Umanistica Digitale – ISSN: 2532-8816 – n. 22, 2026: Special Issue
S. Ferilli, E. Bernasconi, D. Redavid, G.M. Di Nunzio – *Knowledge Graphs for the Web Economy*
DOI: http://doi.org/10.60923/issn.2532-8816/22159

Knowledge Graphs for the Web Economy
The CHiPS&BITS Project on Cultural Heritage

Stefano Ferilli

University of Bari, Bari, Italy
stefano.ferilli@uniba.it

Eleonora Bernasconi

University of Bari, Bari, Italy
eleonora.bernasconi@uniba.it

Domenico Redavid

University of Bari, Bari, Italy
domenico.redavid1@uniba.it

Giorgio Maria Di Nunzio

University of Padua, Padua, Italy
giorgiomaria.dinunzio@unipd.it

Abstract

The growing demand for Cultural Heritage fruition is making it a major economic driver, also in connection with its tourism-related aspects. The current solutions available on the Web are still unable to provide satisfactory support to the various kinds of stakeholders and to the different applications. The History of Computing, as a peculiar, relevant and currently under investigated branch of Cultural Heritage, raises additional challenges and provides new opportunities, also in connection with a significant economic flow it generates. The variety and complexity of issues connected to this domain call for even more advanced solutions. In this paper we introduce the CHIPS&BITS project, which tackles these problems using a knowledge-based approach and leverages a novel framework that can meet the needs associated to this specific domain in a better way compared to standard Semantic Web approaches. We describe the contributions of the project to overcome the limitations of the state of the art, and focus on some of its more peculiar and original features, among which the definition of suitable ontologies to describe the complexity of the domain and the use of Artificial Intelligence algorithms for giving the users an advanced and personalized experience.

Keywords: Knowledge Graphs, Semantic Web, Web Economy, Cultural Heritage

La crescente domanda di fruizione dei Beni Culturali ne fa un importante motore economico, anche in relazione agli aspetti turistici. Le soluzioni attualmente disponibili sul Web non sono ancora in grado di fornire un supporto soddisfacente ai vari tipi di stakeholder e alle diverse applicazioni. La Storia dell'Informatica, in quanto branca peculiare, rilevante e attualmente poco indagata del Patrimonio Culturale, pone ulteriori sfide e offre nuove opportunità, anche in relazione al significativo flusso economico che genera. La varietà e la complessità delle questioni legate a questo dominio richiedono soluzioni ancora più avanzate. In questo articolo presentiamo il progetto CHIPS&BITS, che affronta questi problemi utilizzando un approccio basato sulla conoscenza e sfruttando un framework innovativo in grado di soddisfare le esigenze associate a questo specifico dominio in modo più adeguato rispetto agli approcci standard del Semantic Web. Descriviamo i contributi del progetto per superare i limiti dello stato dell'arte e ci concentriamo su alcune delle sue caratteristiche più peculiari e originali, tra cui la definizione di ontologie adatte a descrivere la complessità del dominio e l'uso di algoritmi di intelligenza artificiale per offrire agli utenti un'esperienza avanzata e personalizzata.

***Parole chiave**: grafi della conoscenza, web semantico, web economy, beni culturali*

Introduction

The demand for Cultural Heritage (CH) fruition has grown steadily in recent years, making it a major component in the economic balance of many countries and regions. CH does not only generate direct revenues through access fees to cultural items, but also drives tourism-related income through travel, accommodation and subsistence expenditure. It is often tightly linked to local art, traditions and folklore, which visitors are keen to experience once they reach their destination. All these activities, each with its own issues and complexities, are nowadays mediated by digital services; however, they are usually supported only at a superficial level, with separate platforms for booking travel, accommodation, meals and tickets. From the tourist's perspective, by contrast, these elements are part of a single, integrated experience in which each component depends on and influences the others.

So, it is fundamental to tackle this landscape in a holistic way, in order to provide the final user with an approach that resembles more closely his perspective and can help him in a more effective way, also boosting the associated economic value as a side effect. This objective requires connecting, combining and integrating all the pieces of information from the different aspects, which is what we call `knowledge'. Hence, our knowledge-based approach to the task. Specifically, we propose a Knowledge Graph for storing and handling all of the above information, based on semantic technologies that ensure interoperability between different systems and platforms, make the procedures and results personalized and understandable to the final users, and enable the application of advanced Artificial Intelligence (AI) solutions to overcome the limitations of the current solutions. In particular, we propose the use of a framework we are currently developing that can support the needs of this application better with respect to standard solutions based exclusively on Semantic Web technologies.

Some examples are the following:

- The full LOD perspective is not very suitable for preserving privacy (e.g., in auction, rental or exchange platforms); in CHIPS&BITS, the items owned by the users, their selling price and other sensitive information is not to be disclosed unless the users give permission.
- Suppose user *u* owns device *d*, then sells it and then buys it again. This requires to set two instances of relationship *owned* between *u* and *d*, having each the initial and final date of ownership as attributes. This is not immediately feasible using the RDF representation.
- Storing the knowledge on one DB enables the use of analysis, mining, and reasoning algorithms that would not be applicable on the `unbounded' data organization of the Semantic Web. Still, the knowledge in CHIPS&BITS can interoperate with the knowledge stored in the Semantic Web.

The specific branch of CH we approach in this paper is the history of computing, due to cultural, technological and economic motivations:

- Culturally, the increasing importance but rapid obsolescence of computers recently raised much interest in vintage computing knowledge and artifacts as Cultural Heritage. A wealth of precious information is being discovered, retrieved and preserved, also thanks to the computing pioneers. It is an opportunity that will soon vanish to collect, store, organize, deliver, and pass to future generations all this information.
- From the technological viewpoint, the intrinsic complexity of the domain makes it a challenging problem, requiring novel solutions that go beyond the state-of-the-art. In fact, it requires to tightly and seamlessly integrate several aspects that are currently dealt with by different communities: hardware, software, documentation, and even intangible aspects of computing, plus the aspects concerning tourism and local traditions.
- From the standpoint of (Web) economy, the size of the economic flow revolving around this topic, coming not so much and only from the CH fruition aspects, but from the preservation, restoration and exchange activities flourishing in a world-wide community of enthusiasts, collectors and practitioners; and from the possibility of leveraging this non-standard kind of CH to open new, season-independent touristic revenue opportunities also for territories that are outside of the mainstream touristic flow connected to traditional CH.

Our efforts in this direction resulted in the starting of CHIPS&BITS (acronym of 'Computing Heritage Intelligent Platform for Searching & Browsing Information Technology and Science'), a self-funded project aimed at developing a knowledge-based platform for the history of computing, and advanced/innovative Web-based tools to allow its users to effectively use and experience, search and discover its content, and generate value from it, dealing with the complexity of the subject matter. The objective is manyfold: spreading awareness of, and supporting education on, such important and still mostly ignored branch of CH; collecting and preserving knowledge and materials that are at stake of being lost forever, and supporting all the different stakeholders interested in this topic in their activities, including those generating economic value.

CHIPS&BITS proposes innovative solutions and tools for all the above problems and, to the best of our knowledge, is among the first platforms to address the history of computing through an integrated knowledge graph:

- where stakeholders can put, interrelate, exchange, search information on the history of computing both from a scholarly perspective and for fostering its fruition in a touristic perspective, possibly generating money exchange and job opportunities as a side effect;
- featuring a unified and coherent knowledge base and repository collecting information and documents on the history of computing;
- featuring a data schema/ontology that can describe this field in all of its complexity, seamlessly dealing with the different facets of computing: hardware, software, documentation, intangible aspects, economic exploitation;
- exploiting advanced AI solutions to ensure more comprehensive, accurate, diverse and personalized results to its users;
- compared with existing initiatives such as ArCo, the Steve Jobs/Pirates of Silicon Valley graph, and the Computer History Museum ontology (see the Related Work section), CHIPS&BITS explicitly combines hardware, software, documentation, intangible aspects and economic information in a single, extensible knowledge base.

