



Research paper

KD SENSO-MERGER: An architecture for semantic integration of heterogeneous data

Yoan Gutiérrez^{a,*}, José I. Abreu Salas^a, Andrés Montoyo^a, Rafael Muñoz^a,
Suilan Estévez-Velarde^b

^a University Institute for Computing Research, University of Alicante, Carretera San Vicente del Raspeig s/n, 03690, Alicante, Spain

^b Artificial Intelligence and Computing Systems, University of Havana, San Lázaro y L. Edificio Felipe Poey. Plaza de la Revolución, Havana, Cuba



ARTICLE INFO

Keywords:

Heterogeneous data
Knowledge discovery
NERC
Natural language processing
Ontology and knowledge representation
Semantic data integration

ABSTRACT

This paper presents KD SENSO-MERGER, a novel Knowledge Discovery (KD) architecture that is capable of semantically integrating heterogeneous data from various sources of structured and unstructured data (i.e. geolocations, demographic, socio-economic, user reviews, and comments). This goal drives the main design approach of the architecture. It works by building internal representations that adapt and merge knowledge across multiple domains, ensuring that the knowledge base is continuously updated. To deal with the challenge of integrating heterogeneous data, this proposal puts forward the corresponding solutions: (i) knowledge extraction, addressed via a plugin-based architecture of knowledge sensors; (ii) data integrity, tackled by an architecture designed to deal with uncertain or noisy information; (iii) scalability, this is also supported by the plugin-based architecture as only relevant knowledge to the scenario is integrated by switching-off non-relevant sensors. Also, we minimize the expert knowledge required, which may pose a bottleneck when integrating a fast-paced stream of new sources. As proof of concept, we developed a case study that deploys the architecture to integrate population census and economic data, municipal cartography, and Google Reviews to analyze the socio-economic contexts of educational institutions. The knowledge discovered enables us to answer questions that are not possible through individual sources. Thus, companies or public entities can discover patterns of behavior or relationships that would otherwise not be visible and this would allow extracting valuable information for the decision-making process.

1. Introduction

The current challenges of a digital society require automatic knowledge representation and discovery, which explains why it is one of the most rapidly growing fields of research in computer and information science. Today, misinformation causes serious problems, because decision-making procedures require analysis of full data, and not partially, “data is the new oil”. The information needed by organizations is scattered all over the Web in an unstructured format, making it difficult to integrate into decision making and data analysis platforms and to feed predictive models about peoples’ behavior, for example (i.e., risk detection, prevention of cyberbullying, terrorist warnings, etc.). Incorporating this knowledge into information systems requires the development of methodologies that discover relevant information and continuously update it, by enriching and integrating knowledge from various sources of structured and unstructured data. Moreover, the speed at which new information is produced poses new challenges,

in terms of the continuous processing, evaluation, and refinement of the acquired knowledge.

This surplus of information – news, email, social media, blogs – greatly exceeds the human capacity for processing and consuming data and is one of the challenging byproducts of the internet’s accelerated growth. Thus, building automatic systems that can extract knowledge from this flow of information has become one of the most actively researched fields in Computer Science. This process, involving Big Data (John Walker, 2014) has attracted the attention of diverse research communities, such as business intelligence, engineering, entertainment, e-commerce, and social media, among others (Wu et al., 2014; Chen et al., 2012; Shah et al., 2015). At the same time, it is evident that exploiting big data is a complex problem given the decentralization and non-natural integration of the data.

Research in machine learning, knowledge discovery, data mining, and more recently natural language processing has enabled the techniques and tools to handle the huge amount of information on the

* Corresponding author.

E-mail addresses: ygutierrez@dlsi.ua.es (Y. Gutiérrez), ji.abreu@ua.es (J.I.A. Salas), montoyo@dlsi.ua.es (A. Montoyo), rafael@dlsi.ua.es (R. Muñoz), sestev@matcom.uh.cu (S. Estévez-Velarde).

<https://doi.org/10.1016/j.engappai.2024.107854>

Received 17 January 2023; Received in revised form 13 November 2023; Accepted 3 January 2024

0952-1976/© 2024 Elsevier Ltd. All rights reserved.

Internet. Tasks such as the construction of search engines (Brin and Page, 1998) and recommender systems Davidson et al. (2010) are some examples. These systems could be used to improve business, health care, and policy decisions (Ferrucci et al., 2013).

As for the different approaches relevant to knowledge discovery, a continuous spectrum of techniques can be identified, based on how much expert knowledge is used. Knowledge-based techniques use rules handcrafted by domain experts (Chandrasekaran, 1986). These approaches have a great degree of reliability and precision and generally allow for more complexity in the extracted knowledge. However, applying knowledge-based techniques to large amounts of data does not guarantee accurate results because the knowledge base is limited.

By contrast, the statistical approaches consist of techniques based on pattern recognition with statistical and probabilistic models (Kevin, 2012). They perform better with large amounts of data (Le, 2013), providing better recall, but are often limited to extracting simple knowledge models, and can be more sensitive to noisy, fake, or biased information (Bolukbasi et al., 2016).

Given these mutually complementary characteristics inherent to both approaches, several hybrid approaches have been proposed. Recently, new research has emerged related to ontology learning (Cimiano et al., 2009), learning by reading (Barker et al., 2007) and entity embedding (Hu et al., 2015), where researchers combine techniques from machine learning, natural language processing and knowledge representation to solve more complex problems that cannot be dealt with by using only classical tools.

In addition, the design of non-monolithic learning systems, built as a set of modular components that can be combined in different ways, provides a novel approach to addressing the challenge of continuously improving learned knowledge. This challenge needs the development of computational tools that are capable of building ontologies that have automated or semi-automated processes.

This composability would allow a continuous learning system to not only improve the quality of the extracted knowledge but also to learn to tune its internal parameters - i.e. the selection of suitable technologies for the Natural Language Processing task - to perform better knowledge extraction in the future. It is conceivable that such a system could gradually learn what types of basic processes -e.g., entity recognition, POS tagging, etc. - are most useful for a given domain or a given corpus. Similarly, such a system could learn which types of probabilistic models provide the best results on a particular dataset.

Given this previous scenario, the continuous learning technologies fed by updated internet content may offer a potential solution for companies or public entities by discovering behavior patterns or relationships that would not otherwise be visible. This would allow for extracting valuable information for decision-making processes. In such a scenario, standardizing data becomes a pre-requisite for effective and accurate analysis.¹

Some of the major challenges faced by organizations when integrating heterogeneous data sources include:

- **Knowledge Extraction:** a complex and time-consuming task when data sources have different formats, structures, and types. Moreover, it requires the development of coordinated automatic tools to identify relevant information for each scenario. This work contributes to dealing with the challenge by proposing a plugin-based architecture of knowledge sensors that are capable of processing heterogeneous data.
- **Data Integrity:** Integrating data from a variety of sources can introduce contradictory and incomplete data. So, preventing this is crucial. Data quality is a primary concern in every data integration strategy. Poor data quality can be a compounding problem that can affect the entire integration cycle.

In this sense, our proposal deals with this challenge by developing the methodology and technologies to ensure the coexistence of multiple domains of knowledge dealing with contradictory facts. Also, our solution allows the exploration and dynamic interpretation of knowledge. Finally, we applied a set of metrics to ensure the integration quality.

- **Scalability:** The phenomenon of surplus of information is addressed via knowledge extraction sensors, i.e. NERCs (Li et al., 2022) can specialize the KD systems for specific domains. Data heterogeneity leads to the inflow of data from diverse sources into a unified system, which can ultimately lead to exponential growth in data volume.^{2,3}

Our plugin-based architecture tackles this challenge because only relevant knowledge of the scenario is integrated by switching off the non-relevant sensors. Also, we minimize the expert knowledge required, which may pose a bottleneck when integrating a fast-paced stream of new sources.

Through the use of Natural Language Processing (NLP) technologies, we can contribute to alleviating the major difficulties and challenges of managing knowledge and integrating heterogeneous sources. NLP techniques can be used to automatically extract relevant information from unstructured data sources such as text documents, emails, social media posts, etc. Thus, this can help to standardize and integrate data from various sources, making it easier to access and analyze.

The advantage of our proposal is the combination of techniques to group knowledge into the same structure from structured and unstructured data sources. We leverage techniques such as *field matching* to determine the pieces of the structured data that refer to the same concepts or relations. Also, different *information extraction* procedures may be used for unstructured data to discover relevant knowledge. Thus, our proposal creates a unique representation of knowledge by the integration of different heterogeneous sources such as tabular data or social networks, without the requirement of a predefined schema of data representation. In consequence, a Knowledge Graph (KG) is built from the data provided.

The **goal of our proposal** is to design an architecture for Knowledge Discovery that is able to semantically integrate heterogeneous data - both structured and unstructured - and build internal representations that can be adapted and integrated across multiple domains.

One of the main features of our proposal is its explicit handling of separate knowledge from structured and unstructured data, unlike other proposals in the literature, such as Never-Ending Learning (NELL) (Mitchell et al., 2018), (Jin et al., 2022) or Bootstrapping ontology evolution with multimedia information extraction (Petasis et al., 2011).

This approach enables the application of different techniques based on the specific domain. Additionally, it allows for the existence of contradictions or unreliable information over time, which can be verified in the future.

The paper is organized into the following sections to provide a comprehensive description of our proposal: Section 2 gives a brief overview of the state of the art in the automatic extraction of knowledge from different sources of information. Section 3, describes the proposed architecture for knowledge discovery. In Section 4, we present the methodologies that can be used to evaluate KD SENSOMERGER. In Section 5, we present a case study - combined demographic data, socio-economic data, and user reviews - where the architecture is applied. Finally, in Section 6, we present the main conclusions of the research and outline possible future lines of research.

¹ <https://www.dataversity.net/challenges-of-integrating-heterogeneous-data-sources/>

² <https://dzone.com/articles/3-challenges-of-integrating-heterogeneous-data-sou>

³ <https://ieeexplore.ieee.org/document/8396165>

2. State of the art

Early works such as Litwin et al. (1990) and Sheth and Larson (1990) illustrate that the inherent heterogeneity of schemes, sources, and formats for data and knowledge have driven interest in their integration for a long time.

In general, this process has to solve issues related to (i) collecting the data, (ii) extracting the knowledge, (iii) matching schemes, and (iv) merging (Fusco and Aversano, 2020). In this section, we discuss different proposals that address the aforementioned issues. To enhance the presentation, we differentiated approaches that explicitly handle the integration of heterogeneous data 2.1 from other approaches that are able to solve one or several issues but do not handle the whole process of heterogeneous integration.

2.1. Heterogeneous data and knowledge integration

The *mediator architecture* together with data warehousing, are two well-known approaches to integrating heterogeneous information. Typically, in mediation setups, the data remains in its sources, described by the so-called local schemes. Then, a global-mediated schema is designed and mapped to local schemes, to provide the unique entry point for queries. Usually, the interface between the mediator and the sources is provided by the so-called wrappers. Two main approaches have been studied. In global-as-view (GAV), the global schema is conceived as views over the sources. Meanwhile, local-as-view (LAV) defines the global schema independently of the sources, with local schemes as views of the global (Lenzerini, 2002). Regardless of the case, it is necessary to define a scheme of the data. The mediator provides the user with an abstraction of the underlying models of the sources, allowing working with a well-understood interface.

An architecture with similar characteristics but not described as a mediator is proposed in Adams et al. (2000) to integrate legacy databases from specific domains within Boeing corporation. A knowledge-based system serves as a common data model that hosts the inference engine, integrating the different databases. As the authors explained, technically skilled personnel are desirable to create and match from the global ontology to the legacy knowledge representations. Despite the authors acknowledging the convenience of supporting open information scenarios where new information sources may appear, the work did not cover knowledge extraction from sources other than the legacy databases studied.

Manually creating the schemes and mappings may pose a bottleneck to the integration process. Thus, it is desirable to automate it. In Reynaud et al. (2001) authors explored the automated integration of heterogeneous XML data sources using the mediator architecture. The SAMAG system finds mappings between Document Type Definitions (DTD) based on semantic and structural criteria. They leveraged WordNet (Fellbaum, 1998) to define semantic links between DTD terms through relationships such as synonymy, hyponymy, and meronymy.

