big Data'25, December 8–11, 2025, Macau, China*

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-XXXX-X/2018/06

<https://doi.org/XXXXXXXX.XXXXXXX>

mapped to SDG terms such as low capacity. To handle this, we model the indicator in a knowledge graph. It defines concepts, attributes, and allowed values from official SDG metadata. Using LLMs, we align attributes and map values, even when formats or labels differ. Once the graph is populated, we can run queries that retrieve only the relevant data. This enables clean and accurate computation, even from fragmented and inconsistent sources.

In this paper, we present SDG-KG, a framework that supports the structured computation of SDG indicators from open and heterogeneous data sources. Our approach leverages a *dual knowledge graph model* that integrates both indicator semantics and dataset metadata, enabling automated query driven selection, alignment, and execution.

We summarize our contributions as follows:

- We introduce a dual-graph based framework that models indicators, data sources, and metadata to enable transparent and reproducible indicator computation from open data by generating source code indicators.
- We address three key challenges: schema mismatches across sources, conflicting or redundant datasets, and the absence of formal computation rules. We show how our dual graph architecture, combined with LLM-assisted mappings and a trust based filtering mechanism, addresses each.
- We evaluate SDG-KG on real world queries for SDG indicators (e.g., 11.2.1 and 11.6.2), using datasets from six countries. Results show accurate value computation, interpretable source selection, and robustness across metadata profiles.

The remainder of this paper is structured as follows: Section 2 reviews related work on data integration, indicator computation, and the use of knowledge graphs and language models. Section 3 introduces the key components of our modeling approach, including the SDG Graph and Metadata Graph. Section 2.2 presents the SDG-KG framework, detailing its architecture, data flow, and trust based source selection strategy. Section 5 reports the experimental evaluation on indicator computation, mapping quality, and metadata filtering. Section 6 concludes the paper.

## 2 RELATED WORK

### 2.1 Knowledge Graphs for Data Integration

Knowledge Graphs (KGs) are widely used for integrating heterogeneous data sources. They represent data as a set of entities and relationships, enabling semantic alignment across domains. KGs provide a flexible schema and support reasoning, making them useful for unifying structured and semi-structured data [8]. However, most KGs focus on entity level integration and do not capture the procedural or computational logic required for tasks like indicator computation.

Traditional integration methods struggle [5, 18] with schema heterogeneity and context loss. KGs address this by incorporating ontologies and shared vocabularies, which help map concepts across datasets [22]. This makes KGs effective for open data integration, especially when sources vary in format and semantics.

Virtual Knowledge Graphs (VKGs) extend this idea by providing an abstraction layer over existing databases without requiring data duplication [23]. Queries are rewritten dynamically into the

sources, enabling access to distributed data through a unified semantic interface [2]. VKGs are increasingly used in industrial settings where data is too large or sensitive to move. Yet, they assume high metadata quality and lack mechanisms to rank or prioritize among multiple sources

Recent works also explore the combination of KGs with Large Language Models (LLMs) to improve schema matching and data discovery across sources. These hybrid systems leverage the structure of KGs and the contextual understanding of LLMs, offering promising results for open data integration [15]. All these recent techniques have not yet been applied to the studied domain because of the lack of structure in both open data and indicator documentation. It implies a high complexity for the computation of indicators or even the query itself that requires to be tuned manually.

### 2.2 Schema Matching and Computation

Knowledge Graphs are well suited for computing indicators, especially in the context of the SDGs. They can model indicators, data sources. Santos et al. [19] built a city level KG to structure urban mobility indicators and automate dashboard generation. Their work shows how KGs support dynamic computation from diverse inputs. An other approach [ANONYMIZED] combined KGs with Large Language Models (LLMs) to improve schema matching and entity linking. This approach enhances alignment between open datasets and SDG definitions, supporting scalable and accurate indicator computation. Together, these works show that KGs, especially when combined with LLMs, are a promising solution for modeling indicators and fragmented data.

Matching schemas across datasets is a core challenge in data integration. Traditional methods rely on rules or string similarity [18]. Recent work uses embeddings and language models to better capture semantic alignment [25]. In KGs, entity linking is used to match attributes and values to graph nodes. These techniques are critical when working with open data, where metadata is often missing or inconsistent [8].

While KGs and LLMs are effective for schema alignment, most approaches stop at the mapping stage and don't cover the full pipeline from source selection to indicator computation. They often ignore metadata like time or space, and rarely assess how mapping quality affects final results. There's still a gap in building end-to-end, explainable workflows.

## 3 PRELIMINARIES

Our work focuses on computing indicators using a KG based representation. In this section, we present the core modeling components that support this process. We define the structure of the indicator model, the representation of data sources, and the mapping.

### 3.1 SDG Graph

The Sustainable Development Goals (SDGs) are organized into a hierarchy of goals, targets, and indicators, each designed to measure global progress. While prior models like LinkedSDGs and SustainGraph [7] capture SDG concepts and target evolution, SDG Graph focuses on structuring indicators for computation using multiple and heterogeneous data sources from open data.