Since the proposed approach is general, it may act as a pilot for other branches of CH.

This paper is organized as follows. After discussing the motivations and contributions of this work in the next section, the ontology for the history of computing is presented in Section Technical Details. The section Personalization and Privacy discusses how the system implements user profiling policies, how these impact knowledge extraction and how privacy is managed. Then, Section Current Status of the Project briefly describes the features and interactive interface of GraphBRAIN, and its current content. Finally, after discussing some Related Work, Section Conclusions and Future Work concludes the paper and outlines future work issues.

Motivations & Contributions

The CHIPS&BITS project aims at contributing under several aspects, as described in the following sub-sections. The motivations and the specific key points for each aspect can be summarized as follow:

- Cultural Heritage
  - Computing's rapid advancement makes it a significant cultural heritage item.
  - Urgency to document and preserve vintage computing knowledge while contributors are still active.
  - Collaborative tools to build a knowledge base which align with Open/Citizen Science principles.
- Domain Complexity
  - Interdependency of hardware, software, documentation, and intangible factors.
  - Need for a holistic ontology to capture and manage the complexity of computing heritage.
  - Current cataloging standards inadequate for the field's uniqueness.
- Web Economy
  - Increasing vintage computing market with high-value exchanges and auctions.
  - Crowdfunding and events generating economic flow and public interest.
  - Enhancing awareness, connections, and personalized stakeholder support.
- Technology
  - Ontology-driven knowledge representation to manage domain complexity.
  - Proprietary database (GraphBRAIN) ensures data protection and flexible access control.
  - Advanced AI techniques enable user-friendly and diverse knowledge delivery.

Cultural Heritage

While Computer Science and Engineering are quite young areas in the landscape of human knowledge, their incredibly rapid advances in the last decades, and the relevance computers gained in every aspect of our lives, recently raised significant interest in the study of the history of computing and in the preservation of knowledge and artifacts related to it. So, computing is now not just a means to support and foster all activities related to CH (known as *Digital Cultural Heritage)* [1], but also it is becoming the very object of study and fruition as CH items, as well. Museums and private collectors started popping up all around the world. Such an ongoing excitement for vintage computing is recently pushing the discovery, retrieval and preservation of huge amounts of precious information on this topic. Many people that by first-hand contributed to the development and to the milestones of the field are still alive, and can be included in the loop, which makes this is an unprecedented opportunity, that probably will soon vanish and never happen again, to collect, safely store and sensibly organize all this wealth of information, so that it may be made available to future generations for research, study or education purposes. Unfortunately, there are some obstacles that may tamper this objective. First, the field is so new that there is no established research and scholarship yet. Second, knowledge in this field is spread across many people, each of which has just a partial (incomplete and biased) view of the whole story. What is worse, technology in this field suffers from extremely rapid obsolescence, and thus new practitioners tend to ignore important technological information needed to understand and handle items of just a few years earlier. Hence, the need for education about the history of computing and the way in which vintage hardware and software worked.

Leveraging the enthusiasm of practitioners in this field, CHIPS&BITS provides a (set of) tool(s) for collaborative building and enrichment of a knowledge base covering all aspects of the history of computing from personal, scientific, industrial sources. In line with the EU idea of Open/Citizen Science[1], putting together all the `partial' pieces should result in a picture that is more complete, more understandable and, ultimately, more valuable than the simple sum of the pieces. Every new piece of information opens new knowledge paths, supporting new ways for discovering information and completeness, accuracy and diversity of the search results. Also, comparing different sources it can spot controversial cases and support trustworthiness.

Domain Complexity

The production of computing devices is much more complex than other kinds of manufactures. First of all, it inextricably connects to each other aspects that are usually in the realm of different communities, using different platforms and standards: hardware (museums), software (Internet repositories), documentation (libraries and archives), and intangible aspects (people). Indeed, hardware is just dead matter if software does not give it life, and software is meaningless and useless if detached from the hardware platform it is intended for; both require documentation to understand and properly run them; finally, why hardware, software and documents are as they are can often be understood only in connection with events, customs, anecdotes, etc. that (sometimes fortuitously) determined a course of history rather than another. The development of these objects involves many contributors, at different levels and with different roles.

Even considering hardware and software only, complexity is taken to the extreme, and stands apart from all other CH items (see [2] for a more comprehensive discussion of this).

Devices and software are tightly coupled: each is often meaningless without the other, yet they require very different kinds of metadata (for example source code, libraries and support media). Mass-produced objects coexist with prototypes and handcrafted “homebrew” artefacts, which demand different cataloguing strategies and rich descriptions of manufacturing processes and materials. Both devices and software are compound objects, with configurations, versions and variants, and with secondary objects (manuals, packaging, peripherals) that may be as important as the main item. Finally, many devices are fully understandable only within systems and configurations, where subtle compatibility constraints between hardware, software and peripherals play a crucial role.

Existing cataloguing and description standards proposed for CH items, including those specifically developed for technological objects, were not designed to cope with all these intertwined technical, social and economic dimensions. As a result, they only partially capture the complexity described above and make it difficult to express it in a uniform way. Hence, a pressing need for the definition of a specific new scheme, to be shared and reused by all the stakeholders involved in this area of interest. However, this is a challenging task, because there is no standard, nor precise categorization taxonomy for computing elements, which makes it difficult to define the classes, assign them stable attributes, and organizing them into a hierarchy.

1 https://openscience.eu/article/project/foster

In this landscape, CHIPS&BITS contributes by adopting a knowledge-based approach and defining a 'holistic' ontology, to be used as a data schema, bringing together all these aspects.

Web Economy

There is a lively market for vintage computing elements and memorabilia, with exchanges, sales and auctions ongoing all over the world. Some items were sold for values similar to those of works of art by famous artists (e.g., one of the surviving units of Apple I, the first computer developed by Apple, was sold for $ 375,000; one of the prototypes of Commodore 65 – a computer which was never released –, endowed with a – probably unique – expansion board, was sold for about $ 95,000; some of the first models of personal computers, produced in thousands units, are sold for several thousand dollars, and so on). This market relies on standard, non-domain-aware platforms, providing the users (sellers and buyers) absolutely no support to understand and put in perspective the items they are exchanging.

Dozens crowdfunding campaigns have been successfully run aimed at reproducing old computers (or building new ones based on obsolete technology), developing new software for vintage machines, writing books on the history of computing, producing films and documentaries, establishing museums, etc. The number of supporters is in the order of thousand people, and the money collected for each of these campaigns is typically in the order of tens of thousands of euros, often ending up in the hundreds of thousands. Also, many kinds of events (conferences, seminars, shows) on `vintage computing' have been organized, attracting scholars, researchers, amateurs, collectors or other kinds of enthusiasts. This generates tourism, and associated economic flow, as a side effect. Finally, additional economic flow is associated to job opportunities in the form of rental of items, service or restoration, consultancies, and education.

CHIPS&BITS aims at supporting all these activities by providing the stored knowledge to the various kinds of stakeholders involved, making them more aware, boosting the connections and exchanges, and helping to match requests and offers, all in a personalized way.

Technology

Supporting the previous needs also poses technological challenges that require new solutions to overcome the limitations of the current state-of-the-art.

A collaborative approach in which many people, with different expertise, culture, background and perspective contribute small pieces that together make up the big picture, requires powerful and shared representation and management solutions to organize the knowledge.