The mediator architecture is also the basis of Bergamaschi et al. (2001). The Mediator environment for Multiple Information Sources (MOMIS) introduces the ODL_{j3} language as a means by which to build rich representations of source schemes. They curate a Common Thesaurus to encode inter-schema knowledge. It is made of terminological and extensional relationships between classes or attribute names. Four kinds of relationships are considered. Schema-derived comes from foreign keys in relational sources schemas. Lexically driven approaches are based on lexical relations between classes and attribute names. Both types encode intra-schema knowledge. On the other hand, designer-supplied inter-schema relationships are directly supplied by the expert to capture domain-specific knowledge. Finally, inferred relationships are introduced by subsumption from consistent extensional relationships validated by the expert. They also introduced the concept of affinity to measure the level of similarity between classes in different

source schemes. A core contribution is the use of a clustering algorithm based on affinity to identify candidate classes for integration and as a way of dealing with semi-structured data. The process is semi-automated, requiring expert knowledge to define the ODL_{j3} descriptions and validate relationships. The integration of unstructured data is not covered.

Defining the integration architecture as separated modules or layers has the advantage of providing modularization. By this means, roles within the process are identified and defined by their interfaces allowing easy replacement of components for improved versions. This is the approach taken in Vdovjak and Houben (2001) with their 5-layer RDF-based mediator architecture. The Source Layer acts as a wrapper of the sources being the main requirement for the ability to export the data to XML. The XML Instance Layer handles the inputs of the XML from the previous layer. The next layer, XML2RDF, provides the bridge between XMLs and the Mediator layer. The latter is at the core of the architecture. It hosts the Conceptual Model (CM) created beforehand through ontology engineering. At the top, rests the Application Layer, consisting of different applications leveraging the unified interface provided by the architecture. It is necessary to study how to handle unstructured data or cases where we cannot ensure the ability to export to XML. Moreover, as in the cases previously discussed, one potential drawback is the requirement of conceiving a conceptual model beforehand.

This is also the case of the more recent proposal described in Fusco and Aversano (2020). They generalized the GaV and LaV approaches using ontologies as a conceptual schema to represent both the global view and the data sources. They distinguish the four stages described next. Source Wrapping encompasses the creation of an ontology for each source to be integrated. The schema-matching step automatically searches for mappings between the different models, making it possible for the expert to validate and modify them if necessary. Based on these mappings, a global ontology is created at the Schema Merging stage, which also runs automatically. Finally, at the Query Reformulation step, a query targeted to the global view is reformulated to queries over the local sources. This approach scales easily to new sources since the process is mostly automatic, besides the tuning of some parameters for the matching and merging procedures. However, the creation of local ontologies is a probable bottleneck in cases of fast-paced streaming of new sources.

Ontologies are also a requirement of Obi-Wan, the framework for RDF integration discussed in Buron et al. (2020). They also leverage the mediator architecture following the Ontology-Based Data Access (OBDA) paradigm in a Global-Local-As-View (GLAV) approach, which generalizes both GAV and LAV. Heterogeneous sources are integrated into a virtual RDF graph. It is structured as RDFS⁴ ontology as well as data triplets extracted from the sources through GLAV mappings.

The cases discussed up to here do not fully address the integration of unstructured sources. Moreover, the knowledge discovery from these sources as part of the integration process needs to be studied further. In Nentidis et al. (2020), the authors deal with the topic of integrating unstructured sources and KD. The iASIS Open Data Graph implements a pipeline to automatically retrieve and integrate relevant knowledge in the biomedical domain. First, they integrated parallel harvesters for the unstructured content – literature – and structured knowledge, i.e. ontologies and databases. This addresses issues related to the data collection stage commented on at the beginning of Section 2. The Literature Harvester provides text content as well as article-topic relations that can be derived from text or metadata. Next is the Literature Analysis stage, where new relations are discovered using different text mining tools — a kind of sensor in our proposal. This module outputs *concept–concept* derived from the text as well as *mentioned in* relations. Meanwhile, the Structured Harvester acts as a gateway for *concept–concept* relations. Next, follows a Semantic

⁴ <https://www.w3.org/wiki/RDFS>

Integration component, integrating the concept–concept relations from the Structured Harvester as well as the article–topic relations from the Literature Harvester. This stage outputs *concept–concept* and *has_mesh*⁵ relations. In the end, iASIS Open Data Graph integrates the relevant knowledge. It is important to note that the process is designed to run automatically. They demonstrated the suitability of their approach in different case studies in the biomedical domain. Moreover, their architecture contemplates the automatic discovery of knowledge; however, the integration of new text mining tools, or sensors designed for other kinds of data such as geo-spatial or images are not covered.

In this section, we reviewed proposals addressing the integration of heterogeneous sources. Some interesting conclusions can be drawn from this analysis. First, the mediator seems to be the preferred approach for integration because it reduces redundancy, is easier to maintain, and provides real-time or near-real-time integration in the case of fully automatic approaches. It also tends to scale better in fast-changing environments. However, the necessity for further research in this area is underscored by the need to minimize the requirement of expert knowledge to build global and local ontologies. Moreover, it is convenient to fully automate the whole process. Also, it is key to deal with the scarcity of approaches that consider a knowledge extraction stage, or indeed the lack of studies that focus on open domains

Our proposal addresses these issues. First, we aim to fully automate the knowledge extraction challenge with a plugin architecture that enables the easy incorporation of knowledge extraction tools – sensors – into the system. Sensors can be designed for sources varying from text, images, or geographic data, thereby maximizing the degree of heterogeneity the system can handle. The sensor architecture also contributes to the scalability challenge. Irrelevant knowledge can be avoided by simply switching off the respective sensors. One of the main contributions is that our proposal tries to minimize the requirement of expert knowledge — besides the creation of the sensors. The ontology holding all the knowledge is created while the data is integrated into a fully automated dependency injection mechanism. This makes the approach independent of the domain by design. Of course, the nature and domains of applications will depend on the available sensors. A final comment is that we did not conceive our architecture as a mediator as previously described. It also has elements of the data-warehousing approach, such as in Nentidis et al. (2020), where relevant knowledge for the scenario is extracted and stored in a unified knowledge base.

The rest of the section is devoted to discussing approaches that are related to knowledge representation and reasoning (Section 2.2), the use of machine learning to extract knowledge (Section 2.3) and the KD process (Section 2.4). These areas are highly relevant to our proposal as they address the key issues that need to be solved to integrate information from heterogeneous sources.

2.2. Knowledge representation and reasoning

The problem of discovering, storing, and using knowledge in a computationally effective way has been extensively studied (Mitchell et al., 2018; Rospocher et al., 2016; Cimiano et al., 2009; Jin et al., 2022). This issue has been addressed from two different but complementary research areas: the fields of knowledge representation (Section 2.2) and machine learning (Section 2.3). The knowledge representation community provides the means to computationally represent and manage stored knowledge properly. Conversely, the machine learning community provides tools for deriving useful knowledge from large collections of structured and unstructured data. Also, there is a third field – Knowledge Discovery (Section 2.4) – which involves both knowledge representation and machine learning.

⁵ Medical Subject Headings (MESH) <https://www.nlm.nih.gov/mesh/meshhome.html>

Since the dawn of computer science, one of the problems that has attracted wide attention is that of representing knowledge in a computational format, such that automatic reasoning can be performed to discover new, previously unknown truths (Sowa et al., 2000).

Arguably, the most popular knowledge representation technology is the use of ontologies (Guarino, 1995) that have become the *de facto* standard. Ontologies can be defined as a formal specification of a conceptualization (Corcho et al., 2003). This represents concepts, relations between these concepts, instances of these concepts, and inference rules for deriving new relations.

As such, ontologies can be considered as a combination of two predominant approaches for knowledge representation: those based on formal logic (Brachman et al., 1992); and those based on graphs of semantic relations (Chein and Mugnier, 2008).

The primary objective of constructing these knowledge structures is to facilitate the exploration, comprehension, and reuse of information by both humans and machines across various applications. These applications include answering questions, locating pertinent content, comprehending social structures, and making scientific breakthroughs. However, the considerable scale and intricacy of these knowledge graphs pose a significant challenge, especially when it comes to mining information across diverse topic areas (Aggarwal et al., 2010).

2.3. Machine learning to extract knowledge

The field of machine learning provides tools for the automatic extraction of information and knowledge from different sources of data. This field not only permits the automation of processes and tasks associated with knowledge discovery or text mining but also provides a large improvement in the scalability of these processes (Wu et al., 2014).

Supervised and unsupervised learning are the two most prevalent approaches in machine learning (Kevin, 2012). Supervised learning can be used for recognizing specific elements of knowledge in a data source. For example, tagging pieces of text to indicate that they define an entity (Nadeau and Sekine, 2007) (e.g., a person, organization, or place), recognizing relations between said entities, or assigning a sentiment or opinion score (Liu, 2012) to a fragment of text. On the other hand, unsupervised learning can help identify relevant structures in a large set of data. Clustering algorithms can be used to detect similar concepts or to extract abstract concepts from groups of more concrete elements. Other techniques can be used for reducing the amount of information, for example, to remove noisy, uncertain, or irrelevant pieces of information (Bingham and Mannila, 2001).

In recent years, there has been an increased interest in the problem of automatically learning relevant representations. Word embeddings (Mikolov et al., 2013) and more general entity embeddings (Hu et al., 2015) represent the first step towards powering deep learning approaches with more explainable internal representations. Ontologies are, by definition, representations of a given conceptualization. Therefore, using them – or similar semantic representations as information seeds of a given domain – will enhance the performance of data mining processes based on machine learning

2.4. Knowledge discovery

Recently, a new discipline – ontology learning – has emerged, which draws ideas and techniques from both the knowledge representation and the machine learning fields. This discipline deals with the problem of automatically building ontological representations of knowledge from a variety of data sources. As such, the theory of ontology learning is relevant to the design of this proposal. Ontology learning searches to automate part of the process of creating and maintaining ontologies. This discipline has the potential of reducing the cost of creating and, most importantly, of maintaining large and complex ontologies (Cimiano et al., 2009). As a sub-discipline, Learning by reading (Barker et al.,

2007) is a field that draws techniques from natural language processing and knowledge representation and reasoning research areas. The purpose is to build a formal representation of fields given unrestricted related textual data. This representation must also allow fully automatic reasoning.

In ontology learning, two general high-level tasks can be distinguished: ontology population (Lubani et al., 2019) and ontology enrichment (Petasis et al., 2011).

Ontology population deals with the sub-problem of finding new instances for an already defined ontology, while ontology enrichment deals with adding new concepts and relations to an existing ontology. There is an overlap between these tasks, and most of the existing approaches cannot be classified purely in these terms. In this field, several tools have been proposed, that combine different approaches and solve different subsets of ontology learning tasks. Some of these systems are listed as follows:

- SYNDIKATE (Hahn and Romacker, 2001): an early approach to populating a knowledge base, with a predefined ontological structure (classes and relations).
- ARTEQUAKT (Alani et al., 2003), SOBA (Buitelaar et al., 2006) and, WEBtoKB (Craven et al., 2000): extracts knowledge and exploits the semi-structured format of web resources.
- VIKEF (Stoermer et al., 2006): Extracts knowledge from structured data.
- ADAPTATIVA (Brewster et al., 2002): provides a bootstrapping strategy, where human experts give feedback about the extracted knowledge.
- OPTIMA (Kim et al., 2008) and ISODLE (Weber and Buitelaar, 2006), OntoLT (Buitelaar and Sintek, 2004): describes basic NLP techniques to extract knowledge from text.
- LEILA (Suchanek et al., 2006), Text2Onto (Cimiano and Völker, 2005), KnowItAll (Etzioni et al., 2004): provides NLP techniques based on statistical models.
- OntoGain (Drymonas et al., 2010), ASIUM (Faure and Poibeau, 2000): extracts entities and relations from text, and builds a hierarchy of concepts.
- BOEMIE (Castano et al., 2007): infers abstract concepts from instances, on text, images, and videos.
- NELL (Mitchell et al., 2018): extracts knowledge continuously from a stream of web data.
- CogKGE (Jin et al., 2022): builds knowledge graph embedding.