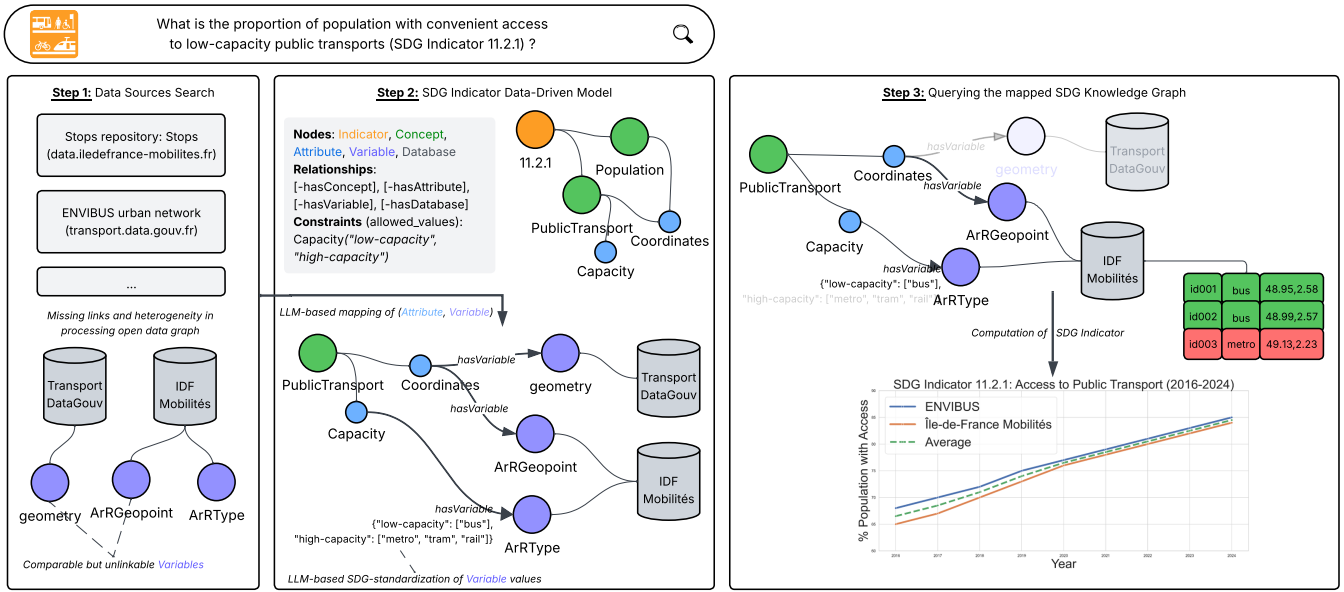


Figure 1: Computing SDG Indicator 11.2.1 from open data. Step 1: collect datasets. Step 2: model the indicator in a knowledge graph. Step 3: align data and compute the result.

The SDG Graph represents the hierarchical structure of the SDGs, where entities and relationships are explicitly defined following the United Nations (UN) *SDG Indicator Metadata Repository*. This structure is automatically generated as part of a data-driven approach and results from a previous expert-validated contribution.

Each node in the graph represents a well defined concept from the SDG documentation. Relationships describe how concepts are connected, from goals to indicators, and from indicators to variables needed for computation. The graph helps formalize indicator logic and connects it to data through the Metadata Graph.

The model consists of the following components:

- **Goals:** High level objectives that address global issues (e.g., *Sustainable Cities and Communities - SDG 11*).

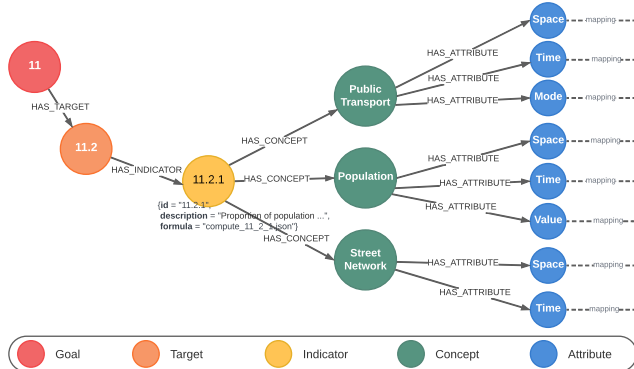


Figure 2: Sample of SDG Graph

- **Targets:** Specific aims within each goal that define measurable outcomes (e.g., *Provide access to safe, affordable, accessible, and sustainable transport systems, Target 11.2*).
- **Indicators:** Quantifiable metrics for evaluating progress toward each target (e.g., *Proportion of the population that has convenient access to public transport, by sex, age, and persons with disabilities, Indicator 11.2.1*).
- **Concepts:** Domain specific entities related to SDG measurement (e.g., *Population, Public Transport, Street Network*).
- **Attributes:** Descriptive properties of concepts used in calculations (e.g., for *Population: Space, Time, Age Group, Sex*).

Concepts and attributes form the basis for computing an indicator. Each concept expects specific input data, which must be linked from real datasets. These inputs include spatial and temporal scopes, disaggregation dimensions, and value types. It worth noting that each concept node is linked to 1 *Space* node, 1 *Time* node and at least 1 *Value* node. This structure is a key structure for SDG computation which requires the alignment of data on space and time all along the process (Section 4.4).

Figure 2 illustrates a sample *SDG Graph*, highlighting the nodes and relationships between goals, targets, indicators, concepts, and attributes. To enable automatic computation, open dataset attributes are mapped to predefined SDG attributes. These mappings are used later in query generation and indicator calculation.

### 3.2 SDG Indicator Formula

To compute an indicator, the corresponding formula is embedded in the SDG Graph's indicator node, represented in a structured JSON format that defines how involved concept nodes are transformed and aggregated. The *formula* is extracted from the UN Metadata




Figure 3: The SDG-KG’s framework of our proposed method.

repository, where it is expressed in natural language or mathematical notation and must be encoded.

Take the following JSON as an example (refers to the *formula* “compute\_11\_2\_1.json” of the indicator node in Figure 2):

```
[
  {"FilteredPop":{"$Filter":["Population","$Qc"]}},
  {"AccessPop":
    "FilteredPop",
    {"$ComputeDistance": ["PublicTransport", "StreetNetwork"]}
  }],
  {"RETURN": {"$div": ["AccessPop", "FilteredPop"]}}
]
```

Here, the JSON document is an ordered list of computation steps composed of documents with variables (keys), ending with a “RETURN” value. The “FilteredPop” variable represents the subset of population (concept node) linked to the use case’s context  $Q_c$  (gender, age, etc.). The “AccessPop” variable is the subset of this population with convenient access to public transport where the link between the population (“FilteredPop”) and transports (“\$ComputeDistance”) is implied by the alignment of space and time values from each piece of data. The “\$ComputeDistance” operation calculates the proximity between population clusters (extracted from the “FilteredPop” variable) and public transport stops using the “StreetNetwork”. Finally, the “RETURN” statement computes the population proportion with access by dividing “AccessPop” by the total “FilteredPop”. The formula guides the combination of concept nodes (string values in red) in order to capture any required attribute. Data linkage through space and time is implied (and treated during the execution plan generation Section 4.4).