CHIPS&BITS adopts a knowledge-based approach based on ontologies as schemas for representing and interconnecting all the different kinds of data. The integration of different ontologies describing the different aspects allows to deal with the complexity and peculiarities of the field. A valuable result by itself, it enforces qualitative approaches, more focused on concepts than on statistics. The ontology itself will be searchable knowledge, useful for understanding the domain. Since CH data may have Intellectual Property-related restrictions, and their economic exploitation may involve private or sensitive information, CHIPS&BITS guarantees data protection, overcoming some limitations of both the collaborative approach and the LOD perspective, while still supporting systems' interoperability. This is obtained by the use of GraphBRAIN, a framework that is based on a proprietary DB to store the data, so that the data owners may decide what, when, how and to whom can be disclosed. The portion meant to be public can be published as LOD, and transparently be perceived as such by standard Semantic Web technologies. In particular, GraphBRAIN was extended to add flexibility to the data visibility settings, by which the data owner may decide at several levels of granularity who is guaranteed access to each kind of information stored in the knowledge base. Finally, the fruition of such a varied and complex knowledge to serve the needs of different kinds of stakeholders and applications, including research, hobby, trade, and education, requires the system to ensure user-friendliness, personalization, diversity, accuracy, and comprehensiveness in knowledge delivery.

CHIPS&BITS uses advanced AI solutions based on Natural Language Processing (e.g., for expressing in natural language the formal knowledge extracted from the graph), Document Image Analysis (to automatically extract knowledge from existing documents), Knowledge Representation and Reasoning (to obtain implicit information in an explainable way), Network Analysis (to understand the relevance and possible connection of knowledge items), and Machine Learning/Data Mining (to learn new knowledge from the knowledge base or profiles from user behaviors) to extract and deliver relevant knowledge by fulfilling all of these requirements. The use of qualitative approaches (not affected by frequency and explainable), in addition to quantitative ones, enforces diversity in the outcomes and trustworthiness of the system and of its outcomes.

At the time of writing, the graph-mining, relevance assessment, subgraph extraction, link prediction, recommendation and natural-language generation functions described in Section "Technical Details" are already available in GraphBRAIN and have been exercised on the CHIPS&BITS knowledge base. Work is ongoing to consolidate these components into a stable service layer and to extend them with large-scale document image analysis and the LLM-based extraction tool mentioned in concluding section.

Technical Details

This section provides details about some of the technical solutions embedded in CHIPS&BITS. A graphical overview of the framework is shown in Figure 1.

GraphBRAIN Framework

GraphBRAIN is a general-purpose knowledge base management system aimed at covering all stages and tasks in the lifecycle of a knowledge base, including knowledge acquisition, organization, and (personalized) fruition[2]. It adopts the Labeled Property Graph (LPG) data model, where nodes (representing individuals) and arcs (representing relationships) may have labels (usually expressing their type) and associated attribute-value pairs, and uses the Neo4j DBMS [3]. Neo4j is schema-less, which ensures great flexibility but does not allow to associate a clear semantics to the graph items. For this reason, GraphBRAIN requires its users to work according to pre-specified data schemes, expressed in the form of ontologies. Thus, a

2 A demo of the system can be found at http://digitalmind.di.uniba.it:8088/GraphBRAIN/

characterizing feature of GraphBRAIN is its bringing to cooperation a database management system for efficiently handling, mining and browsing the individuals, with an ontology level that allows it to carry out formal reasoning on the knowledge. GraphBRAIN can apply several ontologies on the same graph, representing different perspectives on the same knowledge. The classes shared by different ontologies allow the system to connect knowledge across domains: their individuals act as bridges, allowing the users of a domain to reach information coming from other domains. The ontologies are expressed in a proprietary format, specifically tailored for LPGs [4].

In this context, adopting a labeled property graph rather than a pure RDF triple store is motivated by several requirements. Relationships in CHIPS&BITS, such as ownership, need to carry multiple attributes (for example start and end dates, transaction price and privacy flags) and to be queried as first-class citizens, which is more naturally supported in an LPG setting. Graph-mining and path-based personalization algorithms also benefit from the index-free adjacency and native traversal operators of graph databases, such as Neo4j. At the same time, interoperability with the Semantic Web is preserved by the OWL import/export layer: when data can be published, the relevant parts of the knowledge base are exported as RDF/OWL and aligned with standard CH ontologies, while more sensitive fragments remain confined to the protected LPG store. In other words, the LPG back-end and the Linked Data interfaces are complementary rather than competing technologies, and Linked Data in our setting does not imply that all data must be open.

Both ontologies and instances may be imported from, or exported to, the standard Semantic Web format Web Ontology Language (OWL)[3], in order to support their interoperability and reuse as linked open data (LOD) [5]. The ontologies are freely available, and accessible from the on-line GraphBRAIN platform. After logging in, any user can go to the "Schema" tab, then select the ontology of interest from the drop-down menu on the top-right of the page, and finally press the "Export as:" GBS button on the bottom of the page. This will start a download of the overall selected ontology (including ontological items coming from imported schemas).

GraphBRAIN provides a top-level ontology defining very general and highly reusable concepts and relationships (e.g., Person, Place; Person.wasIn.Place; etc.). This top-level ontology plays a crucial role to interconnect the domain-specific ontologies, ensuring an overall connected knowledge graph. Using a suitable tool, GraphBRAIN administrators may create, build and maintain additional ontologies.

After setting up the ontologies, information can be fed into the knowledge base. A form-based interface allows users to manually insert/update/remove instances or to query the knowledge base for instances of entities and relationships: they must select one of the available domains/schemas, and the forms are automatically generated by the system starting from the corresponding ontologies. The knowledge base can also be fed by automatic knowledge extraction from documents and other kinds of resources (e.g., books or the Internet). Instances may also have attachments, making GraphBRAIN a digital library, whose content is organized according to

3 http://www.w3c.org/owl

formal ontologies, fostering interoperability with other systems. Users may add, show, or delete attachments.

In a collaborative spirit, users may add comments on (to provide suggestions or add information), or approve/disapprove, each entity or relationship instance, each single attribute value thereof, and even the ontological items. This feedback is used to assign a trust value to the users.

Another interface allows users to display a portion of the graph, browse it interactively and display detailed information about entity and relationship instances. This allows the user to continue his search in a less structured way, by exploring the graph (expanding or compressing node neighbors) with a predefined goal in mind, or letting the data themselves drive the search, possibly finding relevant information in a serendipitous way.

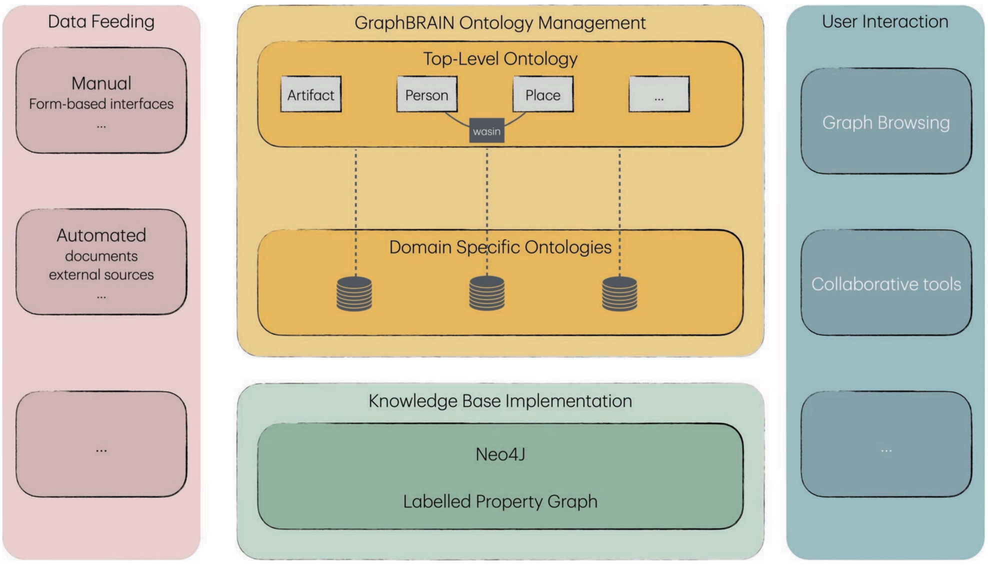

Figure 1: Overview of the GraphBRAIN framework

The users can also run on the available knowledge several algorithms for graph mining, network analysis, information extraction, automated reasoning, natural language processing, etc. Some of these algorithms are reused from the literature; others are original. Currently included functions are:

- assess relevance of nodes and arcs in the graph, and extract the most relevant ones;
- extract a portion of the graph that is relevant to some specified starting nodes;
- extract frequent patterns and associated sub-graphs;
- predict possible links between nodes;
- recommend relevant knowledge items;

- translate into natural language the information content of a portion of the graph.