In general, these tools are focused on the extraction of knowledge and on the task of discovering relevant knowledge. When extracting knowledge from a trustworthy source, even a natural language source, it makes sense to focus on optimizing recall, i.e., obtaining as much information as possible. When the input source consists of medical papers or the main web page of an institution, there is a high probability that most of the information present in those documents is correct. Hence, an ontology extraction procedure that maximizes recall will obtain good results.

However, when the input source is of lesser quality, such as blogs or social media posts, there is a greater likelihood that some, or even most, of the information is fake or incorrect. In this case, the information needs to undergo a deep semantic analysis to establish the writer's stance, i.e. positive, negative, etc. In this context, the problem of extracting useful knowledge from a large internet-based corpus is a problem of filtering and selecting relevant information.

Despite the existence of some general-purpose systems, so far no proposal has been identified that can learn from the multitude of sources of information by activating relevant and specialized extraction and mapping sensors (i.e., Named Entity Recognition, Geolocation mapping, Sentiment Analysis, Text Categorization, aspect-based semantic analysis, and many others) that will help to determine the intrinsic characteristics of the analyzed data.

Another challenge is to obtain a computational representation of this knowledge, independent of the domain, source, and format.

In previous works (Estevez-Velarde et al., 2018; Estévez-Velarde et al., 2020; Fernández et al., 2015; Balahur et al., 2011; Montoyo et al., 2012; Hermida et al., 2012) authors explored different NLP tasks related to the proposal, such as semantic knowledge discovery, semantic enrichment, Automatic Machine Learning (AutoML), the creation of a knowledge base to detect emotions and other related ones. In this work, we leverage these advances to solve the aforementioned challenges. We propose a novelty architecture for the semantic integration of heterogeneous data. Also, we present a case study to evaluate our approach.

3. Methodology

In this section, we present KD SENSO-MERGE, an architecture for the semantic merging of knowledge from heterogeneous data.

This architecture deals with the main issues described in Section 2, attempting not only to recognize information but also to reduce irrelevant information and obtain information with more relevance. Additionally, we design an architecture where new knowledge is discovered and incorporated into the existing knowledge base.

This process occurs in a feedback cycle, where new sources of information are processed and the extracted knowledge is merged with the existing knowledge. For this purpose, the framework requires three main modules, as shown in Fig. 1.

As shown in Fig. 1, the methodology proposed takes as starting point a document, which can be structured (1) or unstructured (1'), and depending on its typology, the corresponding module (Structured Data Processing or Unstructured Data Processing modules) is activated to convert it into an internal format that enables its processing by the Generation module. This Generation module applies different techniques to identify and extract information and generates a temporal semantic structure from the document (2). The Evaluation module takes as input this structure (3.1) and the previous knowledge existing in the system (3.2). The evaluation module uses field matching and information extraction techniques to link knowledge portions by adding non-existing connections. These new connections, depending on their quality degree,⁶ are returned to the generation module (3.3) to create the final knowledge base (4.1). Once, the Generation module produces this knowledge will be used as Previous Knowledge (4.2) for the next documents, and so on.

The modules responsible for processing structured and unstructured data are designed to extract knowledge from these specific sources. In both cases, the output is an ontology that represents the temporal knowledge acquired. This knowledge is not stored yet but sent to the module for knowledge discovery, where further processing is applied.

In addition to these main modules, three more modules (i.e. Algorithms, Long Term Memory(LTM); Organizational Ontology) are considered, which are interrelated. Each module within the architecture has a specific responsibility to define the inputs and outputs, enabling effective intercommunication with the other modules. Fig. 2 shows a general overview of the architecture.

The use of ontologies provides to the architecture the effective skills for representing knowledge in a wide variety of domains and scenarios (Staab and Studer, 2010). They are flexible enough to adapt to a particular domain and powerful enough to represent complex concepts. However, one of the most complex tasks in this sense is maintaining an ontology up-to-date with respect to the massive amount of unstructured data that is generated and published every day. Therefore, the need arises for computational tools to build ontologies with automated or semi-automated processes. Ontologies present a structure that facilitates the integration of new entities and their relationships.

⁶ Precision or other quality measures.

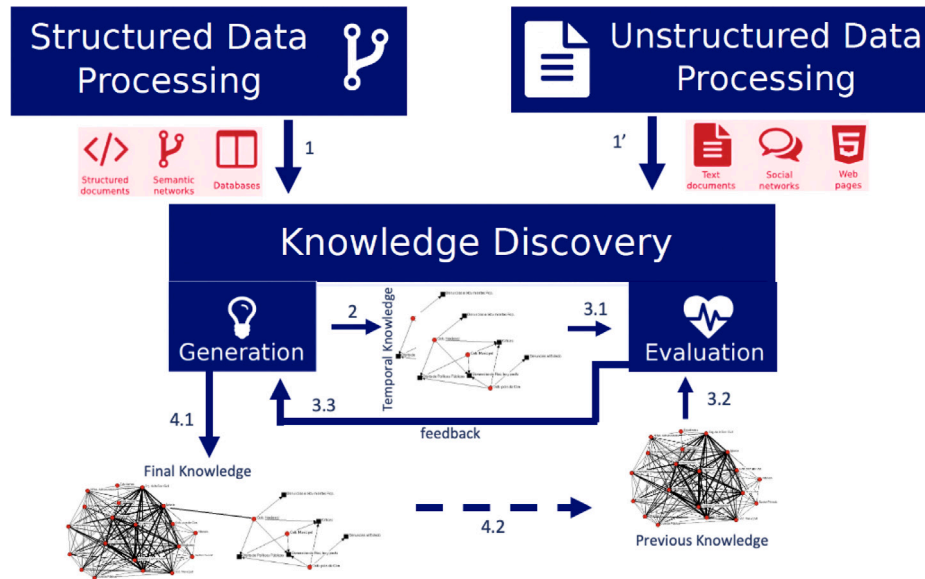


Fig. 1. Simplified architecture of the framework and workflow.

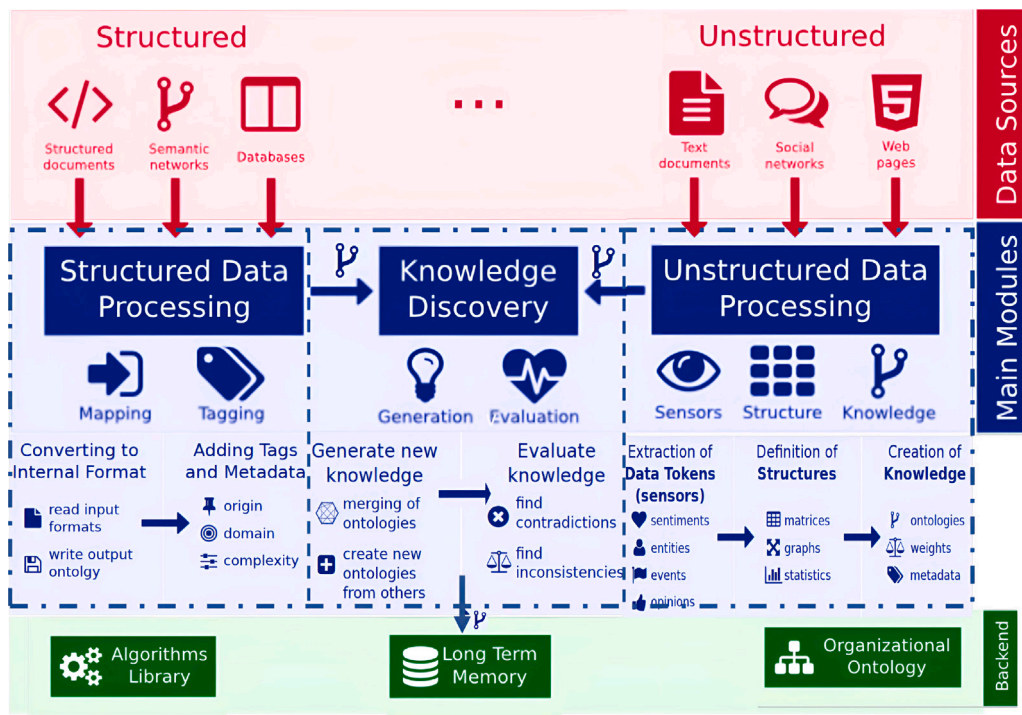


Fig. 2. Overview of the architecture of the learning knowledge. Examples of the sources of information are represented at the top. The three main modules (Sections 3.1–3.3) are highlighted at the center, with the most important processes that occur in each. The utility modules appear at the bottom.

Furthermore, they can expand from previously created knowledge bases. The use of ontologies benefits the expansion of previously generated knowledge and its adaption to specific use cases.

Our architecture has some advantages over those presented in the state of the art. For example, it does not depend on a predefined database schema. Also, it is not necessary the supervision of human experts to give feedback about the extracted knowledge. Our system treats heterogeneous documents as structured or unstructured in the same architecture, enabling specific sensors to be used depending on the problem to be tackled.

In Fig. 2, the top layer (Data Sources) illustrates the input data sources for the architecture. The middle layer consists of the primary modules responsible for processing the input data and extracting valuable knowledge embedded within. Additionally, Fig. 2 shows the sub-processes taking place within each module. These main modules maintain communication with one another by sharing ontologies, which represent the intermediate states of the acquired knowledge. This continual flow through the architecture ensures constant improvement.

The inner workings of the main modules are elaborated in the following Sections: 3.1, 3.2, and 3.3. The bottom layer, referred to

as the Backend, comprises modules that are utilized by the rest of the architecture. These include a library of common algorithms and a centralized storage service:

Algorithm Library: Different algorithms or mathematical models are contained in the Algorithm Library for solving specific problems. This library contains metadata about the algorithms which are analyzed by the architecture, which selects the right tool for a given task. For each algorithm, its performance can be measured according to the task involved, see the evaluation alternatives in Section 4. Shortly, this architecture will automatically manage, measure, and evaluate modules, algorithms, and parts, if the knowledge of how these parts interact is described in an organizational ontology inside the architecture. The technology inside it is AutoGOAL (Estévez-Velarde et al., 2020).

Long Term Memory (LTM): It serves as the repository for all the knowledge accumulated by the other modules, functioning as a storage repository within the architecture. Each stored ontology within the repository contains metadata that describes its content.

Organizational Ontology: It holds an ontological representation of the architecture, allowing it to utilize knowledge of its behavior in all internal processes. The details of this ontology can change according to specific implementations of the architecture's functionality. A convenient implementation of this ontology, in a computational solution, could be the base of a fully automated dependency injection mechanism, which would govern what modules can interact with each other.

3.1. SDPm: Structured data processing module

This module is dedicated to processing structured data, which can be found online in various formats. When it comes to representing information, there are different types of structures available, such as relational databases, concept maps, and knowledge graphs. However, we propose the use of ontologies due to their semantic richness. Ontologies have been chosen over other formats like DTO (Data Transfer Object) because they offer greater expressiveness. Given the abundance of diverse structured formats, the initial stage of this module involves converting any of these representations into a standardized internal format, specifically an ontology. This conversion is achieved through a mapping process, ensuring uniformity for internal use (Choi et al., 2006; An et al., 2006; Noy and Musen, 2003).

The problem of mapping the content of relational databases and other structured data sources into ontologies has received considerable attention since the inception of the Semantic Web (Konstantinou et al., 2008). Several semi-automatic and fully automatic approaches exist, with varying degrees of flexibility. An interesting example is the Relational.Owl framework (de Laborda and Conrad, 2006), which performs fully automatic conversion from a relational database to an Owl schema. On the other hand, other tools such as MAPONTO (An et al., 2006) require the user to provide an initial seed of correspondences before a full mapping can be produced. Using a combination of the techniques with machine learning that power these systems can provide a solution to the mapping process. The second stage of this module is to add metadata tags to the imported sources so that they can be easily integrated later with similar sources.

Formally, this stage can be defined as:

$mapping : o \rightarrow o'$ where:

$o = \langle C, R \rangle$ is an ontology (where C are the space of concept and R are relations),

$o' = \langle C', R' \rangle$ is an ontology in the internal format.

A key characteristic of structured data sources is that they almost always include a manual refinement or evaluation by domain experts. Hence, structured data sources tend to include mostly relevant and reliable information. The overall pipeline of this module can be conceptualized as a traditional Extract, Transform, and Load (ETL) process (Vassiliadis, 2009; Hermida et al., 2012), where structured data is extracted from a variety of formats, transformed into a common representation (in the form of an ontology) and loaded into the architecture for further analysis.