## 4 PROPOSED METHOD

Our method connects open data to SDG indicators through structured modeling and automated alignment. It is based on a dual graph approach. The first graph captures indicator structure; the second describes datasets and their metadata. These graphs are combined and queried to compute the indicator.

Figure 3 presents an overview of the process. It starts by building the Metadata Graph from available datasets. Then, the SDG Graph is generated from official documentation. The two graphs are aligned using LLM-assisted mapping. The resulting graph is filtered based on the user query. Finally, the indicator is computed using a rule-based execution plan.

### 4.1 Metadata Graph

The Metadata Graph (Metadata Graph) describes the structure and content of each dataset [6]. It includes nodes for datasets, their attributes, and metadata. Each attribute is enriched with dimensions such as spatial and temporal coverage, granularity, source reliability, and completeness. The Metadata Graph follows *Open SDG*<sup>1</sup> format. Open SDG is a standard data structure for SDG reporting. It is used by number of countries and institutions to publish their indicator data. The format defines how each dataset should be organized. It includes fields for indicator codes, dimensions, time periods, values and units. It also provides a way to include metadata such as definitions, sources and computation methods.

By following this structure, the Metadata Graph is compatible with official SDG data and makes it easier to connect new datasets.

This graph allows to make reasoning about data quality and relevance. It also enables flexible selection of sources based on user constraints. Similar metadata-driven approaches are used in prior works on data discovery and selection [12].

The Metadata Graph is designed to be queryable and extensible. For example, a user can retrieve all datasets that cover a specific region at a given time resolution. This structure also helps exclude sources that do not meet a minimum reliability threshold. For spatial and temporal descriptors, we rely on controlled vocabularies (e.g., NUTS levels for regions). When these are missing, we extract them using LLM-based parsing from headers or descriptions.

<sup>1</sup><https://open-sdg.readthedocs.io/en/latest/data-format/>

**Table 1: Operators' Definition Involved in Algorithm 1**

GetConcept	$\text{GetConcept}(D \in G^*, Q_I) = R$
Filter	$\theta_{\text{filter}_c}(R) = \{r \in R \mid \sigma_c(r)\}$
ComputeDistance	$\theta_{CD_\tau}(R_1, R_2) = \{(s_i, t_i, v_i) \in R_1 \mid \exists (s_j, t_j, v_j) \in R_2 \text{ such that } \text{dist}(s_i, s_j) \leq \tau\}$
Aggregation	$\theta_{\text{agg}}(R) = r'$ , where $r'$ summarizes $R$ under $\oplus \mid \oplus : \mathcal{P}(R) \rightarrow \mathbb{R}^m, \oplus(R) = \text{agg}(v_1, v_2, \dots, v_n)$
Align	$\theta_{\text{Align}}(R_1, R_2) = \{(s, t, (v_1, v_2)) \mid (s, t, v_1) \in R_1, (s, t, v_2) \in R_2\}$

## 4.2 Automated Mapping

Once both graphs are created, we map dataset attributes to SDG concepts. This step is partially automated using an LLM. It helps resolve naming mismatches and align semantics. For example, it maps geometry to ArRGeopoint, or bus to low capacity.

The mapping involves two levels. First, we align attributes to SDG concepts. Then, we align value vocabularies. The goal is to translate dataset specific codes into SDG compliant terms. LLMs have been shown effective for this task [11], especially when metadata is sparse. All mappings are stored in the graph, with scores (see below) and traceability. They can be reviewed and refined by experts if needed.

To ensure mapping quality, each link generated by the LLM is assigned a confidence score. This score is computed using cosine similarity between the embedding of the predicted label and that of the target concept in the SDG Graph. It captures how close the two terms are semantically, even if they differ lexically. For example, `pop_acc_lc_500m` is correctly matched to Proportion of population with access to low-capacity public transport within 500 meters. Low-confidence mappings (e.g., *score* < 0.7) are flagged for expert review. This semi-automated process balances precision and scalability. In our experiments, LLMs outperform traditional string similarity methods, especially in multilingual cases or when values are abbreviated. This mapping step allows us to align the SDG Graph and Metadata Graph, producing the unified and executable graph: SDG-KG.

## 4.3 Graph Filtering

In order to answer a query from a given indicator, space and time parameters, we need to extract relevant parts of the SDG-KG to focus on proper open data values. The SDG-KG  $\mathcal{G}$  helps to provide a set of solution subgraphs  $\mathcal{G}_I \subset \mathcal{G}$  where the indicator node  $I$  (from use case  $Q_I$ ), in the SDG Graph, is linked to source nodes  $D \subset \mathcal{D}$  in the Metadata Graph.

The goal of this step is to select the most relevant graph  $G^* \subseteq \mathcal{G}_I$  containing the best set of sources  $D$  among the available sources  $\mathcal{D} \subset \mathcal{G}_I$ . To achieve this, we propose the *Trust Score Metric (TSM)* that quantifies the alignment between a dataset's metadata and the user-defined query parameters.

*Definition 4.1.* Let  $D \subseteq \mathcal{D} \subset \mathcal{G}_I$  be a dataset with metadata  $D = \{D_s, D_t, D_c, D_r, D_p\}$  respectively representing its spatial, temporal, contextual granularity, reliability and completeness information. Let  $Q_I = \{q_s, q_t, q_c\}$  be a query specifying spatial, temporal, and

contextual constraints. The *TSM* of  $D$  regarding  $Q_I$  is defined as:

$$TSM(D, Q_I) = \frac{\sum_{i \in \{s,t,c\}} [w_i \cdot \text{Sim}(D_i, q_i)] + w_r \cdot u(D_r) + w_p \cdot v(D_p)}{\sum_{i \in \{s,t,c,r,p\}} w_i}$$

where:

- $w_s, w_t, w_c, w_r, w_p$  represent user-defined normalized weights [10] controlling the importance of each factor,  $\sum w_i = 1$ ,
- $\text{Sim}(D_i, q_i)$  is the similarity function for each dimension (spatial  $s$ , temporal  $t$ , contextual  $c$  -  $q_i \in Q_I$ ),
- $u(D_r)$ : Reliability score, assessing the trustworthiness of the dataset source (e.g., governmental, NGO, crowdsourced data).  $u(D_r) \in [0, 1]$ ,
- $v(D_p)$ : Completeness score, indicating the proportion of expected records available in the dataset.  $v(D_p) \in [0, 1]$ .