If available, a user profile can be used to personalize the results of all these algorithms. This would ensure that each user obtains tailored information, which is another novelty introduced by GraphBRAIN.

While a proprietary technology, all the functions of the GraphBRAIN framework are exposed as services of an API whose code will be released for free use, so that any third-party application and organization will be able to exploit it.

Holistic Ontology

Currently the CHIPS&BITS knowledge graph is organized according to a merger of the following ontologies:

- **general** dealing with very general concepts and relationships that are expected to be present in almost all domains; of interest to CHIPS&BITS are, e.g., entities **Agent** (including persons, organizations and users), **Place** (geographical, administrative, locations, etc.), **Event** (both historical ones and conferences, fairs, shows, etc.), **Collection** (any kind of grouping), **Document** (in its most general definition as "something that serves as evidence or proof", and thus not limited to printed (or printable) documents), **IntellectualWork** (the original result of an intellectual effort, such as inventions, subjects, works of art, etc.), **Award** (any kind of recognition that can be awarded to, or record that can be marked by, persons, companies, devices, documents, or components: e.g., Education achievements, Prizes, Records), **Item** (a specific, identifiable specimen of a serially produced object).
- **lam** concerning galleries, libraries, archives and museums; it includes and is aligned to the standard data models from IFLA (the International Federation of Library Associations) [6], is connected to the general ontology via concepts such as **Document**, **Collection** (for series and libraries), **Item**, and **IntellectualWork**, and provides CHIPS&BITS with many representational components for documentation.
- **education** describing educational aspects such as subjects, learning resources, learning paths, learning accomplishments, etc.; it includes and is aligned to standard data models such as the IEEE LOM (Learning Object Metadata), IntelLEO (an ontology on educational procedures developed within a European research project), and OERschema (an ontology on educational organization at its early stages of development) [7]; it is also connected to the lam ontology via the concepts concerning documents, and obviously to the general ontology via various high-level concepts.
- **computing** The core of CHIPS&BITS, for which more details will be provided below.
- **tourism** concerning touristic-related information, including history, cultural heritage items, points of interest, transportation, hospitality, logistics and services; it was borrowed from another ongoing effort [8] in which GraphBRAIN will underlie an

integrated system aimed at supporting all stakeholders involved in touristic activities (tourists, entrepreneurs and institutions).

- **food** Complements the tourism section by adding information about food&beverage, especially from the perspective of typical dishes and beverages from specific regions of touristic interest.

These shared elements among these ontologies allow CHIPS&BITS to relate knowledge items from different domains, extending in this way the available scope of search beyond the single perspectives. For example, a tourist interested in the history of computing, while in Bari, might be spotted the chance to visit the collection at the Department of Computer Science, in order to see a specimen of the *Olivetti Programma 101* computer.

To get general information about the CHIPS&BITS ontology, Figure 2 shows a protégé screenshot of the GraphBRAIN format ontology exported to OWL where the metrics related to the axioms can be observed.

A sample of domain-specific classes from the computing ontology is the following:

- **ElectronicComponent**: a part, useful or needed to build a **Device** but not operable directly by the final user (e.g., some of its subclasses are **PrintedCircuitBoard**, **IntegratedCircuit**).
- **Device**: a manufact that can be operated directly from the final user (e.g., it includes as subclasses **Calculator**, several kinds of **Computer**, **InputDevice** and **OutputDevice**, **StorageDevice**, etc.).
- **Software** (of various kinds: **DevelopmentSoftware**, **EducationalSoftware**, **EmbeddedSoftware**, **OfficeAutomationSoftware**, **OperatingSystem**, **Videogame**, …).
- **Package**: a specific packaging of a **Device** or of a **Collection** of devices sold together (often very relevant for collectors).
- **Configuration**: a relevant group of **Device**, relevant because typical or determined in order to satisfy specific needs (e.g., a configuration of devices for desktop publishing).
- **System**: a group of **Device** that is functional only as a whole; it differs from a **Configuration** in that, in a **Configuration**, at least one of the **Device** would be functional if taken alone.

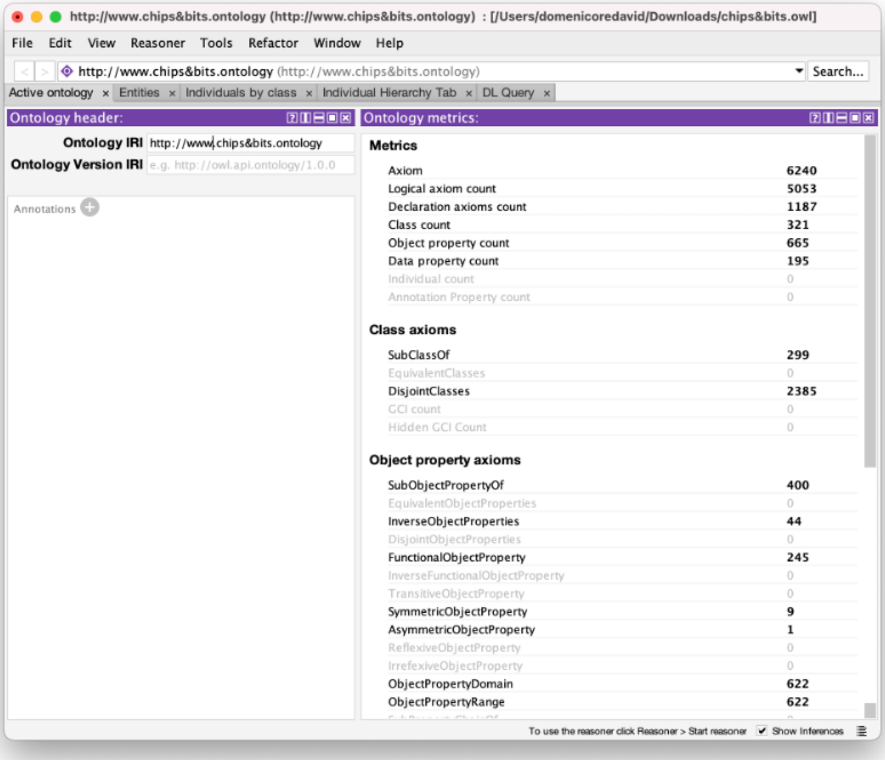

Figure 2: CHIPS&BITS ontology metrics.

Also useful are the concepts of **Award** (that may be useful to note the firsts in the history of computing), **IntellectualWork** (extended to account for several subclasses: e.g., **Algorithm**, **Approach**, **Invention**, **Programming Language**, **Technologies**, **Theorem**, **TheoreticalModel**), **Item** (extended from just documents to components, devices, software, etc.), borrowed from the other ontologies.

A sample of domain-specific relationships are **Agent.acquired.Item** (including donations and purchases), **clones**, **compatibleWith** and **mayReplace** (between pairs of components, devices or software), **owned** (for agents who owned organizations, but also to record notable previous owners of some items), **Component.partOf.{Item,Device}** (some rare parts may dramatically change the value of an item or device), **evolves** (for hardware and software versions), **requires** (software may require another software or a device to work properly), **repaired** (to record who is able to repair what). Some of these are relevant for the economic aspects, others are relevant for repairing and restoration, which is extremely important because computing devices has peculiarities that cannot be expressed in existing ontologies designed for other kinds of CH.