Afterward, the normalized and tagged block of knowledge, stored as an ontology, is passed on to the knowledge processing module for further refinement and storage purposes. To illustrate, Section 5 shows the resulting knowledge graph obtained by our architecture when fed with four datasets (three structured and one semi-structured) related to social and demographic information on educational institutions.

Formally, this stage can be defined as:

$tagging : o \rightarrow \langle o, M \rangle$ where:

$o = \langle C, R \rangle$ is an ontology,

$M = \langle (m_1, v_1), \dots, (m_k, v_k) \rangle$

is a set of metadata labels(m_i)and values(v_i).

3.2. UDPm: Unstructured data processing module

Unstructured data sources come in a wide range of formats and computer representations. While text is a prevalent form for storing and conveying human knowledge, other forms of communication, such as images, sound files, and videos, are also valuable and gaining popularity.

In contrast to structured sources, unstructured sources exhibit significant variation in terms of reliability and completeness. For example, within natural language corpora, sources such as books and encyclopedias tend to be more relevant, self-consistent, and complete than online sources such as social media, blogs, and other comments. However, the latter is broader and more dynamic than the former. A well-known case is the problem of Wikipedia reliability, where different articles may have different quality and reliability (Orlowski, 2005), and even seemingly well-written articles may temporarily contain errors or false information (Hu et al., 2007).

Furthermore, unlike structured sources, unstructured data lacks a predefined structure of concepts and relationships. Depending on the required analysis, it is possible to extract events (Exner and Nugues, 2011), entities (Gattani et al., 2013), facts (Banko et al., 2007), sentiments and opinions (Montoyo et al., 2012), actions (Balahur et al., 2011) and many other types of elements. Considering the diverse range of elements involved, the unstructured data processing module is conceived as a pipeline. This pipeline facilitates the processing and transformation of simple concepts into more complex ones. The unstructured data module at Fig. 2 shows three stages to identify relevant information.

The first stage is called the "sensory level". Within this level, there exists a series of processing units referred to as "sensors." These sensors are responsible for extracting various pieces of data. For example, one type of sensor might extract entities from a natural language corpus supported by NER (Named Entity Recognition) technologies (Gattani et al., 2013). Another sensor could extract sentiments or opinions using Sentiment Analysis tools (Montoyo et al., 2012), while another could perform POS tagging (Bird et al., 2009), recognize events (Exner and Nugues, 2011), detect actions (Giunchiglia and Fumagalli, 2017), etc. A particularly useful sensor for textual information might rely on the construction of semantic representations of text (e.g., using Abstract Meaning Representation trees (Banarescu et al., 2013)), which can then be used to identify many of the semantic modules of a given sentence, such as Subject-Verb-Object (SVO) structures (Estevez-Velarde et al., 2018).

In general, each sensor carries out a specific analysis and generates a stream of *data tokens* of a particular type. These data tokens represent individual units of semantic information, such as the presence of a specific entity or the association between an entity and an event. Importantly, these data tokens are not interrelated with each other.

This is similar to the way humans obtain independent chunks of information from the environment using the senses. At the sensory level, it is possible to add new forms of extracted information and include different sensors. This allows the architecture to be expanded to accommodate additional sensors as needs arise.

Formally, this module can be defined as a set of three functions, one for each stage:

sensorial : $T \rightarrow A$ where:

$T = \langle t_1, \dots, t_n \rangle$ is a set of documents
(e.g., in natural text),

$A = \langle a_1, \dots, a_n \rangle$ is a set of semantic annotations
(e.g. entities extracted from the text).

The second stage is called the “Structural Level”. During this stage, the extracted data tokens from the original source are collectively processed to uncover an underlying structure. For example, starting from entities extracted from a corpus of natural language, this stage would build a graph interrelating them, by connecting those entities that appear together (e.g., in the same sentence) in the original corpus. Some of the techniques that might be used in this process come from the signal analysis domain, such as Latent Semantic Analysis (LSA) (Hofmann, 2017), Principal Components Analysis (PCA) (Guo et al., 2002), Word Embeddings (Turian et al., 2010) and many clustering techniques. Typically, the output of this stage is a representation of the underlying structure of the previously extracted data tokens, which can take the form of a graph, a correlation matrix, or a statistical description. This stage can be considered as an automated rendition of the knowledge construction processes described in the METHONTOLOGY method (Fernández-López et al., 1997). To continue with the human analogy, we can consider the information extracted so far as a form of “temporal memory”. As such, it is noisy and full of spurious or irrelevant information, but it contains all the relevant pieces of knowledge within.

Formally, this stage can be defined as:

structure : $T \times A \rightarrow G$ where

$G = \langle A, E \rangle$ is a graph like-structure,

$E \subseteq \{e, e \in A \times A\}$ is a set of relations between annotations.

The third and final stage is called “Knowledge Level”. During this stage, the structured information that was previously constructed undergoes analysis to refine and eliminate noise, while extracting the pertinent pieces of knowledge. This process facilitates the synthesis of the accumulated knowledge by considering the contextual relationships between the semantic units extracted in the previous stage. Some of the types of analyses that can be performed at this stage include: extracting relevant information through clustering or noise filtering; inferring rules on the underlying structure of the data (i.e. using A-Priori); and, building predictive or generative models. Several tools for machine learning and artificial intelligence can be used for these analyses, such as Neural Networks, Bayesian Classifiers, Logistic Regression, and Hierarchical Clustering (Haykin and Network, 2004).

The output of this stage is a list of triplets, which is subsequently transferred to the Knowledge Discovery module (refer to Section 3.3) for further integration with the stored knowledge. The resulting ontology then becomes a part of the stored knowledge within the architecture. This stored knowledge is subject to iterative refinement, correction, and enhancement, as new knowledge is continually extracted from diverse sources. This iterative process is similar to “long-term memory”

in humans in the sense that only the most relevant information is remembered over time. From this perspective, the architecture simulates this process so that elements that appear consistently and possess a sufficient degree of certainty are stored. In this process, the Ontology Population and Ontology Enrichment tools are used. An example of this process can be found in Section 5.

Formally, this stage can be defined as:

knowledge : $T \times A \times G \rightarrow \langle o, M \rangle$ where

$o = \langle C, R \rangle$ is a classic ontology
(C are the classes and R are the relations),

$M = \langle (m_1, v_1), \dots, (m_k, v_k) \rangle$ is a set of
metadata labels (m_i) and values (v_i).

3.3. KDM: Knowledge Discovery module

The knowledge discovery module receives the output from both unstructured data processing and structured data processing, which are always in the form of ontologies. Each ontology represents a collection of knowledge assets either from a specific domain or a general domain. It is possible for some entities to overlap, meaning they contain the same knowledge facts, despite being labeled as different entities or relationships. There might be instances where contradictions or inconsistencies exist within individual ontologies or between different ontologies, as depicted in the knowledge discovery module in Figure Fig. 2. Since the architecture learns incrementally, a new ontology created in one of the previous modules can be checked against the relevant part of the previously stored knowledge. To accomplish this objective, the knowledge discovery module undertakes two primary tasks: generating new knowledge and evaluating its validity.

The knowledge generation process in this module involves two main processes: ontology merging (Noy et al., 2000; Noy and Musen, 2003) and the creation of new ontologies (or more general-domain ontologies) derived from existing ones (Aussenac-Gilles and Jacques, 2006; Blomqvist, 2009). Ontology merging requires the module to perform entity, relation, and instance matching between two or more ontologies that are considered similar based on specific metadata values, such as the domain (Shvaiko and Euzenat, 2013). After the merger, a new ontology is created that combines the matched entities and instances from the source ontologies into a single item. As a hypothetical example, a database about pharmaceutical products and their uses could be processed and stored. Then, as a new source of unstructured data emerged, providing a set of research papers that contain information about several diseases, these new sources would be processed and a new ontology would be obtained. This new ontology would relate to the diseases and treatments dealt with by the research papers (Ciccarese et al., 2011). By mapping the entities in this hypothetical ontology with the existing knowledge about pharmaceutical products (Chalortham et al., 2008), a richer ontology is created. Another more complex process is the creation of new knowledge (new entities or new relations) based on the inference of general rules extracted from the bulk of knowledge available.

Formally, knowledge generation can be defined as:

generation : $o' \times M \times \Theta \rightarrow \langle o^*, M \rangle$ where:

$o' = \langle C, R \rangle$ is an ontology,

$M = \langle (m_1, v_1), \dots, (m_k, v_k) \rangle$ is the metadata,

$\Theta = \langle \theta_1, \dots, \theta_m \rangle$ is the subset of ontologies
considered relevant for o' .

$o^* = \langle C^*, R^* \rangle$ is the newly created ontology.

Given the incremental nature of the knowledge that is stored, inconsistencies and contradictions are expected to exist between the different ontologies. The evaluation process eventually filters and refines the inconsistent or contradictory knowledge, see Section 4. In this

hypothetical example, two or more ontologies from a new domain may be evaluated with a medium level of reliability and stored. In time, as new ontologies of the same or similar domain are incorporated, the existing knowledge is re-evaluated, and its reliability increases or decreases accordingly. In this way, the architecture is never complete and can be self-aware of the quality of the knowledge that underpins the evolving framework.

Formally, the knowledge evaluation can be defined as:

evaluation : $\mathcal{O}' \times M \times \Theta \rightarrow M_T$ where:

$M_T = \langle m_1, \dots, m_k \rangle$ is a set of evaluation metrics.

4. Quality metrics

In this section, we present a methodology to evaluate KD SENSOMERGE and obtain metrics that can validate its performance in the wide variety of tasks that the architecture should be capable of conducting.

Software engineering metrics are suitable for evaluating the quality of software systems. These systems should be highly modular and extensible so that they can be easily adapted to new input formats, or new algorithms can be easily plugged in and integrated throughout the pipeline. Modular architecture design can help achieve a high degree of extensibility.

In addition to the previously mentioned high-level metrics, each of the tasks performed by the architecture can be evaluated separately. Most of these tasks have a defined performance metric that can be used to rate the degree of correctness of that task. For many of the tasks described in the previous sections, we can find standard performance metrics in the literature that can be used to evaluate each process.

Each of the different tasks performed by the architecture can have a very different baseline performance. A 90% precision can be a very good result in some complex tasks, such as dependency parsing (Alberti et al., 2017), but mediocre for other tasks, such as image classification (Russakovsky et al., 2015). Moreover, this baseline number can vary not only across tasks but in the same task, according to which test suite (or corpus) is used.

There are also several evaluation metrics and methodologies available for the general problem of ontology learning, such as OntoRand (Brank et al., 2006) and OntoMetric (Lozano-Tello and Gómez-Pérez, 2004).

Some of the most commonly described approaches in the literature for evaluating ontologies (Petasis et al., 2011) are: **M1 — Comparison with a gold standard**, where a learned ontology is compared with a baseline ontology for the same domain (Corcoglioniti et al. (2016)). **M2 — Expert evaluation**, where an alternative middle ground to the previous approach is having a domain expert (or several) to simply look at the resulting ontology and evaluate it according to some predefined metrics (Rospocher et al., 2016). **M3 — Evaluation through an application**, where a more practical approach consists of finding an interesting application and evaluating if the use of a learned ontology provides an improvement in the application (Gurevych et al., 2003). **M4 — Data-driven evaluation**, where a data-driven evaluation can be performed, by comparing the entities and relations in an ontology to a corpus of data that is not used during the construction of the ontology but is representative of the same domain (Brank et al. (2005)).

Evaluating a single ontology learning method is a complex task, as demonstrated by the multiple approaches proposed by the research community. Therefore, it is highly unlikely that we can find a single automated metric to measure the overall performance of our architecture. The best approach – adopted by us – is to use a combination of existing methods, adapted to the knowledge discovery scenario, with the added complexity of dealing with multiple ontologies at the same time. In some cases, a gold standard can be found and used to obtain a benchmark comparison. In other cases, provided that a suitable interface is added to easily query the knowledge, a domain expert can

interact with the architecture and give a qualitative assessment of the domain of interest. From a pragmatic perspective, the most interesting and valuable evaluation is that which finds relevant practical problems that can be solved or improved.