The trust score serves as a ranking function over all possible dataset combinations. This is important when several datasets cover the same concept and parameters (space and time from  $Q_I$ ) but differ in quality or relevance. For example, two datasets may both contain population data, but only one may provide disaggregated information by age or sex. Our metric ensures that the selected subgraph  $G^*$  contains the most relevant and complete sources for computing the requested indicator.

The resulting *TSM* score is normalized between 0 and 1, where higher values indicate a closer match to the query. This formulation ensures that selected datasets for indicator computation best align with user-defined spatial, temporal, and contextual constraints with respect to the source reliability and completeness score.

The problem of finding the best candidate graph  $G^*$  can be formalized as the following:

$$G^* = \arg \max_{G \subset \mathcal{G}_I} \prod_{D \subseteq \mathcal{D} \cap G} TSM(D, Q_I) \quad (1)$$

## 4.4 Execution Plan Generation

Once the indicator node  $I$  and a use case  $Q_I$  are chosen, as well as the best subgraph solution  $G^* \subseteq \mathcal{G}_I$ , it requires to translate this subgraph into an execution plan that links open data sources  $D$  to the  $I$  node's formula computation  $f_I$ . The indicator's formula  $f_I$  given as a JSON document (Section 3.2) is composed of explicit (in the formula:  $\$Filter$ ,  $\$ComputeDistance$ ,  $\$Div$ , etc.) and implicit operations (*Align*, *SUM*, etc.) applied on nodes (extracted by *GetConcept* and variables which structures the execution plan.

Each concept from the SDG Graph is linked to at least space ( $s$ ), time ( $t$ ) attributes and value nodes ( $v$ ) (cf. Section 3.1). Thus, the use case  $Q_I$  is applied on each linked source to get corresponding data  $R = [s, t, (v)]$ .

**Algorithm 1** Generated execution plan for Indicator 11.2.1**Input:** A Solution Graph:  $G^*$ , Query Parameters  $Q_I = \{q_s, q_t, q_c\}$ **Step 1: Retrieve concepts from the graph**

- 1:  $PT \leftarrow \text{GETCONCEPT}(G^*, \text{"PublicTransport"}, q_s, q_t)$
- 2:  $Pop \leftarrow \text{GETCONCEPT}(G^*, \text{"Population"}, q_s, q_t)$
- 3:  $SN \leftarrow \text{GETCONCEPT}(G^*, \text{"StreetNetwork"}, q_s, q_t)$

**Step 2: Filter Population**

- 4:  $\text{FilteredPop} \leftarrow \text{FILTER}(Pop, q_c)$

**Step 3: Compute Service Area & Population Within**

- 5:  $\text{ServiceArea} \leftarrow \text{COMPUTEDISTANCE}(PT, SN, 500)$
- 6:  $\text{AccessPop} \leftarrow \text{ALIGN}(\text{FilteredPop}, \text{ServiceArea})$

**Step 4: Compute Indicator Value**

- 7:  $\text{FilteredPopV} \leftarrow \text{SUM}(\text{GETATTRIBUTE}(\text{FilteredPop}, \text{"Value"}))$
  - 8:  $\text{AccessPopV} \leftarrow \text{SUM}(\text{GETATTRIBUTE}(\text{AccessPop}, \text{"Value"}))$
- RETURN** ( $\text{AccessPopV} / \text{FilteredPopV}$ )

The execution plan  $\theta_I$  is composed of operations  $\theta \in \Theta$  to compute the indicator value (defined in Table 1). Each operator  $\theta$  is defined as a transformation  $\theta : R \rightarrow R$  and the execution plan is a composition of such operators:

$$\theta_I = \theta_n \circ \theta_{n-1} \circ \dots \circ \theta_1$$

Our approach aims at generating execution plans like Algorithm 1 relying on the SDG-KG graph. The execution plan is generated by applying the following steps:

- First each concept node is identified from Indicator nodes (graph traversal) and apply *GetConcept* to apply use case’s space and time constraints on  $G^*$  (Step 1);
- For each key value pairs (except RETURN) generates a variable with the corresponding operations (Steps 2 and 3). This step is recursive (e.g., lines 5 and 6). Notice that implied alignment operators (list of variables & operations) link data on related space areas  $s$  and time period  $t$  (Line 6 & Table 1);
- Finally, the RETURN step seeks for Value data (*GetAttribute*) with a graph traversal to mapped nodes in the Metadata Graph. By default a sum<sup>2</sup> of values is computed. The final computation formula is applied with previously computed variables.

The output of our approach is a source code like in Algorithm 1 that computes the indicator on a given use case  $Q_I$ .

Notice that the *GetAttribute* operation handles sources  $D \subset G^*$  from the Metadata Graph with various capabilities (SQL, NoSQL, CSV). Implementation details are out of the scope this article.

Thanks to our approach composed of operators and the final execution plan, we can infer some properties: **compositionality** (each  $\theta_i$  is closed over  $R$ , enabling modular and traceable pipelines), **reproducibility** ( $\theta_I$  is fully determined by  $G^*$  and the indicator formula  $f_I$ ), **inspectability** (intermediate results at each  $\theta_i$  can be extracted for validation or adjustment).

## 5 EXPERIMENTS

In this section, we evaluate our framework by answering the following Research Questions (RQs):

- **RQ1:** Does the SDG-KG capture the computation logic and data structure better than existing SDG models?
- **RQ2:** Does the SDG-KG accurately map dataset attributes and values to SDG concepts using LLMs?
- **RQ3:** Does the TSM select the most relevant datasets based on user constraints?
- **RQ4:** Can the SDG-KG framework compute indicator values from aligned open datasets?