The proposed conceptualization may be a starting point for the definition of cataloguing standards for cultural heritage material related to the history of computing. Indeed, existing standards for cultural heritage, even those developed for technological and scientific elements (e.g., [9]), cover only parts of this specific field and do not provide a single coherent model for the full range of complexity and subtleties that CHIPS&BITS needs to represent.

For reproducibility and reuse, the full set of ontologies used in the current CHIPS&BITS deployment is available through the GraphBRAIN schema export interface described in Section "Technical Details". A versioned snapshot of these ontologies will also be mirrored in an open repository (http://digitalmind.di.uniba.it:8088/GraphBRAIN) so that they can be cited and reused independently of the platform.

Personalization and Privacy

While designed according to the linked data perspective, the GraphBRAIN knowledge graph is not available in its entirety as Linked Open Data. Indeed, it is not directly accessible to the public. Access is available only through the query, graph browsing and knowledge extraction facilities provided by the API, that the CHIPS&BITS application must use. This is a fundamental feature for CHIPS&BITS: as regards personalization, satisfying the needs of many different users and kinds of stakeholders requires the ability to return tailored results to each of them. Concerning privacy, some users might not want to disclose some details about the knowledge they provided (e.g., economic values of transactions, ownership or place of storage of some items, etc.).

Each function in the API, exposed as service, returns relevant portions of the graph based on the input parameters, on the user's profile (if available), and on the privacy specifications for the various knowledge items. Personalization and Privacy are handled with a system of registered users and with the help of a (relational) database, which was inherited and extended from the GraphBRAIN framework. Privacy is strictly related to personalization, in that each user must decide what parameters can be exploited to build his model, and which functions can access which of these parameters.

User Profiles

The user profile takes the form of a set of weights, associated to the ontological elements (domains, entities, relationships or attributes) or to specific entity or relationship instances (i.e., nodes or arcs in the graph). The weight is formed and continuously updated by taking into account both explicit preference indications by the users and implicit preferences computed on usage. Many kinds of user interaction contribute to the latter part: selection of an entity or relationship, access to a specific node, arc or attribute (to display its value or to add, modify or delete it), approval/disapproval or comment on a node, arc or attribute value. These parameters are stored in 3 tables of the relational DB, that associate users to the following information:

- **Preferences** explicit preferences, in the form of fixed weights for different ontological items, domains or even specific instances, also depending on the schema in which such an item is used and even on the application that is being used for accessing GraphBRAIN;

- **Authors** logs all Create, Update or Delete operations on the entity or relationship instances, including their associated domains and attributes, reporting for each: the instance being changed, the type of information being changed, the type of operation, the old value (if any), the user who performed the change, and the datetime of the operation;
- **Readers** logs all selections of a domain and all Read operations to entity or relationship instances, also noting the domain and application under which they were accessed, the user who performed the access, and the datetime of the operation;
- **Evaluations** stores all approval/disapproval and comments made on the instances of different kinds of ontological items, along with the user who performed the access, and the datetime of the operation.

Additional information useful to build the user's model can be gathered from other applications. Indeed, CHIPS&BITS uses an external user management system (called Mondo) that serves different applications. If other applications served by this system are based on GraphBRAIN and use the same ontologies, but embed different KGs, Mondo can ask them to provide information about their profile for the same user and forward it to another application. Of course, these other applications will also adopt the same approach to user modeling, so in the following we will only need to describe how the mentioned DB tables can be used to determine a user's profile. Then, the profile can be used for several purposes. Currently considered are the following:

- to notify the user about changes in the knowledge associated to his preferred elements;
- to drive knowledge extraction and graph mining algorithms so that they can focus more on his preferred elements and less on the others;
- for recommendation purposes;
- more in general, for assessing the degree of relatedness between users, items, or users and items;
- for clustering graph items.

Based on the data in the mentioned tables, a profile assigns partial weights in [0,1] to KG items, one for each perspective (i.e., DB table):

- w_p associated to explicit preferences of the user;
- w_a computed for knowledge modifications made by the user;
- w_r computed for knowledge accesses performed by the user;
- w_e computed for the evaluations entered by the user.

The overall weight for an item is computed as a weighted combination of the partial weights, using the following weights:

- $\alpha_p \in [0,1]$ the relevance of the explicit preferences in the overall weight;

- $\alpha_a \in [0,1]$ the relevance of knowledge modifications in the overall weight;
- $\alpha_r \in [0,1]$ the relevance of the knowledge accesses in the overall weight;
- $\alpha_e \in [0,1]$ the relevance of the evaluations in the overall weight.

All the above (partial or overall) weights can be computed by referring to a specific domain and/or application, by just filtering the rows involving that domain or application before counting them.

Knowledge Extraction

Most knowledge extraction functions in GraphBRAIN are based on the computation of a subgraph. So, plugging a personalization strategy into the subgraph extraction algorithm results in consequent personalization of all functions based on this algorithm. For this purpose, we considered a modification of the Spreading Activation algorithm, proposed in psychology research [10] and then borrowed from AI with Semantic Networks [11]. In a nutshell, the idea underlying this algorithm is that an activation value can be associated to each node in the graph. Starting from an initial set of relevant nodes with a pre-defined activation value, activation is propagated to adjacent nodes, and the process is iterated, with a general decay factor and a specific decay factor associated to each arc (expressing how much of the source activation budget must be brought to the sink node), until some stop condition is met (e.g., a given number of nodes has been activated, or the activation falls below a given threshold). In the original algorithm, the activation value propagated from a source node *i* to a sink node *j* through an arc having decay factor *w(i,j)*, with general decay factor *d,* is

$$a(j) = a(j) + a(i)\, w(i,j)\, d$$

Note that in LPGs different arcs carrying exactly the same information can still be distinguished by their identifier, and will be processed separately by this loop.

In our approach, not only the arcs, but also the nodes may have an associated decay value: so, the propagation from a source node along an arc considers the decay values associated to both the arc and the sink node. These values are in [0,1], and instead of being fixed for nodes and arcs in the graphs they are computed based on the user profile's weights for each user (or group of users). Also, we assume activation values value to be real numbers in [0,1], with 0 meaning no activation (i.e., irrelevance) and 1 meaning maximum activation/relevance. Initially, all nodes have activation equal to 0, except for the starting nodes which take an initial activation value defined by the user (to reflect different importance for them) or automatically computed, again, using the user profile's weights. As to the starting nodes, again, they can be explicitly selected by the user, or automatically computed for him based on his profile's weights.

Privacy Management

Privacy is obtained by adding to each ontological element (entity, relationship or attribute) in the ontology definition a *privacy* attribute, which can be set to True or False. When True, each user can set the privacy value of that ontological item, or of specific instances thereof, to one of the values:

- **Private** visible only to its owner,

- **Restricted** only selected users can access it, or
- **Public** visible to all users and publishable as open data.

This information is stored in a 'Privacy' table in the relational DB. When Restricted, another table 'Permission' in the DB reports, for each node or arc as a whole, or for its labels and attributes, which users may access it, specifying for each of them if they can see, add, modify or delete it (in case of a node or of an arc) or its value (in the case of an attribute).

This privacy value will be applied to all the pieces of information (i.e., nodes, arcs or attribute values in the graph) owned by that user (i.e., that were added to the graph by that user). Given a piece of information, the system will first look for a specific privacy setting for it and apply it; if missing, it will apply the privacy level set by the user to its kind (entity, relationship or attribute).

Current Status of the Project

CHIPS&BITS currently involves 11 ontologies, describing various contextual or application domains related to the history of computing (general, libraries/archives/museums, education, tourism). The dependencies among such ontologies are graphically shown in Figure 3, where lower-level ones import the concepts from the higher-level ones. The figure shows that ontologies purposely defined in GraphBRAIN all import the `general' ontology, while ontologies obtained by translating existing schemes are stand-alone. However, a connection still exists between these classes and the other ones, because in the translation an alignment was carried out, and the same names were used for the classes referring to the same concept.