5. Case study: Social and demographic analysis of educational institutions

In this section, we present a case study designed as a proof-of-concept of our contributions. We aim to illustrate how the proposed architecture handles the semantic integration of heterogeneous data such as text, geospatial, and tabular. We choose a domain that corresponds to a requirement of the project “Study of the technological needs to generate a Geo Artificial Intelligence system in public administration”,⁷ developed by the Center of Digital Intelligence,⁸ Alicante, highlighting the potential of the KD SENSOMERGE architecture to leverage publicly available data. We analyzed the socio-demographic context of educational institutions (EI) by merging data from the different sources described in Table 1.

However, it is worth noting that the quality of the integrated knowledge will depend on factors such as the type of data, its availability, and the performance of the sensors. Also, other domains may be studied as long as they allow for demonstrating the integration of heterogeneous sources.

We consider educational institutions officially recognized in Spain and, for the sake of simplicity, only up to the high school level. An EI has attributes such as name, latitude, longitude, type (whether public, private, or mixed), phone number, etc. We also gathered data from the census sections (CS), which are the smaller territorial administrative units for statistical data collection and management within a municipality in Spain. The CS has a delimited cartography that allows operations such as the location verification of a place (e.g. an educational institution). Also, the CS has attributes linking statistical data such as age distribution, average income, etc. In addition, user reviews about EIs are considered an unstructured source.

The next section describes the realization of the main elements of the architecture to prepare and process the data 5.1 as well as the knowledge generation stage 5.2.

5.1. Case study: Data processing

As shown in Table 1, we considered several datasets. Here, we discuss their details as well as the core responsibilities and results of the data preprocessing components of the architecture.

From the *cs_geo* dataset we obtained the following fields:

- the CS official identifier (*cusec*). A 10-character string encoding the municipality, district, and section of the CS.
- the cartography *geometry* of the CS. A polygon is given by a set of geographic coordinates.

The other datasets related to CS are *cs_income*, and *cs_population*. Average income by CS, as well as total population and inhabitants by place of birth. The relevant fields are:

- section identifier (*section*). This dataset uses a different encoding for the CS identifier since it comes as a string together with the municipality name. We need to preprocess this field to remove the unnecessary data to enable the mapping to the *cusec* field from *cs_geo*.
- the average income (*average_income*) of people living within a CS.

⁷ <https://cenid.es/proyectosgeointeligencia-artificial-podcast/>

⁸ <https://cenid.es>

Table 1

Data sources used for the case study. Within Section 5 we will use the name as a fancy identifier to easily refer to the dataset.

Description	Type	Source
Census Section Cartography. Specifies the shapes delimiting the CS. Includes the province, sector area, and other location data but no statistical information about the CS. name: <i>cs_geo</i>	structured	Spanish Government Open Data Portal ^a
Census demographic and economic data. Provides different indices for the CS such as total population, or disaggregated by gender, place of birth, and average income among others. It does not contain any cartographic data. name: <i>cs_income</i> , <i>cs_population</i>	structured	Spanish Statistics Institute ^b
Dataset of educational institutions. Includes the name, location, whether it is public, private, or mixed, as well as other identification data. It does not contain data about the CS in which the EI is located or opinions about the EI. name: <i>ei_data</i>	structured	Valencian Community Geospatial Data Catalog ^c
Google Maps Reviews. Google user opinions about educational institutions. This dataset has a structure in the sense that it contains well-defined fields such as EI identifier, user identifier, as well review text. However, the later field is plain text, thus unstructured. name: <i>ei_reviews</i>	semi-structured	Google Maps ^d

^a <https://datos.gob.es/es/catalogo/ea0010587-cartografia-digitalizada-de-secciones-censales>

^b <https://www.ine.es/jaxiT3/Tabla.htm?t=31097>

^c https://geocatalog.gva.es/#/search?uuid=spa_icv_educacion_centros_educativos

^d <https://www.google.com/maps>

- the total population *population_total* of the CS, as well the number of people born in Spain *population_spain* and abroad *population_abroad*.

The source *ei_data* includes detailed information for an EI. We retain the following data:

- the official name (*name*) of the EI. A string of arbitrary length.
- the levels of education (*level*) covered by the EI program, e.g. secondary, high school, etc.
- the (*regime*) of the EI. A nominal field indicating whether the EI is public, private, or hybrid.
- geographic coordinates *latitude* and *longitude*

Finally, the *ei_reviews* dataset collects user reviews from Google Maps related to the EI. It includes:

- the (*name*) of the EI. It depends on the API used but it is possible to search reviews using the *name* from *ei_data* and the target region, e.g. Alicante, Valencian Community, Spain. Otherwise, we may need to map this field to *name* in EI to improve the coverage of the discovered knowledge.
- the screen name (*author*) of the user who wrote the review.
- the (*score*), a numeric value from 1 to 5 representing the score given by the *author* of the review.
- the content (*text*) of the review.

The next section discusses how the knowledge discovery loop proposed by KD SENSO-MERGE is carried out.

5.2. Case study: Knowledge discovery

This section describes the process of Knowledge Discovery from the datasets in Section 5.1 and the sensors used.

5.2.1. Sensors

For the case study, we used two different sensors, one designed to extract information from plain text and the other from structured geographical data. In this section, we cover their main aspects.

The dataset *ei_reviews* contains reviews as well as the score the reviewer assigned to the EI being analyzed. With the appropriate sensors, KD SENSO-MERGE can provide very valuable information not explicitly given in the dataset. In our case, we leverage a *topic classification sensor* to find out if the review is talking about one or more of the following topics within the domain:

- *staff*: captures opinions related to the professors, management, service staff as well other relevant stakeholders within the EI.
- *education*: the comment is about the values, goals, curricula, educational methodology of the EI, etc.
- *infrastructure*: refers to the conditions of the building, sports areas, and gardens of the EI among others.
- *environment*: the review talks about the surrounding neighborhood of the EI, parking facilities, parks, people, etc.

We used a text categorizer⁹ designed for the NLP library spaCy.¹⁰ We tuned the model using 10 to 20 examples of sentences encompassing the topics, extracted from *ei_reviews*. When analyzing the text, the model assigns a score from 0 to 1 representing the likelihood of the topic being covered by the text. If the *score* ≥ 0.3 we considered the topic is included.

The other sensor handles location data, i.e. latitude and longitude, as well as cartography data. Commonly, certain entities such as the CS, neighborhoods, or countries have associated information about their geographical limits or perimeters. Other EIs have their location as an attribute. The *located in sensor* can detect this kind of information, creating a new relation between entities, in our case, indicating the CS

⁹ <https://spacy.io/universe/project/classyclassification>

¹⁰ <https://spacy.io>

section	geometry:attr
3014.01001	POLYGON[(-0.47830111036416273, 38.34417517039852), (-0.4783027334266349, 38.34417158767038), (-0.4783171416358447, 38.34414525411608)]
3014.01008	POLYGON[(-0.49036680459935167, 38.34615206413387), (-0.4903536056030057, 38.34611764483372), (-0.4903398128258657, 38.346082893849605)]
3014.01017	POLYGON[(-0.484327156033468, 38.34639956829743), (-0.48424292828128623, 38.346106273605564), (-0.4841660275010249, 38.34597489534194)]

Fig. 3. Snippet of the *cs_geo* dataset in tabular format.

section	average_income:att	population_total:att	population_spain:att	population_abroad:att
3014.01001	17759	1146	915	231
3014.01008	21884	774	709	65
3014.01017	18767	721	552	169

Fig. 4. Snippet of the dataset containing income and population in tabular format.

where the EI is located. The dataset *cs_geo* is provided in shapefile.¹¹ We use GeoPandas¹² to handle this data. This library has functions to manage the perimeter of the CS as well as to verify whether an EI location is within it. However, other libraries/approaches can be used to implement this sensor.

5.2.2. Steps involved in applying the process

We started loading the four mentioned datasets into the architecture whereby all the data were transferred to the Un/Structured Data Processing modules (UDPm and SDPm, respectively) depending on the type of data involved. The architecture can infer the data types automatically or from the user's column labeling, notifying the respective modules for appropriate data management before transferring the temporal memory (i.e. triplets extracted) to the Knowledge Discovery module (KDM). Noticeably, when a field of the dataset contains string data, NLP sensors can be activated by the Unstructured Data Processing module to detect entities, events, or any category to be discovered or matched with semantic pieces already existing in the Long Term Memory (LTM). In the case of a field containing categories, it is labeled as *:rel* to treat it as concepts to be matched.

5.2.3. Loading dataset one: *cs_geo*

First, we loaded *cs_geo* dataset generating temporary triplet-based structures. For this dataset, the triplets had the following format: subject–predicate–object. Here are some examples: *cusec(value)-municipality-municipality(value)*; *cusec(value)-district-district(value)*; and so on. For this dataset, the triplet is focused on linking the first column (*cusec*) as an identifier of entities to the rest of the attributes. Fig. 3 shows a snippet of this data in tabular format.

Since all the data is structured, the SDPm takes care of this first step. This process to generate triplets is applied to both modules UDPm and SDPm.

Next, the output of the SDPm, a list of triplets, is transferred to the KDM to be stored (considering a merging process if necessary). Due to no knowledge having been recorded in advance, no conflicts or merging issues were handled by the KDM to store these triplets in the LTM (i.e. neo4j for example).

At this point, we can query our knowledge base (KB) only to get information about the CS, i.e. their *cusec* or *geometry*. Fig. 7(a) shows a view of the KB after this step and contains only instances of the CS entity.

5.2.4. Loading dataset two: *cs_income* and *cs_population*

We proceeded to load, and depending on the type of data, UDPm or SDPm is activated. We tagged all fields, but the *section* as attributes

to avoid the discovery of other entities. Fig. 4 shows an example of the data.

In this case, the architecture, through the KDM, recognizes information that refers to the same CS, producing the merging process that updates the graph. This stage does not introduce changes to the structure of the KB, but CS instances have additional attributes now.

5.2.5. Loading dataset three: *ei_data*

SDPm deals also with the structured data of the *ei_data* dataset, an excerpt of the data is shown in Fig. 5. However, when the KDM receives the list of triplets generated from the *ei_data*, it activates the *located in sensor*, which is capable of finding the area, i.e. CS, to which another entity belongs. This enables the detection of location-related fields as well as areas given by fields like *geometry*. Next, the appropriate relations between CS and EI are created. Other fields from *ei_data* are included as attributes for EI entities.

Up to this point, the module used for this process is the SDPm since no unstructured data is involved. The resulting list of triplets that encodes the pieces of knowledge is reviewed by checking for consistencies or missing data, and there were none in our case study. Once the checks are done, the list of triplets is stored in the LTM via the KDM. Also, this module is in charge of the merging process, i.e. updating attributes and relations of existing entities. New entities are stored and linked to existing ones. Fig. 7(b) illustrates a view of the KB after loading the data. As it can be observed, it contains new instances of the EI entity as well the relations to the EI introduced by the *located in sensor*. See Fig. 7(a) for a previous state of the KB.

5.2.6. Loading dataset four: *ei_reviews*

Finally, *ei_reviews* were fed, and the fields *name*, *author*, and *score* were processed by the SDPm while the field *text* was handled by the UDPm. Then, the KDM maps the field *name* to its match *ei_data* and is leveraged by the architecture for linking the *reviews*, seen as an attribute, to the corresponding EI. Fig. 6 shows a fragment of this data in tabular format. It is worth noting that the dataset did not include the topics covered within the topic.

The field *text* is analyzed at the sensory level. As a proof of concept, we activated the *topic classification sensor* presented in Section 5.2.1. This is useful to get a deeper understanding of not only how people scored the EI but also the most relevant topics for them. As a result, a list of triplets provided by both modules is merged into the LTM, ensuring the coexistence and improvement of the information at each loading round. Fig. 7(c) illustrated the structure of the KB after this step. Instances of the CS, EI, and reviews are presented as well as the relations between them.