### 5.1 Experimental Setup

We evaluate our framework on a series of experiments which simulate the computation of a real SDG indicator using open data. We structure the evaluation around the main components of the framework: the SDG-KG modeling, the LLM-based mapping, the source filtering using metadata, and the execution pipeline.

*Datasets.* We collect open datasets from various national and local data portals (6 countries), covering themes such as mobility, demography, and environment from national state sources. Datasets are provided in formats such as CSV or GeoJSON and typically contain attributes related to space, time, and values. Metadata is extracted when available or manually annotated if missing.

To provide context on the underlying data used for each country, Table 2 summarizes the volume and source of the datasets. The number of rows, attributes, and file sizes vary but remain tractable for automated computation.

*Indicators.* We rely on the official SDG Indicator Metadata Repository<sup>3</sup> to define indicator structures in the SDG Graph. This documentation provides detailed descriptions of indicators, including formulas, disaggregation dimensions, and reporting requirements. These definitions are used to guide the modeling, mapping, and computation processes, ensuring compatibility with SDG standards.

*Environment.* All experiments are run on a machine with an 8-core CPU and 32 GB RAM. The graph is stored and queried using Neo4j 5.0. The entire framework is implemented in Python. For the LLM-assisted mapping, we tested several models, including OpenAI’s GPT-3.5 and GPT-4, to evaluate robustness and consistency. Cypher queries integrating the TSM function are applied in Neo4j to extract subgraph candidates and produce  $G^*$  (best score). Then it traverses this subgraph during the execution plan generation in order to define corresponding operators and the generated indicator computation. The output is a Python script that can be evaluated.

*Evaluation Tasks.* We organize the evaluation around the four research questions introduced earlier. To assess RQ1, we compare our SDG-KG with existing SDG models. We analyze their ability to represent indicator formulas, required variables, and disaggregation dimensions, as well as their potential for mapping data. For RQ2, we evaluate the accuracy of LLM-based mapping by comparing predicted attribute and value alignments to a manually curated gold standard. We report precision, recall, and F1-scores. For RQ3, we validate the effectiveness of the *Trust Score Metric* (TSM) by issuing user queries with different spatial, temporal, and contextual constraints, and checking if the selected datasets match expert expectations. We also test the sensitivity of the metric to the weight

<sup>2</sup>The number of rows in the Open SDG format depends on the level of granularity of geographic areas (national, regional, city, etc.)

<sup>3</sup><https://unstats.un.org/sdgs/metadata/>

**Table 2: Overview of open data sources and volume per country. Format: rows / columns / size**

France (fr)	USA (us)	Germany (de)	UK (uk)	South Korea (kr)	Australia (au)
INSEE / data.gouv.fr	Census / data.gov	Destatis / govdata.de	ONS / data.gov.uk	KOSTAT / KOSIS	ABS / data.gov.au
5.2K / 120 / 80MB	6.1K / 140 / 95MB	4.5K / 85 / 70MB	4.9K / 90 / 72MB	5.8K / 110 / 100MB	3.2K / 100 / 65MB

**Table 3: Coverage of SDG-related Models (focus on SDG-11)**

Model	#Total Terms	#SDG-11 Terms	#11.2.1 Terms
LinkedSDGs	285	29	10
SDGIO	463	35	2
PPMI	1,062	102	4
Science4SDGs	3,040	217	12
SDG Ontology	4,122	316	27
SDG-Pathfinder	4,456	225	8
FP-4-SD	5,594	112	8
EuroStat	26,741	1,023	22
SDG-KG	<b>52,753</b>	<b>2,977</b>	<b>43</b>

parameters. Finally, for RQ4 we test whether the framework can compute correct indicator values using aligned datasets. We select a set of indicators in the SDG Graph and compare the computed results to verified outputs. More details are available on Github<sup>4</sup>.

## 5.2 Comparison of existing SDG Models (RQ1)

We evaluate whether SDG-KG captures the computation logic and structure of SDG indicators better than existing models. Most previous ontologies were built for text classification or documentation not for operating indicator formulas over data.

We analyze eight major SDG-related models: *SDG Ontology*, *SDGIO*, *FP-4-SD*, *LinkedSDGs*, *SDG-Pathfinder*, *Science4SDGs*, *PPMI*, and *EuroStat*. Table 3 compares them with SDG-KG based on their number of total concepts, SDG-11 terms, and terms related specifically to Indicator 11.2.1. A higher number of terms suggests greater semantic coverage, which may improve mapping potential, though this depends on term quality.

Most existing models focus on goals or targets. Few go deeper to support actual indicator level variables or computations. For example, *SDGIO* and *OSDG 2.0* are mainly used for text classification. They include keyphrases, but not structured entities or formulas. *LinkedSDGs* and *FP-4-SD* provide some mappings but lack operational concepts like walking distance or transport capacity.

General purpose ontologies like Wikidata, Schema.org, and DBpedia or even SustainGraph [7] are not designed for SDG computation. They contain many relevant concepts, but lack structure. We tested a hybrid approach: we used 62 SDG-11 keyphrases from *OSDG 2.0* to prompt LLMs over Wikidata and Schema.org. The output often included vague or high-level matches (e.g., “urban mobility”) with no operational attributes.

In contrast, SDG-KG is modeling concepts, attributes, and disaggregation variables, along with formulas and constraints for computing indicators. It supports spatial buffers, value thresholds, and demographic filters. As far as we know, this makes it the only model that is fully computation-ready.

<sup>4</sup><https://anonymous.4open.science/r/SDG-KG-CED3/>

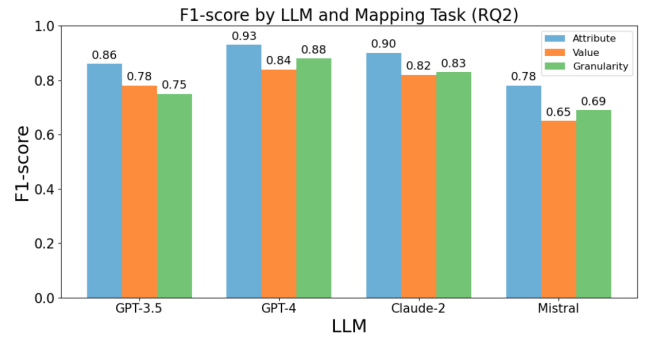
## 5.3 LLM-based Mapping Accuracy (RQ2)

We evaluate whether SDG-KG can accurately map dataset attributes and values to SDG concepts using Large Language Models (LLMs) [24]. This step is critical to align heterogeneous open data to the SDG Graph without manual engineering.