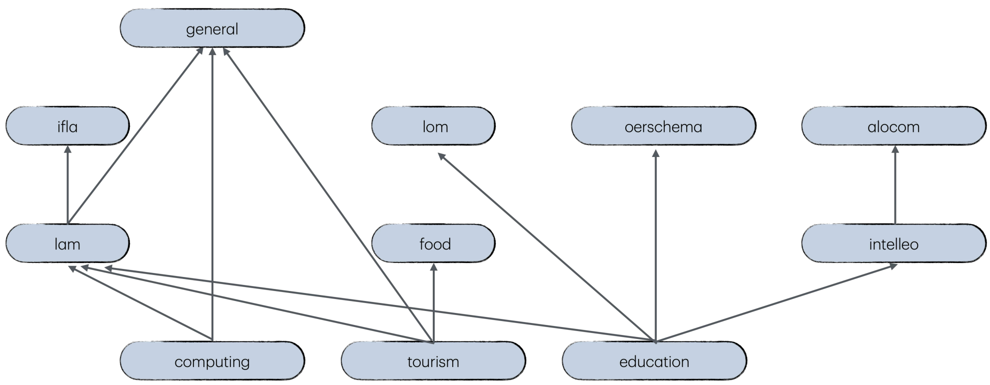

Figure 3: Dependencies among ontologies involved in CHIPS&BITS

Table 1 reports statistics on the number of ontological elements (classes, relationships, attributes) in these ontologies. The number of top classes is reported separately from the number of their subclasses, in order to distinguish the specializations from the core concepts. Since the model underlying GraphBRAIN is an LPG, also relationships may have attributes. Note that in GraphBRAIN one relationship name may represent many specific relationships, distinguished by different subject-object pairs (which would require different names in OWL). So, the `Relationships' column reports the number of different relationship names followed, in parentheses, by the number of actual relationships. Considering both its own elements and those inherited by the imported ontologies, the core ontology for the history of computing currently includes 267 classes and 540 relationships.

Population of the CHIPS&BITS knowledge graph has already started, based on the current version of the ontology. Concerning the specific domain, the catalogues of a private collection and of the Museum of Computer Science at UniBA are being used as sources of knowledge, entered partly manually and partly using a batch upload procedure from their catalogues. These two collections were chosen because their owners and maintainers share our holistic view, according to which a seamless integration among hardware, software, documentation and intangible aspects. The private collection includes about 500 units between devices and rare components, plus hundreds of books, magazine issues, videos, software packages and other memorabilia. The Museum of Computer Science has on show a mainframe, several minicomputers and terminals, many personal computers and workstations, plus many items in the deposits and a huge collection of several thousand historical books, magazines and software, a large part of which is still to be catalogued. Moreover, the catalogues of 3 libraries specifically oriented toward the computing and related knowledge, for a total of 16000 records, were uploaded.

Domain	Ontology	Classes + Subclasses	Attributes	Relationships	Attributes
Top-level	*general*	13+128	107	33 (157)	43
Galleries, Libraries, Archives & Museums	*ifla*	12+60	128	23 (77)	22
	glam	7+45	88	32 (143)	26
Education	*lom*	12+43	88	14 (22)	11
	oerschema	6+14	15	11 (31)	2
	alocom	5+35	6	5 (8)	0
	intelleo	33+102	156	33 (101)	6
	education	4+16	6	17 (35)	26
Tourism & Traditions	*food*	8+21	13	13 (34)	3
	tourism	4+118	29	7 (104)	12
Core	*computing*	7+101	66	22 (81)	33

Table 1: Statistics on the elements in the ontologies involved in CHIPS&BITS

Other instances uploaded in the knowledge base concern the contextual part, expressed by the general ontology. For example, instances of class **Place** have being uploaded from an atlas, and instances of class Event have been manually entered by hobbyists and enthusiasts. To provide a wide range of labelling items, classes **Category** and **Word** were populated from WordNet [12][13], together with the associated relationship instances, from the standard part of the Dewey Decimal Classification (DDC) system [14], and from the ACM Computing Taxonomy and IEEE taxonomy. These instances can be linked to individuals of other classes (e.g., **Document**, **Person**, **Place**) and used as semantic or lexical tags (respectively) to express information about them (e.g., 'Alan Turing' might be linked with 'Computer Science', 'World War II', etc.). This will be another important way to contextualize the knowledge items and support personalization of behavior and outcomes.

After uploading these resources, a group of enthusiasts and hobbyists started creating connections between the different items (e.g., books or documents and devices or people, categories or words and books or documents, etc.).

The current content of the CHIPS&BITS knowledge base consists of 347295 entity instances, of which 1858 related to the computing domain and 4822 from the libraries/archives/museum domain, and 641293 relationship instances. The education and tourism sections were started recently, so they include just a few hundred instances. Conversely, the general ontology contributes with a few hundred thousand instances, thanks to the batch ingestion of various resources: the atlas for places, WordNet for words and concepts, and some subject taxonomies.

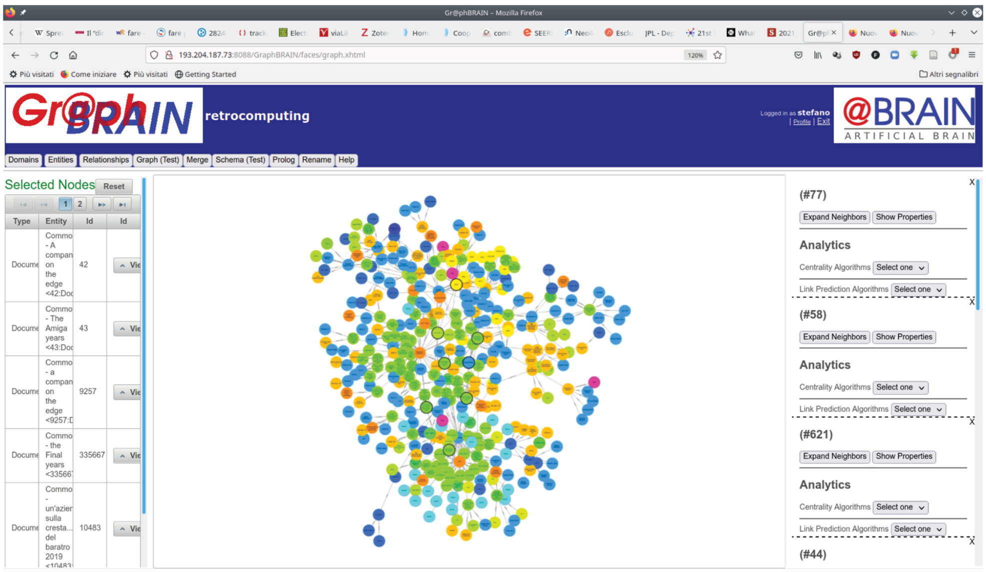

Figure 4: A sample subgraph from the computing knowledge graph

Let us now report some use cases we ran on the knowledge base. Figure 4 shows a subgraph extracted from the overall knowledge based on starting nodes that are relevant to the user's information need (shown with thicker borders in the graph and listed on the left-hand side of the interface). The subgraph extraction algorithm is aware of the user's profile and uses it to drive the extraction. Most of the functions used in the following use cases rely on a preliminary subgraph extraction like this, which is the starting point for carrying out the different tasks described.

As a preliminary indication of system behavior, we analyzed the log files produced during several batch ingestions of existing resources (city names, libraries and archives, museums, taxonomies, and similar datasets). For each batch, GraphBRAIN's ontology-driven validation checked whether the records being imported conformed to the corresponding ontologies with respect to mandatory fields; in these runs no hard validation errors were reported, and the only warnings concerned possible duplicate nodes or relationships, which were then inspected and, where appropriate, merged by curators. Although these checks currently focus on required fields and do not yet cover more sophisticated consistency or completeness constraints, they suggest that the data model and ingestion pipeline can accommodate heterogeneous legacy catalogues. From a user-experience perspective, curators who interacted with the system particularly appreciated (i) the usefulness of cross-domain links between objects, institutions and people, which make implicit relations explicit, and (ii) the extraction of personalized views of the knowledge graph, in the form of subgraphs built by expanding user-defined relevant nodes. The latter functionality is currently implemented as an experimental Prolog-based algorithm that has been tested off-line on CHIPS&BITS data but is not yet fully integrated into the production interface; nonetheless, initial trials indicate that it produces meaningful and manageable subgraphs that can support both exploration and curatorial work.