¹¹ <https://www.esri.com/content/dam/esrisites/sitecore-archive/Files/Pdfs/library/whitepapers/pdfs/shapefile.pdf>

¹² <https://geopandas.org/en/stable/>

institution	level:attr	regime:att	longitudo:att	latitudo:att	section:re
IES JAIME II	INSTITUTO DE EDUCACIÓN SECUNDA	Público	-0.470363	38.362228	03014.03037
CEIP EUSEBIO SEMPERE	COLEGIO DE EDUCACIÓN INFANTIL Y	Público	-0.503569	38.370572	03014.06020
CEIP MANJÓN-CERVANTES	COLEGIO DE EDUCACIÓN INFANTIL Y	Público	-0.482196	38.35745	03014.03006

Fig. 5. Snippet of the ie_data dataset in tabular format. Content in Spanish Language.

comment	institution:rel	author:attr	rating:att	text:attr
review:00025	CEIP MANJÓN-CERVANTES	anónimo	5	Buenos profes y buenas instalaciones. Mis hijos van muy contentos al colegio.
review:00478	CEIP SAN GABRIEL	anónimo	5	Muy buena escuela. Con una línea educativa compleja y firme.
review:00002	IES JAIME II	anónimo	5	Centro educativo muy completo con profesores muy profesionales.

Fig. 6. Snippet of the ie_reviews dataset in tabular format. Content in Spanish Language.




Fig. 7. View of the structure of the Knowledge Base after each dataset is loaded. The yellow nodes represent CS instances, the green nodes EI instances, the gray ones the reviews, and the cyan the topics.

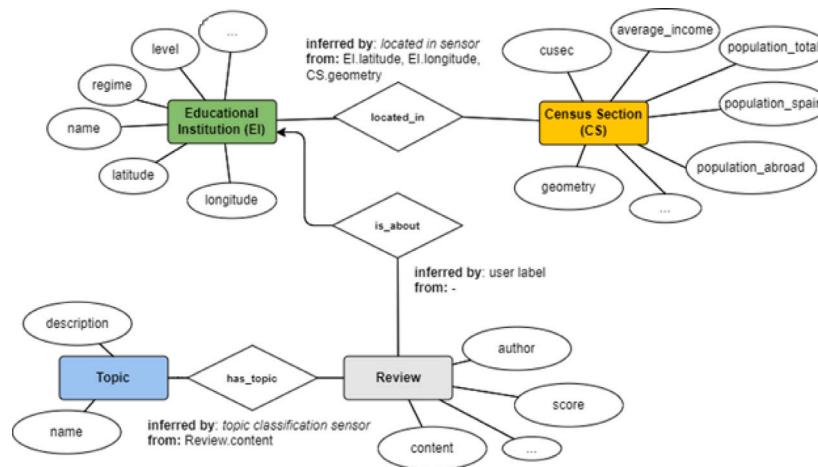


Fig. 8. Concepts and relations discovered by KD SENSO-MERGE, Near relations, a note about how they were inferred.

Following these steps, KD SENSO-MERGE managed to build a KB populated with instances from the entities CS, EI, Reviews, and Topic, as well as the semantic relations between them. Fig. 8 shows the concepts and relations discovered.

5.3. Case study: Knowledge exploitation

After KD SENSO-MERGE has ingested the datasets, we obtained the distribution shown in Table 2 which collects the counts of instances in each dataset and identifies the number of discovered relations.

Instances are the basic piece of data, for example, in ei_data it comprises a record with all that is registered for an EI. Relations is the number of relations discovered from the dataset, and KB Size is the number of relations in the KB.

Dataset	Instances	KB nodes	KB relations
cs_geo	243	243	0
cs_stats	243	243	0
ei_data	133	376	133
ei_review	1811	1948	4213
TOTAL	2430	2810	4346

Automatically merging data – geospatial, entity, and topic linking – and discovering knowledge enables questions to be answered that would not be possible through individual sources. For example, we would want to know “the CS with average income below 10.000 euros, with at least one public EI that has reviews with a score below 3 and

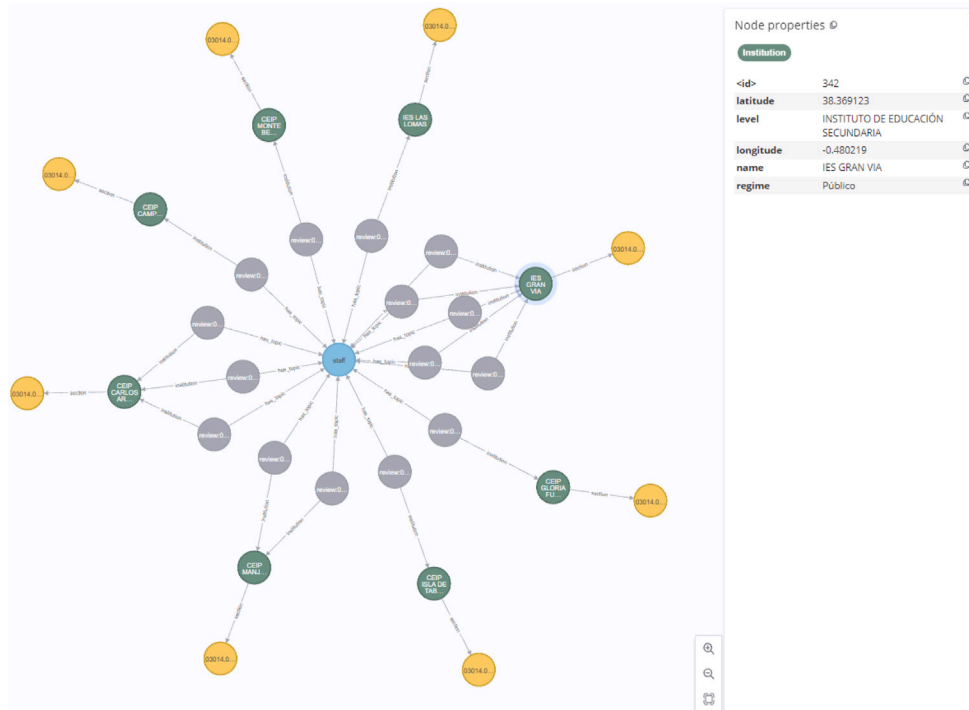


Fig. 9. Subset of the KB showing the CS with average income below 10.000 euros, with at least one public EI that has reviews with a score below 3 and that talks about the staff. Results from the integration of *cs_geo*, *cs_stats*, *ei_data* and *ei_reviews*.

that talks about the staff ”. Fig. 9 shows a subset of the KB answering this query.

Using the geographic data, we may also analyze visually the localization of the EI together with the average income. Fig. 10 shows this information for the city of Alicante, in Spain. This is a view of a Geographic Information System (GIS) that was developed leveraging the knowledge from KD SENSO-MERGE which saved us from the tiresome work of merging the data, modeling the entities, and creating the relations.

Some approaches have dealt with challenges that are similar to our case study. For example (Al-Sulttani et al., 2021) faces superficial area challenges by applying machine learning for surface water quality prediction; and (Afan et al., 2021) deals with the fluctuations of groundwater level which are modeled by employing ensemble deep learning techniques. The two cited researches have in common with our proposal that they handle heterogeneous data (i.e. numbers and geolocations) to make predictions. However, for both of them, geolocation is only descriptive data, and therefore not integrated into the predictive model. By contrast, in our use case, machine learning (i.e. soft or deep) is applied to enrich the data for further semantic and statistical analysis, and all data is integrated into the KB.

5.4. Case study: Evaluation

The dataset was transformed into triplets to represent entities, attributes, and relations with other entities. Many of these entities were automatically recognized and generated as new concepts; others were linked to each other via geolocalization or topic. All of these procedures were evaluated by the architecture to ensure quality knowledge would be stored in the Long Term Memory (LTM).

As was mentioned in Section 4 we have different evaluation metrics. The user should pick the most suitable metric in each case. Here, we analyze the cases of M1 — Comparison with a gold standard and M2 — Expert Evaluation.

Table 3
Results of M1 evaluation respect to the relation *has_topic*.

Topic	Precision	Recall
staff	0.68	0.69
education	0.22	0.24
infrastructure	0.21	0.20
environment	0.02	0.02

In our case, relations produced by sensors are the source of errors. For other relations, the map can be perfectly inferred given the datasets share common features that are used for linking between them. For the *located_in* relation, we can determine with precision where each EI is located. Thus, we can use M2 to check if the relations were inferred as expected. In this case, 100% of the expected 133 relations were created correctly.

The *has_topic* relation is created by the *topic classification sensor*. We used M1, however, a gold standard needed to be prepared beforehand. We randomly picked 100 reviews and annotated the topics. Table 3 shows the specific quality metrics, as part of M1 (See Section 4) applied to this Case Study. It is worth noting that modeling a state-of-the-art topic classifier is beyond the scope of this work whose main purpose is to illustrate the role of the sensory level. As expected, the results are discrete in terms of precision and recall, but suitable to illustrate the procedure users can follow to evaluate their deployment of KD SENSO-MERGE. Work on optimizing the sensor, choosing another approach, or optimizing hyper-parameters may lead to notably better results.

Recall and precision can be translated into an estimation of the quality. Fig. 11 shows different effects of these metrics on the extracted knowledge. The recall accounts for the capability of the system to identify the entities and relations in the content. Low recall could mean we will miss some entities or relations, such as M_1 in 11. In turn, low precision implies the introduction of spurious elements, such as W_1 .

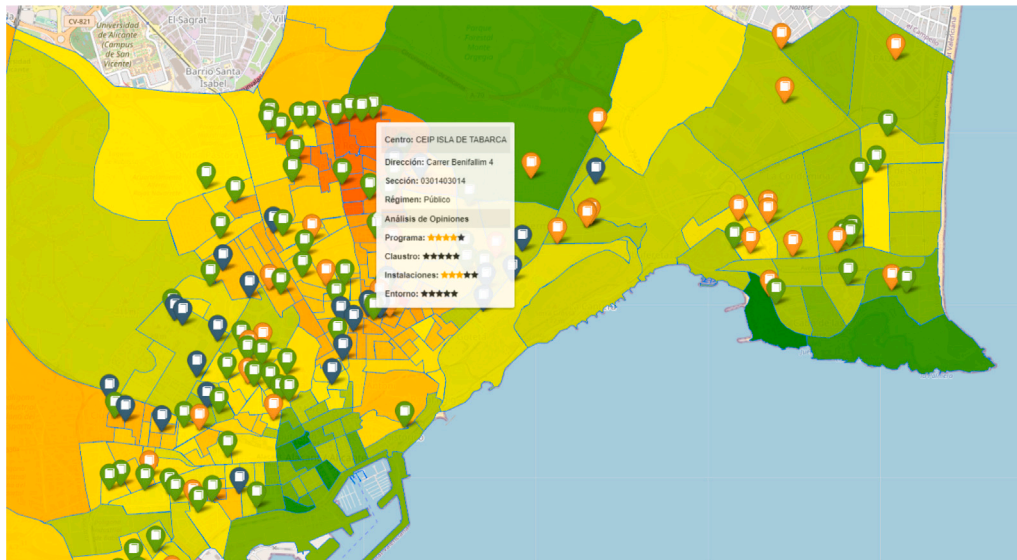


Fig. 10. Map showing the CS where the EIs are located. Public EI is represented with the green marker, blue for private, and blue for hybrid. For the CS background, green represents higher incomes. Labels are in Spanish since we use a dataset in this language.




Fig. 11. Different cases for nodes and relations. C_i correct nodes, M_1 missing entity, W_1 and spurious entity discovered by a sensor. The black arrow represents a correct relation, gray dashed missing ones, and, red dotted incorrect relations.

6. Discussion and conclusions

This research resulted in the design of KD SENSOMERGE, a knowledge discovery modular architecture that semantically integrates heterogeneous data, both structured and unstructured. The architecture builds internal representations that are adaptable across multiple domains. As the modules perform specific tasks, they communicate with each other and continuously merge different knowledge fragments.

We designed a case study that shows the capabilities of the KD SENSOMERGE. It represents a particular knowledge discovery scenario for which the architecture provided a solution. The case study illustrates each of the main modules of the architecture, as well as their interconnections.