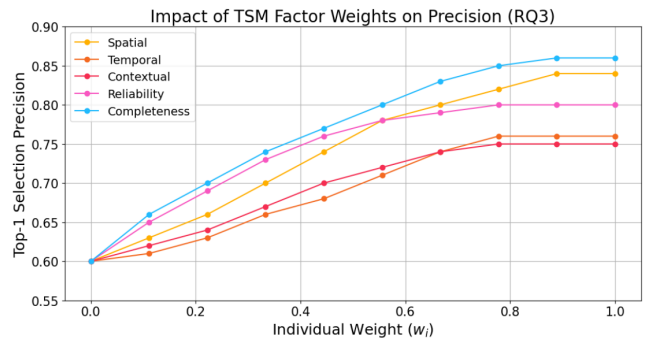
We manually curated a gold standard of 70 data source mappings covering different types of alignment: attribute names (e.g., geometry), values (e.g., bus → low capacity), and demographic or spatial granularity. Each mapping was validated by domain experts and used to benchmark several mappings using LLMs.

We tested four LLMs: GPT-3.5, GPT-4, Claude-2, and Mistral. Each model was evaluated under a few-shot setting with 3 guiding examples. We measured the F1-score for each mapping task.

Figure 4 shows that GPT-4 consistently outperforms the other models, especially for granular alignment and semantic value mapping. Attribute name mappings are generally easier for all LLMs.



**Figure 4: F1-score by LLM and mapping task (RQ2). GPT-4 performs best across all dimensions**



**Figure 5: Influence of individual TSM weights on Top-1 selection precision for SDG 11.2.1. Each curve shows the precision when progressively increasing the weight of a single factor from 0 to 1.**

Claude-2 performs well but struggles slightly with ambiguous values. Mistral, while open source, shows limited performance due to its smaller scale.

These results confirm that advanced LLMs, when guided by a structured SDG Graph and minimal prompting, can reliably perform semantic alignment across heterogeneous datasets. All mappings are recorded in the graph along with a confidence score, estimated from output probabilities or semantic similarity to the target concept, and can be revised by experts when needed.

#### 5.4 Metadata-based Source Selection (RQ3)

We now assess the effectiveness of the Trust Score Metric (TSM), introduced in Section 3.3, for automatically selecting the most relevant subgraph linking to mapped datasets based on user-defined metadata constraints. This experiment focuses specifically on one indicator (SDG 11.2.1: access to public transport) to isolate the behavior of the TSM under different conditions.

Given a pool of 20 datasets, each describing transport-related variables with distinct metadata profiles, we simulate 12 user queries ( $Q_1$  to  $Q_{12}$ ). Each query specifies different spatial, temporal, and contextual requirements (e.g., “urban scale”, “2020–2022”, “low capacity transport”). For each query, the TSM ranks the datasets according to how well their metadata align with the specified constraints. The dataset with the highest rank is then compared to the expert preferred choice.

To evaluate the quality of the ranking, we report the best accuracy, and the average TSM score gap between the selected and the optimal candidate. These metrics allow us to quantify the robustness of metadata-based filtering.

We analyze how the weighting of each TSM dimension influences selection accuracy. Each factor – spatial granularity ( $w_s$ ), temporal coverage ( $w_t$ ), contextual scope ( $w_c$ ), source reliability ( $w_r$ ), and completeness ( $w_p$ ) – can be weighted according to user preferences. We vary one weight from 0 to 1 while proportionally reducing the others, and measure the resulting Top-1 precision for indicator 11.2.1.

As shown in Figure 5, the completeness and reliability dimensions have the strongest impact on source selection. Prioritizing completeness significantly improves accuracy, as complete datasets typically align better with user needs. Similarly, datasets from trusted sources (e.g., official statistical agencies) tend to be more relevant when reliability is emphasized.

Spatial and contextual dimensions improve precision but reach a saturation point. This suggests that fine-grained granularity helps, but is not sufficient on its own. The temporal factor has a more moderate influence, likely due to overlapping time spans among datasets.

Overall, these results demonstrate that metadata alone can guide effective and interpretable dataset selection. The TSM provides a tunable scoring mechanism that allows users to express their priorities explicitly.

Moreover, in Figure 5 we observe that Top-1 precision never reaches 100%. This may be due to the large number of candidate terms in the graph (over 53K), increasing ambiguity in selection.

#### 5.5 Indicator Computation Accuracy (RQ4)

To evaluate whether SDG-KG can compute accurate indicator values from open datasets, we simulate the full execution of selected SDG indicators using real world sources. This includes data alignment, query generation, execution, and comparison to reference values. The goal is to validate both the formal expressiveness and practical reliability of the framework. For that, we have used the previous results for the TSM score in order to maximize the computed value.

We focus on two widely used indicators from SDG 11 (Sustainable Cities and Communities): indicator 11.2.1, which measures the proportion of population with access to public transport, and indicator 11.6.2, which assesses the average level of  $PM_{2.5}$  air pollution in urban areas. These indicators require spatial filtering and integration of heterogeneous datasets, making them ideal to evaluate SDG-KG in realistic scenarios.

After aligning the SDG Graph structure for those indicators, the framework automatically generates a computation plan and executes it on the mapped data.

The computed results are then compared against the ground truth: reference values published by the World Bank<sup>5</sup>. We select six countries that offer reliable open data and are covered by World Bank indicators: France, USA, Germany, UK, South Korea, and Australia.