Use Case 1

User *X* is a collector, with broad interests but a special focus on personal computing and specifically the Commodore brand. He misses a few models from Commodore, and the system knows that another user, *Y*, owns two units of one of these models. The system prompts *Y* in case he would like to sell one of his units to *X*, or exchange it with one of *X*'s units. He accepts, and the system establishes a contact between *X* and *Y*, still hiding their real identities for privacy reasons, until they decide they can disclose their identities to each other.

Use Case 2

Institution *X* owns a unit of a computer model, but knows almost nothing about it. The system extracts relevant knowledge in the form of a subgraph centered in that model, and returns a natural language description of the model, its history and relevance, and its connections to other models, hardware and software. The system also retrieves other users who own(ed) or use(d) that model, so that they may be contacted for more information. Since the unit is not working, the system also provides information on how to fix it, and points to users who can repair or restore it. It also determines the value of the unit based on previous transactions, and retrieves users potentially interested in purchasing the unit in case the institution is not interested in it.

Use Case 3

Researcher *X*, interested in peripherals, comes to know about a company he was not aware of, that had a role in the history of computing and is relevant to his research. Actually, the system itself pointed *X* to that company, based on *X*'s interests. The system delivers a report on that company, describing its history, products and achievements, the people who worked for it and their roles and tasks, all from the specific perspective of peripherals (albeit the company may also have produced computers and software).

Use Case 4

Association *X* is organizing an exhibition about a pioneer of computer science and engineering, *Y*. The system finds the set of devices and software he developed, the companies he worked for, his history and some anecdotes that might make the exhibition more interesting. It also points *X* to users who own and might rent units of the devices and software developed by *Y*, to documents, audios and videos involving *Y* (with information on where they deal with *Y*), to other relevant memorabilia, and to previous events about *Y*.

Use Case 5

User *X* is interested in the Olivetti brand, and specifically in the firsts that it marked in the computing history. The system identifies the models developed by Olivetti which are landmarks, provides an explanation for the choice, identifies museums or exhibitions where units of these models are on show, and provides a description of touristic points of interests, typical food and folklore of the regions where such units are hosted, taking into account *X*'s interests and preferences.

Related Work

Work on Cultural Heritage knowledge representation is more oriented toward data schemas and vocabularies rather than ontologies. However, there are some noteworthy initiatives also about the development of ontologies, and associated knowledge graphs, for CH.

A relevant initiative in this direction is supported by the Italian Ministry of Culture: ArCo (Architecture of Knowledge)[4], is an ontology for, and a knowledge graph of, Italian CH. It currently reuses, and is aligned to standard CH-related schemes and ontologies: CIDOC-CRM, EDM, Cultural-ON, and OntoPiA. The resulting knowledge graph includes, and provides as LOD, data from records of the General Catalogue of Cultural Heritage, a database of Italian cultural heritage entities. Like ArCo, CHIPS&BITS organizes several ontologies in a network for modeling different kinds of cultural properties and their corresponding catalog records. Some branches of knowledge described by ArCo overlap with those of CHIPS&BITS (e.g., Archive, Catalogue & catalogue records, Cultural events & exhibitions). In principle, any part of ArCo

4 http://wit.istc.cnr.it/arco/?lang=en

(and of the ontologies it includes) can be reused in CHIPS&BITS, as well as content might be included from many resources such as Geonames, Wikidata, VIAF, etc. Describing in detail what is or might be reused from these and other existing resources in CHIPS&BITS is out of the scope of this paper, where the aim is just to say that this can be done and that interoperability with them is supported.

Related specifically to the computing domain, [15] applied knowledge-based approaches to an open dataset with persons and relationships extracted from the official biography of Steve Jobs and the 1999 film Pirates of Silicon Valley. CHIPS&BITS aims at being much wider and more comprehensive than this initiative, and at having impact on practical aspects of the discipline.

The Computer History Museum (CHM)[5] has developed an ontology to organize and represent knowledge of computer history. This ontology helps structure information about artifacts, people, events, and technologies in a machine-readable format, facilitating research, exhibitions, and digital archives. It should be noted that the CHM ontology is not freely available on the Web, but some of its parts related to ontological structure are detectable from the API made available for those who want to collaborate in its development[6]. However, this is contrary to the very definition of ontologies as "shareable" resources, and prevents any kind of reuse and interoperability.

At the European level there is an ongoing effort to strengthen tools related to cultural heritage. The Cultural Heritage Cloud[7] is a European Union initiative to create a digital infrastructure that will connect cultural heritage institutions and professionals across the EU. In particular in [16] are reported the characteristics that the cloud infrastructure designed to support the broad spectrum of applications and uses of cultural heritage data should exhibit in order to: 1) fully support the relevant types of data; 2) make the content FAIR (Findable, Accessible, Interoperable, and Re-usable); and 3) integrate the appropriate approaches and technologies needed to interact with cultural heritage data. CHIPS&BITS is compliant with these indications, and although it was not born from a CLOUD perspective, it could be easily transportable as it uses some of their core technologies.

From the technological viewpoint, some works exist that analyze the possibilities for cooperation between ontologies and graph DBs. [17] discusses technical issues that might limit the impact of symbolic Knowledge Representation on the Knowledge Graph area, and summarizes some developments towards addressing them in various logics. Most works tried to merge research on RDF and LPG knowledge representations, but always giving the RDF perspective priority and predominance. GraphBRAIN was the first to push for a native LPG-oriented approach [4]. After

5 https://computerhistory.org/

6 https://dev.chmlab.org/

7 https://research-and-innovation.ec.europa.eu/research-area/social-sciences-and-humanities/cultural-heritage-and-cultural-and-creative-industries-ccis/cultural-heritage-cloud_en

the publication of GraphBRAIN, other initiatives investigated the possibility of developing suitable schemas for LPGs specifically [18].

Conclusions and Future Work

Computing is starting to be considered part of Cultural Heritage, not only a means, due to the rapid obsolescence and advancements in technology calling for preservation. Luckily, we still have the opportunity to get knowledge from the pioneers who are still alive. The CHIPS&BITS project's aim is to collect, interrelate, preserve and deliver knowledge on the history of computing. In this sense, CHIPS&BITS can be situated within a broader line of digital humanities research that adopts semantic web and knowledge-graph technologies to preserve and explore complex cultural objects. Examples in this are the migration of a rich, domain-specific repertory into a Linked Open Data knowledge base as a way to safeguard specialised scholarly knowledge, while enabling new forms of access and reuse [19]. Likewise, in [20], experiments in narrative knowledge graphs show how ontologies and visual interfaces support interpretive work on a literary corpus. In [21], the authors show a methodology for semantic knowledge extraction from archival finding aids which uses Linked Open Data and language technologies to move from document-centric to data-centric representations of historical collections. Together, these convergences suggest that the history of computing, archival sources, rzepertories and even literary narratives can all be treated as graph-structured, interoperable cultural knowledge, and that CHIPS&BITS contributes a new domain and a new set of use cases to this evolving ecosystem.

The peculiarities and complexity of the domain of CHIPS&BITS are such that all available data models are inadequate for it. We proposed a new data schema in the form of an ontology, that enables the use of advanced AI solutions, and might be the core of a new cataloguing standard. This domain is also generating significant economic flow (due to publishing and educational initiatives, collectors' exchanges, event organization and job opportunities): CHIPS&BITS aims at leveraging this in the perspective of the Web economy. CHIPS&BITS proposes flexible solutions to provide tailored behaviors for the variety of potential users and stakeholders involved, and guarantees privacy for sensitive information associated to the economic exploitation. Most of the proposed solutions are general and applicable to any kind of CH.