This work identifies three challenges resolved with promising results. Regarding **Data Integrity Challenge**, KD SENSOMERGE is capable of dealing with conflicting and contradictory facts when performing inference. Unlike most ontology learning systems, which assume that the final ontology obtained is correct, our architecture explicitly models a degree of reliability for each extracted knowledge fragment, using different evaluation metrics. The architecture was designed from the outset to deal with uncertain or noisy information. Over time, older knowledge is reevaluated, reinforced, or discarded, as new

facts are discovered. This makes the architecture dynamic, in the sense that stored knowledge is always evolving, mirroring actual knowledge in human brains. In the future, we plan to manage the ontology data as digital entities, to create profiles and model their behavior over time.

The **Knowledge Extraction Challenge** has been addressed through the evaluation of the case study, which corroborated that sensor performance is fundamental for the quality of integrated knowledge. This is of particular importance for NLP sensors or machine learning-based sensors in general. The values obtained from our case study – about 0.2 recall and precision for some topics – may render a sensor useless for most applications. Nevertheless, we did not aim to develop a state-of-the-art topic detector as the main goal was to demonstrate the sensor’s role within the architecture. However, in a real-life scenario, users need to pay attention to metrics such as recall and precision since they are indicative of missing or spurious knowledge. To tackle this, we plan in future to implement AutoML (Estévez-Velarde et al., 2020) strategies to automatically manage, measure, and evaluate modules, and algorithms.

Finally, the **Data Scalability challenge** has been dealt with in the case study. We can conclude that developing new knowledge extraction sensors is mandatory for opening up new possibilities to create new knowledge. This enhances the deployment of KD SENSOMERGE to tackle the challenges of the knowledge society. This is a worthwhile direction for future research as there is a wide spectrum of knowledge that can be extracted from unstructured sources such as named entities, relations, or topics. For example, a Name Entity Recognition sensor would enable the identification of person entities such as staff members, events, or other educational institutions. Also, improving field matching sensors is important since they guarantee that knowledge related to the same concept is correctly connected.

Another issue to be addressed in future research is to better profile the computational requirements of our proposal. As an architecture, it is a high-level specification of components, their interfaces, and interactions (Pressman, 2005). As a consequence, overall computational requirements will depend on the implementation of the specific modules. Nonetheless, it is important to consider this issue since some approaches may introduce a prohibitive computation burden into the system.

CRedit authorship contribution statement

Yoañ Gutiérrez: Conceptualization, Formal Analysis, Writing – original draft. **José I. Abreu Salas:** Data Curation, Validation, Writing – original draft. **Andrés Montoyo:** Methodology, Investigation, Writing – review & editing. **Rafael Muñoz:** Conceptualization, Validation, Writing – review & editing. **Suilan Estévez-Velarde:** Software, Visualization, Writing – Original draft.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data is publicly available. The paper describes where the reader can obtain the data.

Acknowledgments

This research is supported by the University of Alicante, Spain, the Spanish Ministry of Science and Innovation, the Generalitat Valenciana, Spain, and the European Regional Development Fund (ERDF) through the following funding: At the national level, the following projects were granted: TRIVIAL (PID2021-122263OB-C22); and CORTEX (PID2021-123956OB-I00), funded by MCIN/AEI/10.13039/501100011033 and, as appropriate, by “ERDF A way of making Europe”, by the “European Union” or by the “European Union NextGenerationEU/PRTR”. At regional level, the Generalitat Valenciana (Conselleria d’Educacio, Investigacio, Cultura i Esport), Spain, granted funding for NL4DISMIS (CIPROM/2021/21). Moreover, it was backed by the work of two COST Actions: CA19134 - “Distributed Knowledge Graphs” and CA19142 - “Leading Platform for European Citizens, Industries, Academia, and Policymakers in Media Accessibility”.

References

- Adams, T., Dullea, J., Clark, P., Sripada, S., Barrett, T., 2000. Semantic integration of heterogeneous information sources using a knowledge-based system. In: Proc 5th Int Conf on CS and Informatics. CS&I’2000.
- Afan, H.A., Ibrahem Ahmed Osman, A., Essam, Y., Ahmed, A.N., Huang, Y.F., Kisi, O., Sherif, M., Sefelnasr, A., Chau, K.-w., El-Shafie, A., 2021. Modeling the fluctuations of groundwater level by employing ensemble deep learning techniques. *Eng. Appl. Comput. Fluid Mech.* 15 (1), 1420–1439.
- Aggarwal, C.C., Wang, H., et al., 2010. *Managing and Mining Graph Data*. Vol. 40. Springer.
- Al-Sulttani, A.O., Al-Mukhtar, M., Roomi, A.B., Farooque, A.A., Khedher, K.M., Yaseen, Z.M., 2021. Proposition of new ensemble data-intelligence models for surface water quality prediction. *IEEE Access* 9, 108527–108541.
- Alani, H., Kim, S., Millard, D.E., Weal, M.J., Hall, W., Lewis, P.H., Shadbolt, N.R., 2003. Automatic ontology-based knowledge extraction from web documents. *IEEE Intell. Syst.* 18 (1), 14–21.
- Alberti, C., Andor, D., Bogatyry, I., Collins, M., Gillick, D., Kong, L., Koo, T., Ma, J., Omernick, M., Petrov, S., Thanapirom, C., Tung, Z., Weiss, D., 2017. SyntaxNet models for the CoNLL 2017 shared task. *CoRR abs/1703.04929*.
- An, Y., Borgida, A., Mylopoulos, J., 2006. Building semantic mappings from databases to ontologies. In: 21st National Conference on Artificial Intelligence. AAAI 06.
- Aussenac-Gilles, N., Jacques, M.-P., 2006. Designing and evaluating patterns for ontology enrichment from texts. In: International Conference on Knowledge Engineering and Knowledge Management. Springer, pp. 158–165.
- Balahur, A., Hermida, J.M., Montoyo, A., Muñoz, R., 2011. EmotiNet: A Knowledge Base for Emotion Detection in Text Built on the Appraisal Theories. In: *Natural Language Processing and Information Systems: 16th International Conference on Applications of Natural Language to Information Systems, NLDB 2011, Alicante, Spain, June 28-30, 2011. Proceedings*. Springer Berlin Heidelberg, pp. 27–39.
- Banarescu, L., Bonial, C., Cai, S., Georgescu, M., Griffitt, K., Hermjakob, U., Knight, K., Koehn, P., Palmer, M., Schneider, N., 2013. Abstract meaning representation for sembanking. In: Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse. pp. 178–186.
- Banko, M., Cafarella, M.J., Soderland, S., Broadhead, M., Etzioni, O., 2007. Open information extraction from the web. In: *IJCAI*. Vol. 7. pp. 2670–2676.
- Barker, K., Agashe, B., Chaw, S.Y., Fan, J., Friedland, N., Glass, M., Hobbs, J., Hovy, E., Israel, D., Kim, D.S., et al., 2007. Learning by reading: A prototype system, performance baseline and lessons learned. In: *AAAI*. Vol. 7. pp. 280–286.
- Bergamaschi, S., Castano, S., Vincini, M., Beneventano, D., 2001. Semantic integration of heterogeneous information sources. *Data Knowl. Eng.* 36 (3), 215–249, Heterogeneous Information Resources Need Semantic Access.
- Bingham, E., Mannila, H., 2001. Random projection in dimensionality reduction: Applications to image and text data. In: *Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, pp. 245–250.
- Bird, S., Klein, E., Loper, E., 2009. *Natural Language Processing with Python: Analyzing Text with the Natural Language Toolkit*. " O'Reilly Media, Inc."
- Blomqvist, E., 2009. Ontocase-automatic ontology enrichment based on ontology design patterns. In: *International Semantic Web Conference*. Springer, pp. 65–80.
- Bolukbasi, T., Chang, K.-W., Zou, J.Y., Saligrama, V., Kalai, A.T., 2016. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. In: *Advances in Neural Information Processing Systems*, pp. 4349–4357.
- Brachman, R.J., Levesque, H.J., Reiter, R., 1992. *Knowledge Representation*. MIT Press.
- Brank, J., Grobelnik, M., Mladenić, D., 2005. A survey of ontology evaluation techniques.
- Brank, J., Mladenic, D., Grobelnik, M., 2006. Gold standard based ontology evaluation using instance assignment. In: *Workshop on Evaluation of Ontologies for the Web, EON*. Edinburgh, UK.
- Brewster, C., Ciravegna, F., Wilks, Y., 2002. User-centred ontology learning for knowledge management. *Nat. Lang. Proc. Inform. Syst.* 203–207.
- Brin, S., Page, L., 1998. The anatomy of a large-scale hypertextual web search engine. *Comput. Netw. ISDN Syst.* 30 (1–7), 107–117.
- Buitelaar, P., Cimiano, P., Racioppa, S., Siegel, M., 2006. Ontology-based information extraction with SOBA. In: *Proceedings of the International Conference on Language Resources and Evaluation*. LREC.
- Buitelaar, P., Sintek, M., 2004. Ontolt version 1.0: Middleware for ontology extraction from text. In: *Proc. of the Demo Session At the International Semantic Web Conference*.
- Buron, M., Goasdoué, F., Manolescu, I., Mugnier, M.-L., 2020. Obi-Wan: Ontology-based RDF integration of heterogeneous data. *Proc. VLDB Endow.* 13 (12), 2933–2936.
- Castano, S., Espinosa, S., Ferrara, A., Karkaletsis, V., Kaya, A., Melzer, S., Möller, R., Montanelli, S., Petasis, G., 2007. Ontology dynamics with multimedia information: The Boemie evolution methodology. In: *International Workshop on Ontology Dynamics*. IWOD-07, p. 41.
- Chalortham, N., Leesawat, P., Buranarach, M., Supnithi, T., 2008. Ontology development for pharmaceutical tablet production expert system. In: *Electrical Engineering/Electronics, Computer, Telecommunications and Information Technology, 2008. ECTI-CON 2008. 5th International Conference on*. Vol. 1. IEEE, pp. 205–208.
- Chandrasekaran, B., 1986. Generic tasks in knowledge-based reasoning: High-level building blocks for expert system design. *IEEE Exp.* 1 (3), 23–30.
- Chein, M., Mugnier, M.-L., 2008. *Graph-Based Knowledge Representation: Computational Foundations of Conceptual Graphs*. Springer Science & Business Media.
- Chen, H., Chiang, R.H., Storey, V.C., 2012. Business intelligence and analytics: From big data to big impact. *MIS Q.* 1165–1188.
- Choi, N., Song, I.-Y., Han, H., 2006. A survey on ontology mapping. *ACM Sigmod Rec.* 35 (3), 34–41.
- Cicarese, P., Ocana, M., Garcia Castro, L.J., Das, S., Clark, T., 2011. An open annotation ontology for science on web 3.0. *J. Biomed. Semant.* 2 (2), S4.
- Cimiano, P., Mädche, A., Staab, S., Völker, J., 2009. Ontology learning. In: *Handbook on Ontologies*. Springer, pp. 245–267.
- Cimiano, P., Völker, J., 2005. text2onto. In: *International Conference on Application of Natural Language to Information Systems*. Springer, pp. 227–238.
- Corcho, O., Fernández-López, M., Gómez-Pérez, A., 2003. Methodologies, tools and languages for building ontologies. Where is their meeting point? *Data Knowl. Eng.* 46 (1), 41–64.
- Corcoglioniti, F., Rospocher, M., Aprosio, A.P., 2016. Frame-based ontology population with pikes. *IEEE Trans. Knowl. Data Eng.* 28 (12), 3261–3275.
- Craven, M., DiPasquo, D., Freitag, D., McCallum, A., Mitchell, T., Nigam, K., Slattery, S., 2000. Learning to construct knowledge bases from the world wide web. *Artif. Intell.* 118 (1–2), 69–113.
- Davidson, J., Liebal, B., Liu, J., Nandy, P., Van Vleet, T., Gargi, U., Gupta, S., He, Y., Lambert, M., Livingston, B., et al., 2010. The YouTube video recommendation system. In: *Proceedings of the Fourth ACM Conference on Recommender Systems*. ACM, pp. 293–296.
- de Laborda, C.P., Conrad, S., 2006. Bringing relational data into the semantic web using SPARQL and relational.OWL. In: *Proceedings of the International Conference on Data Engineering Workshops*. ICDEW 2006.
- Drymonas, E., Zervanou, K., Petrakis, E.G., 2010. Unsupervised ontology acquisition from plain texts: The OntoGain system. In: *International Conference on Application of Natural Language to Information Systems*. Springer, pp. 277–287.