Figure 6 presents a visual comparison of computed values by SDG-KG versus official values for the two indicators across both years. Subfigures (a) and (b) correspond to 11.2.1, while (c) and (d) present results for 11.6.2. In each plot, computed values (SDG-KG) are shown alongside World Bank references.

In most cases, computed values are consistent with official data. For indicator 11.2.1, the relative error is below 3% across all countries. For indicator 11.6.2, which is more sensitive to data granularity and completeness, the relative error remains below 7%. The absolute error is also generally low, staying under 2 units. It is worth noting that official statistics often include domain-specific corrections or estimation models that adjust for local coverage gaps or reporting inconsistencies. These corrections, applied by national institutes, can further refine the reported values. SDG-KG does not replicate these adjustments but provides a transparent and reproducible alternative based only on open data.

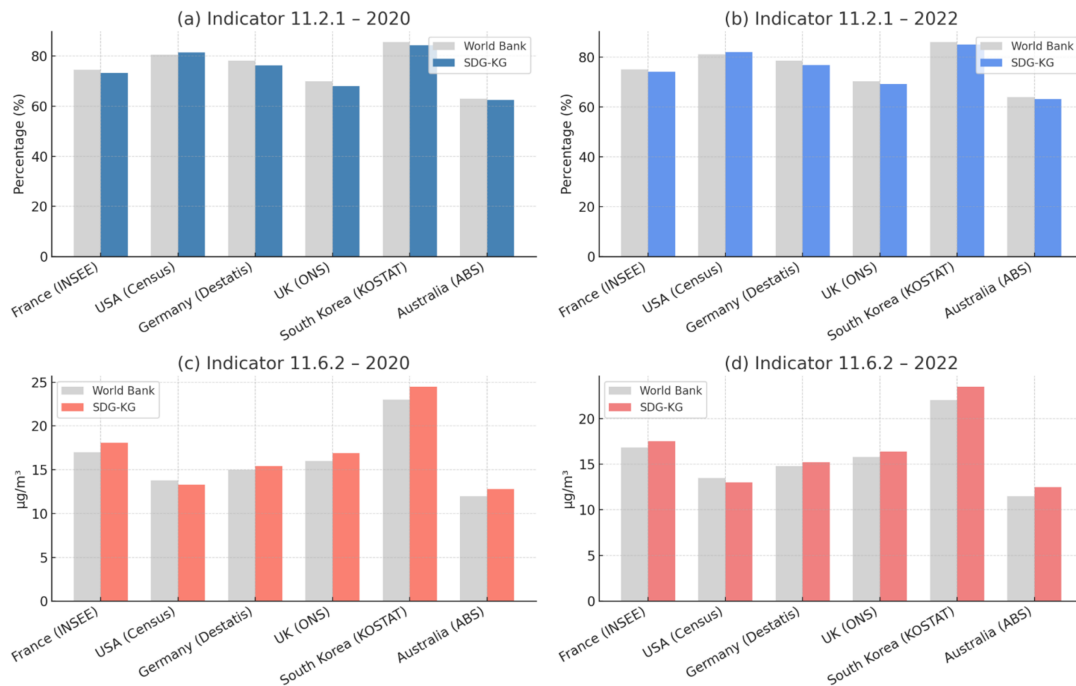
These results confirm that SDG-KG can operate reliably on open, heterogeneous data, and compute SDG indicator values that are both formally correct and quantitatively close to international references. The framework scales to different data formats, structures, and volumes with minimal human intervention.

## 6 CONCLUSION

In this work, we introduced SDG-KG, a knowledge graph based framework for computing SDG indicators from open and heterogeneous data. The approach relies on a dual graph model: one for indicator semantics and one for dataset metadata. It supports automatic mapping, trust based source selection, and indicator execution. Experiments on real world data show that SDG-KG produces accurate values, enables interpretable source filtering, and scales across diverse contexts. These results confirm that open data and

<sup>5</sup>[https://databank.worldbank.org/source/sustainable-development-goals-\(sdgs\)](https://databank.worldbank.org/source/sustainable-development-goals-(sdgs))

## Comparison of SDG-KG Computed Values and World Bank References – 2020 &amp; 2022 (RQ4)



**Figure 6: Comparison of SDG-KG computed values and World Bank references for indicators 11.2.1 and 11.6.2 in 2020 and 2022 across six countries.**

structured metadata can support scalable and explainable indicator computation.

Future work includes integrating domain specific statistical corrections, such as those applied by national agencies, to refine computed values. We also plan to support indicator computation at multiple spatial and temporal granularities that represent an underexplored aspect in current literature. Finally, we aim to model and explain the impact of major disruptive events (e.g., pandemics, policy shifts) on indicator trends, using knowledge propagation within the graph.