Currently, CHIPS&BITS lacks of a specific interface. The development so far focused on building the ontologies and filling the knowledge base according to them. The support for this comes from the GraphBRAIN framework and platform, that provides all the basic CRUD functionalities on the knowledge graph and also several advanced exploitation, information extraction and automated reasoning functions. In fact, the needs of the CHIPS&BITS project were also inspirational for further expanding the functionality, feature and representation of GraphBRAIN, such as the personalization and privacy management strategies described in this paper. Thanks to this extension, we are currently experimenting with the personalized extraction of knowledge. We are also about to complete an extension of the ontology with a specific section for archives, based on internationally adopted data models.

Additional current and future work relates to further refinement and extension of the ontologies, automated and manual population of the knowledge base, and development of tools that can serve the specific needs of the different kinds of CHIPS&BITS stakeholders. In particular, the design of a platform for providing professional support to enthusiasts and collectors (e.g., e-commerce and auction functions, matching of offers and requests for items and for competences, etc.) has been started. Also, for improving knowledge acquisition and fruition, an automated tool based on the Large Language Models technology is under development that can fill entity, relationship and attribute instances by extracting them from books and documents, and that can express the extracted knowledge in natural language.

Acknowledgements

This research was partially supported by projects CHANGES "Cultural Heritage Active Innovation for Sustainable Society" (PE00000020), Spoke 3 "Digital Libraries, Archives and Philology" and FAIR "Future AI Research" (PE00000013), spoke 6 "Symbiotic AI", funded by the Italian Ministry of University and Research NRRP initiatives under the NextGenerationEU program.

References

[1] N. Barbuti, G. D. Felice, A. D. Zanni, P. Russo, A. Valentini, Creating Digital Culture by co-creation of Digital Cultural Heritage: the Crowddreaming living lab method, Umanistica Digitale (2020) 19–34. URL: https://umanisticadigitale.unibo.it/article/view/9956. doi:10.6092/issn.2532-8816/ 9956, number: 9.

[2] S. Ferilli, D. Redavid, An ontology and a collaborative knowledge base for history of computing, in: 1st International Workshop on Open Data and Ontologies for Cultural Heritage (ODOCH-2019), at the 31st International Conference on Advanced Information Systems Engineering (CAiSE 2016), Rome, Italy, volume 2375 of Central Europe (CEUR) Workshop Proceedings, 2019, pp. 49–60.

[3] I. Robinson, J. Webber, E. Eifrem, Graph Databases, 2nd ed., O'Reilly Media, 2015.

[4] S. Ferilli, Integration strategy and tool between formal ontology and graph database technology, Electronics 10 (2021). URL: https://www.mdpi.com/2079-9292/10/21/2616.

[5] T. Heath, C. Bizer, Linked Data: Evolving the Web into a Global Data Space, Synthesis Lectures on the Semantic Web, Morgan & Claypool Publishers, 2011.

[6] S. Ferilli, E. Bernasconi, D. Di Pierro, D. Redavid, A holistic ontology for digital libraries – the ifla-compliant core, in: Proceedings of the 2nd Workshop on Artificial Intelligence for Cultural Heritage (IAI4CH@AIIA2023), co-located with the XXII International Conference of the Italian Association for Artificial

Intelligence (AI*IA 2023), volume 3536 of CEUR Workshop Proceedings, 2023, pp. 39–52.

[7] S. Ferilli, E. Bernasconi, D. D. Pierro, D. Redavid, L. Loop, A holistic ontology for digital libraries, in: Proceedings of the 20th Conference on Information and Research science Connecting to Digital and Library science (formerly the Italian Research Conference on Digital Libraries), Bressanone, Brixen, Italy, 22-23 February 2024, volume 3643 of CEUR Workshop Proceedings, CEUR-WS.org, 2024, pp. 74–86.

[8] S. Ferilli, B. De Carolis, P. Buono, N. Di Mauro, S. Angelastro, D. Redavid, Una piattaforma intelligente per la gestione integrata del settore turistico, in: Primo Convegno Nazionale CINI sull'Intelligenza Artificiale – Workshop on AI for Cultural Heritage, 2019, p. 2 pages. In Italian.

[9] F. Vannozzi, Catalogare il patrimonio scientifico e tecnologico: da sic a sts a pst, storia di un percorso (e di una collaborazione), in: G. Pratesi, F. Vannozzi (Eds.), I valori del museo. Politiche di indirizzo e strategie di gestione, Economia e management della cultura e della creatività, Franco Angeli Edizioni, 2015, pp. 98–101. URL: https://core.ac.uk/download/pdf/53672194.pdf, in Italian.

[10] Collins, A. and Loftus, E. (1975). A spreading-activation theory of semantic processing. Psychological Review, 82:407–428.

[11] Quillian, M. (1966). Semantic Memory. AFCRL (Series). Air Force Cambridge Research Laboratories, Office of Aerospace Research, United States Air Force.

[12] G. A. Miller, Wordnet: A lexical database for english, Communications of the ACM 38 (1995) 39–41.

[13] C. Fellbaum (Ed.), WordNet: An Electronic Lexical Database, MIT Press, 1998.

[14] M.Dewey,Aclassification and subject index for cataloguing and arranging the book sand pamphlets of a library, Amherst, Mass., 1876.

[15] G. Drakopoulos, A. Kanavos, P. Mylonas, S. Sioutas, D. Tsolis, Towards a framework for tensor ontologies over NEO4JRepresentations and operations, in: 8th International Conference on Information, Intelligence, Systems & Applications, IISA 2017, Larnaca, Cyprus, August 27-30, 2017, IEEE, 2017, pp. 1–6.

[16] European Commission: Directorate-General for Research and Innovation, Brunet, P., De Luca, L., Hyvönen, E., Joffres, A. et al., Report on a European collaborative cloud for cultural heritage – Ex – ante impact assessment, Publications Office of the European Union, 2022, https://data.europa.eu/doi/10.2777/64014

[17] M. Krötzsch, Ontologies for knowledge graphs?, in: A. Artale, B. Glimm, R. Kontchakov (Eds.), Proceedings of the 30th International Workshop on

Description Logics, Montpellier, France, July 18-21, 2017., volume 1879 of CEUR Workshop Proceedings, CEUR-WS.org, 2017.

[18] R. Angles, A. Bonifati, S. Dumbrava, G. Fletcher, A. Green, J. Hidders, B. Li, L. Libkin, V. Marsault, W. Martens, F. Murlak, S. Plantikow, O. Savkovic, M. Schmidt, J. Sequeda, S. Staworko, D. Tomaszuk, H. Voigt, D. Vrgoc, M. Wu, D. Zivkovic, Pg-schema: Schemas for property graphs, Proc. ACM Manag. Data 1 (2023).

[19] Bonora, Paolo, and Angelo Pompilio. 2022. "RePIM in LOD: Semantic Technologies to Manage, Preserve, and Disseminate Knowledge about Italian Secular Music and Lyric Poetry from the 16th-17th Centuries." Umanistica Digitale, no. 14: 71–90. https://doi.org/10.6092/issn.2532-8816/15568.

[20] Bruno, Enrica, Valentina Pasqual, and Francesca Tomasi. 2024. "Italo Calvino's 'Destini Incrociati'. An Experiment of Semantic Narrative Modelling and Visualisation." Umanistica Digitale, no. 17 (May): 47–69. https://doi.org/10.6092/issn.2532-8816/19013.

[21] Giagnolini, Lucia, Andrea Schimmenti, Paolo Bonora, and Francesca Tomasi. 2025. "Expliciting Contexts: Semantic Knowledge Extraction from Traditional Archival Descriptions." Umanistica Digitale, no. 20 (July): 115–44. https://doi.org/10.6092/issn.2532-8816/21229.