- Estévez-Velarde, S., Gutiérrez, Y., Almeida-Cruz, Y., Montoyo, A., 2020. General-purpose hierarchical optimisation of machine learning pipelines with grammatical evolution. *Inform. Sci.*
- Estevez-Velarde, S., Gutierrez, Y., Montoyo, A., Piad-Morffis, A., Munoz, R., Almeida-Cruz, Y., 2018. Gathering object interactions as semantic knowledge. In: *Proceedings on the International Conference on Artificial Intelligence*. Icai, The Steering Committee of The World Congress in Computer Science, Computer ..., pp. 363–369.
- Etzioni, O., Cafarella, M., Downey, D., Kok, S., Popescu, A.-M., Shaked, T., Soderland, S., Weld, D.S., Yates, A., 2004. Web-scale information extraction in knowitall:(Preliminary results). In: *Proceedings of the 13th International Conference on World Wide Web*. ACM, pp. 100–110.
- Exner, P., Nugues, P., 2011. Using semantic role labeling to extract events from wikipedia. In: *Proceedings of the Workshop on Detection, Representation, and Exploitation of Events in the Semantic Web (DeRiVE 2011)*. Workshop in Conjunction with the 10th International Semantic Web Conference. pp. 23–24.
- Faure, D., Poibeau, T., 2000. First experiments of using semantic knowledge learned by ASIUM for information extraction task using INTEX. In: *Proceedings of the ECAI Workshop on Ontology Learning*.
- Fellbaum, C., 1998. *WordNet: An Electronic Lexical Database*. MIT Press.
- Fernández, J., Gutiérrez, Y., Gómez, J.M., Martínez-Barco, P., 2015. Social rankings: Análisis visual de sentimientos en redes sociales. *Procesamiento del Lenguaje Nat.* 55, 199–202.
- Fernández-López, M., Gómez-Pérez, A., Juristo, N., 1997. Methontology: From Ontological Art Towards Ontological Engineering. *American Association for Artificial Intelligence*.
- Ferrucci, D., Levas, A., Bagchi, S., Gondek, D., Mueller, E.T., 2013. Watson: Beyond Jeopardy!. *Artificial Intelligence* 199, 93–105.
- Fusco, G., Aversano, L., 2020. An approach for semantic integration of heterogeneous data sources. *PeerJ Comput. Sci.* 6, e254.
- Gattani, A., Lamba, D.S., Garera, N., Tiwari, M., Chai, X., Das, S., Subramaniam, S., Rajaraman, A., Harinarayan, V., Doan, A., 2013. Entity extraction, linking, classification, and tagging for social media: A Wikipedia-based approach. *Proc. VLDB Endow.* 6 (11), 1126–1137.
- Giunchiglia, F., Fumagalli, M., 2017. Teleologies: Objects, actions and functions. In: *International Conference on Conceptual Modeling*. Springer, pp. 520–534.
- Guarino, N., 1995. Formal ontology, conceptual analysis and knowledge representation. *Int. J. Hum.-Comput. Stud.* 43 (5–6), 625–640.
- Guo, Q., Wu, W., Massart, D., Boucon, C., De Jong, S., 2002. Feature selection in principal component analysis of analytical data. *Chemometr. Intell. Lab. Syst.* 61 (1–2), 123–132.
- Gurevych, I., Malaka, R., Porzel, R., Zorn, H.-P., 2003. Semantic coherence scoring using an ontology. In: *Proceedings of the 2003 Conference of the North American Chapter of the Association for Computational Linguistics on Human Language Technology-Volume 1*. Association for Computational Linguistics, pp. 9–16.
- Hahn, U., Romacker, M., 2001. The SYNDIKATE text knowledge base generator. In: *Proceedings of the First International Conference on Human Language Technology Research*. Association for Computational Linguistics, pp. 1–6.
- Haykin, S., Network, N., 2004. A comprehensive foundation. *Neural Netw.* 2 (2004), 41.
- Hermida, J.M., Meliá, S., Martínez, J.-J., Montoyo, A., Gómez, J., 2012. Developing semantic rich internet applications with the S m 4RIA extension for OIDE. In: *International Conference on Web Engineering*. Springer, pp. 20–25.
- Hofmann, T., 2017. Probabilistic latent semantic indexing. In: *ACM SIGIR Forum*. Vol. 51, no. 2. ACM, pp. 211–218.
- Hu, Z., Huang, P., Deng, Y., Gao, Y., Xing, E., 2015. Entity hierarchy embedding. In: *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Vol. 1. pp. 1292–1300.
- Hu, M., Lim, E.-P., Sun, A., Lauw, H.W., Vuong, B.-Q., 2007. Measuring article quality in Wikipedia: Models and evaluation. In: *Proceedings of the Sixteenth ACM Conference on Conference on Information and Knowledge Management*. ACM, pp. 243–252.
- Jin, Z., Men, T., Yuan, H., He, Z., Sui, D., Wang, C., Xue, Z., Chen, Y., Zhao, J., 2022. CogKGE: A knowledge graph embedding toolkit and benchmark for representing multi-source and heterogeneous knowledge. In: *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*. Association for Computational Linguistics, pp. 166–173.
- John Walker, S., 2014. *Big Data: A Revolution That Will Transform How We Live, Work, and Think*. Taylor & Francis.
- Kevin, M., 2012. *Machine Learning: A Probabilistic Perspective*. The MIT Press.
- Kim, S.-S., Son, J.-W., Park, S.-B., Park, S.-Y., Lee, C., Wang, J.-H., Jang, M.-G., Park, H.-G., 2008. OPTIMA: An ontology population system. In: *3rd Workshop on Ontology Learning and Population (July 2008)*.
- Konstantinou, N., Spanos, D.-E., Mitrou, N., 2008. Ontology and database mapping: A survey of current implementations and future directions.
- Le, Q.V., 2013. Building high-level features using large scale unsupervised learning. In: *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*. IEEE, pp. 8595–8598.
- Lenzerini, M., 2002. Data integration: A theoretical perspective. In: *Proceedings of the Twenty-First ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*. pp. 233–246.
- Li, X., Wang, T., Pang, Y., Han, J., Shi, J., 2022. Review of research on named entity recognition. In: *Advances in Artificial Intelligence and Security*. Springer International Publishing, pp. 256–267.
- Litwin, W., Mark, L., Roussopoulos, N., 1990. Interoperability of multiple autonomous databases. *ACM Comput. Surv.* 22 (3), 267–293.
- Liu, B., 2012. Sentiment analysis and opinion mining. *Synth. Lect. Hum. Lang. Technol.* 5 (1), 1–167.
- Lozano-Tello, A., Gómez-Pérez, A., 2004. Ontometric: A method to choose the appropriate ontology. *J. Database Manag.* 2 (15), 1–18.
- Lubani, M., Noah, S.A.M., Mahmud, R., 2019. Ontology population: Approaches and design aspects. *J. Inf. Sci.* 45 (4), 502–515.
- Mikolov, T., Yih, W.-t., Zweig, G., 2013. Linguistic regularities in continuous space word representations. In: *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. pp. 746–751.
- Mitchell, T., Cohen, W., Hruschka, E., Talukdar, P., Yang, B., Betteridge, J., Carlson, A., Dalvi, B., Gardner, M., Kiesel, B., Krishnamurthy, J., Lao, N., Mazaitis, K., Mohamed, T., Nakashole, N., Platanios, E., Ritter, A., Samadi, M., Settles, B., Wang, R., Wijaya, D., Gupta, A., Chen, X., Saparov, A., Greaves, M., Welling, J., 2018. Never-ending learning. *Commun. ACM* 61 (5), 103–115.
- Montoyo, A., MartíÑez-Barco, P., Balahur, A., 2012. *Subjectivity and Sentiment Analysis: An Overview of the Current State of the Area and Envisaged Developments*. Elsevier.
- Nadeau, D., Sekine, S., 2007. A survey of named entity recognition and classification. *Lingvist. Invest.* 30 (1), 3–26.
- Nentidis, A., Bougiatiotis, K., Krithara, A., Paliouras, G., 2020. iASIS Open Data Graph: automated semantic integration of disease-specific knowledge. In: *IEEE 33rd International Symposium on Computer-Based Medical Systems*. CBMS, pp. 220–225.
- Noy, N.F., Musen, M.A., 2003. The PROMPT suite: Interactive tools for ontology merging and mapping. *Int. J. Hum.-Comput. Stud.* 59 (6), 983–1024.
- Noy, N.F., Musen, M.A., et al., 2000. Algorithm and tool for automated ontology merging and alignment. In: *Proceedings of the 17th National Conference on Artificial Intelligence (AAAI-00)*. Available As SMI Technical Report SMI-2000-0831.
- Orlowski, A., 2005. Wikipedia founder admits to serious quality problems. In: *The Register*. Vol. 18.
- Petasis, G., Karkaleitis, V., Paliouras, G., Krithara, A., Zavitsanos, E., 2011. Ontology population and enrichment: State of the art. In: *Knowledge-Driven Multimedia Information Extraction and Ontology Evolution*. Springer-Verlag, pp. 134–166.
- Pressman, R.S., 2005. *Software Engineering: A Practitioner's Approach*. Palgrave macmillan.
- Reynaud, C., Siro, J.-P., Vodislav, D., 2001. Semantic integration of XML heterogeneous data sources. In: *Proceedings 2001 International Database Engineering and Applications Symposium*. pp. 199–208.
- Rospoher, M., van Erp, M., Vossen, P., Fokkens, A., Aldabe, I., Rigau, G., Soroa, A., Ploeger, T., Bogaard, T., 2016. Building event-centric knowledge graphs from news. In: *Web Semantics: Science, Services and Agents on the World Wide Web*. Vol. 37–38. pp. 132–151.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A.C., Fei-Fei, L., 2015. ImageNet large scale visual recognition challenge. *Int. J. Comput. Vis.* 115 (3), 211–252.
- Shah, D.V., Cappella, J.N., Neuman, W.R., 2015. Big data, digital media, and computational social science: Possibilities and perils. *Ann. Am. Acad. Polit. Soc. Sci.* 659 (1), 6–13.
- Sheth, A.P., Larson, J.A., 1990. Federated database systems for managing distributed, heterogeneous, and autonomous databases. *ACM Comput. Surv.* 22 (3), 183–236.
- Shvaiko, P., Euzenat, J., 2013. Ontology matching: State of the art and future challenges. *IEEE Trans. Knowl. Data Eng.* 25 (1), 158–176.
- Sowa, J.F., et al., 2000. *Knowledge Representation: Logical, Philosophical, and Computational Foundations*. Vol. 13. Brooks/Cole Pacific Grove.
- Staab, S., Studer, R., 2010. *Handbook on Ontologies*. Springer Science & Business Media.
- Stoermer, H., Palmisano, I., Redavid, D., Iannone, L., Bouquet, P., Semeraro, G., 2006. Contextualization of a RDF knowledge base in the VIKEF project. In: *Digital Libraries: Achievements, Challenges and Opportunities: 9th International Conference on Asian Digital Libraries, ICADL 2006, Kyoto, Japan, November 27-30, 2006*. Proceedings. Springer Berlin Heidelberg, pp. 101–110.
- Suchanek, F.M., Ifrim, G., Weikum, G., 2006. LELA: Learning to extract information by linguistic analysis. In: *Proceedings of the 2nd Workshop on Ontology Learning and Population: Bridging the Gap Between Text and Knowledge*. pp. 18–25.
- Turian, J., Ratinov, L., Bengio, Y., 2010. Word representations: A simple and general method for semi-supervised learning. In: *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, pp. 384–394.
- Vassiliadis, P., 2009. A survey of extract–transform–load technology. *Int. J. Data Wareh. Min. (IJDW)* 5 (3), 1–27.
- Vdovjak, R., Houben, G.-J., 2001. RDF based architecture for semantic integration of heterogeneous information sources.
- Weber, N., Buitelaar, P., 2006. Web-based ontology learning with isolate. In: *Proc. of the Workshop on Web Content Mining with Human Language At the International Semantic Web Conference*. Vol. 11. Athens GA, USA.
- Wu, X., Zhu, X., Wu, G.-Q., Ding, W., 2014. Data mining with big data. *IEEE Trans. Knowl. Data Eng.* 26 (1), 97–107.