## REFERENCES

- [1] Christian Bizer, Tom Heath, and Tim Berners-Lee. 2009. Linked Data—The Story So Far. *International Journal on Semantic Web and Information Systems (IJSWIS)* 5, 3 (2009), 1–22.
- [2] Diego Calvanese, Benjamin Cogrel, Simon Komla-Ebri, Roman Kontchakov, Davide Lanti, Maria Rezk, Mariano Rodriguez-Muro, and Guohui Xiao. 2017. Ontop: Answering SPARQL queries over relational databases. In *Semantic Web*, Vol. 8. IOS Press, 471–487.
- [3] Tim Davies and Stephen Walker. 2019. The state of open data: Histories and horizons. *African Minds* (2019).
- [4] Xin Dong, Evgeniy Gabrilovich, Jeremy Heitz, Wilko Horn, Ni Lao, Kevin Murphy, Thomas Strohmman, Shaohua Sun, and Wei Zhang. 2014. Knowledge vault: a web-scale approach to probabilistic knowledge fusion. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (New York, USA) (KDD '14)*. ACM, 601–610. <https://doi.org/10.1145/2623330.2623623>
- [5] Xin Luna Dong and Divesh Srivastava. 2015. *Motivation: Challenges and Opportunities for BDI*. Big Data Integration. Springer International Publishing, 1–29. [https://doi.org/10.1007/978-3-031-01853-4\\_1](https://doi.org/10.1007/978-3-031-01853-4_1)
- [6] Rebecca Eichler, Corinna Giebler, Christoph Gröger, Holger Schwarz, and Bernhard Mitschang. 2021. Modeling metadata in data lakes—A generic model. *Data Knowledge Engineering* 136 (2021), 101931. <https://doi.org/10.1016/j.datak.2021.101931>
- [7] E. Fotopoulou, I. Mandilara, A. Zafeiropoulos, C. Lapidou, Gi. Adamos, P. Koundouri, and S. Papavassiliou. 2022. SustainGraph: A knowledge graph for tracking the progress and the interlinking among the sustainable development goals' targets. *Frontiers in Environmental Science* 10 (2022).
- [8] A. Hogan, E. Blomqvist, M. Cochez, C. D'amato, G. De Melo, C. Gutierrez, S. Kirrane, J. E. L. Gayo, R. Navigli, S. Neumaier, A.-C. N. Ngomo, A. Polleres, S. M. Rashid, A. Rula, L. Schmelzeisen, J. Sequeda, S. Staab, and A. Zimmermann. 2021. Knowledge Graphs. *ACM Comput. Surv.* 54, 4 (jul 2021), 37 pages.
- [9] Marijn Janssen, Yannis Charalabidis, and Anneke Zuiderwijk. 2012. Open data as a tool for innovation and accountability: A typology of governance models. *Government Information Quarterly* 29, 4 (2012), 258–268.
- [10] R. L. Keeney, H. Raiffa, and David W. Rajala. 1979. Decisions with Multiple Objectives: Preferences and Value Trade-Offs. *IEEE Transactions on Systems, Man, and Cybernetics* 9, 7 (1979), 403–403.
- [11] Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. In *Proceedings of the 36th International Conference on Neural Information Processing Systems (New Orleans, LA, USA) (NIPS '22)*. Curran Associates Inc., Red Hook, NY, USA, Article 1613, 15 pages.
- [12] Fatemeh Nargesian, Erkang Zhu, Chen Pu, Renée J Miller, and Xin Luna Dong. 2018. Data lake management: challenges and opportunities. *Proceedings of the VLDB Endowment* 11, 12 (2018), 2022–2025.
- [13] United Nations. 2015. Transforming our World: The 2030 Agenda for Sustainable Development. (2015).
- [14] N. Noy, Y. Gao, A. Jain, A. Narayanan, A. Patterson, and J. Taylor. 2019. Industry-scale knowledge graphs: lessons and challenges. *Commun. ACM* 62, 8 (July 2019), 36–43. <https://doi.org/10.1145/3331166>

- [15] Shirui Pan, Linhao Luo, Yufei Wang, Chen Chen, Jiapu Wang, and Xindong Wu. 2024. Unifying Large Language Models and Knowledge Graphs: A Roadmap. *IEEE Transactions on Knowledge and Data Engineering* (2024), 1–20. <https://doi.org/10.1109/tkde.2024.3352100>
- [16] M. Parciak, B. Vandervoort, F. Neven, L. M. Peeters, and S. Vansummeren. 2024. Schema Matching with Large Language Models: an Experimental Study. In *TaDA Workshop at VLDB'24*. VLDB.org.
- [17] László Pintér, Peter Hardi, and Peter Bartelmus. 2021. Use of Indicators in Strategic Environmental Assessments of Urban-Planning Instruments: A Case Study. *Sustainability* 13, 22 (2021), 12639.
- [18] Erhard Rahm and Philip A Bernstein. 2001. A survey of approaches to automatic schema matching. *The VLDB Journal* 10, 4 (2001), 334–350.
- [19] Henrique Santos, Victor Dantas, Vasco Furtado, Paulo Pinheiro, and Deborah L. McGuinness. 2017. From Data to City Indicators: A Knowledge Graph for Supporting Automatic Generation of Dashboards. In *The Semantic Web: 14th International Conference, ESWC'17* (Portorož, Slovenia). Springer-Verlag, 94–108. [https://doi.org/10.1007/978-3-319-58451-5\\_7](https://doi.org/10.1007/978-3-319-58451-5_7)
- [20] Patrick Schopf, Lorenz M. Hilty, and Bastian Wepner. 2020. Big Data Management: Data Governance Principles for Big Data and Open Data. *IEEE Access* 8 (2020), 193–201.
- [21] Rajib Singh, Hrudaya N. Murty, Sunil K. Gupta, and Anil K. Dikshit. 2017. The Role of Indicator-Based Sustainability Assessment in Policy and Decision-Making. *Sustainability* 9, 6 (2017), 1018.
- [22] Gerhard Weikum. 2021. Knowledge Graphs 2021: A Data Odyssey. *Proceedings of the VLDB Endowment* 14, 12 (2021), 3233–3238.
- [23] Guohui Xiao, Linfang Ding, Benjamin Cogrel, and Diego Calvanese. 2019. Virtual Knowledge Graphs: An Overview of Systems and Use Cases. *Data Intelligence* 1, 3 (06 2019), 201–223. [https://doi.org/10.1162/dint\\_a\\_00011](https://doi.org/10.1162/dint_a_00011)
- [24] Hui Yin, Amir Aryani, and Nakul Nambiar. 2024. Evaluating the Performance of Large Language Models for SDG Mapping (Technical Report). arXiv:2408.02201 [cs.LG] <https://arxiv.org/abs/2408.02201>
- [25] Yunjia Zhang, Avriela Floratou, Joyce Cahoon, Subru Krishnan, Andreas C. Müller, Dalitso Banda, Fotis Psallidas, and Jignesh M. Patel. 2023. Schema Matching using Pre-Trained Language Models. In *2023 IEEE 39th International Conference on Data Engineering (ICDE)*. 1558–1571. <https://doi.org/10.1109/ICDE55515.2023.00123>
- [26] Anneke Zuiderwijk, Marijn Janssen, Sunil Choenni, Ronald Meijer, and Rachel SC Alibaks. 2014. Innovation with open data: Essential elements of open data ecosystems. *Information Polity* 19, 1-2 (2014), 17–33.