Expanding the Virtual Record Treasury of Ireland Knowledge Graph

Beyza Yaman [a,*], Alex Randles [a], Lucy McKenna [a], Lynn Kilgallon [b], Diego Rincon-Yanez [a,d], Neil Johnston [c], Peter Crooks [b] and Declan O'Sullivan [a]

[a] *ADAPT Centre, School of Computer Science and Statistics, Trinity College Dublin, Dublin, Ireland*
[b] *Department of History, Trinity College Dublin, Dublin, Ireland*
[c] *The National Archives, London, UK*
[d] *Vienna University of Economics and Business (WU Wien), Vienna, Austria*

Abstract.
The Virtual Record Treasury of Ireland (VRTI) is an online resource which digitally reconstructs the archival collections lost with the destruction of the Public Office of Ireland in 1922. This resource includes a distributed Knowledge Graph (KG), which employs Semantic Web principles and Linked Data to support discoverability of historical entities across multiple archival collections. CIDOC-CRM serves as the core ontology, extended with bespoke types to address domain-specific needs. VRTI-KG was deployed successfully in June 2022, at the end of the first phase of the VRTI's development. The second phase of VRTI-KG development witnessed challenges arising from expanding data sources, advances in technology, an expanding user base, and advanced user requirements. This article focuses on the technical solutions developed to overcome these challenges, which include: developing a robust URI structure; updating and expanding the ontology; uplifting authoritative geospatial data; and creating a map interface to facilitate public engagement. Finally, two bespoke user interfaces were created for the KG; the KG Editor is designed to enable domain experts to interact intuitively with the KG, while the KG Explorer allows public users to navigate seamlessly between historical data in VRTI-KG and the reconstructed archival records in the VRTI document database.

Keywords: Geospatial linked data, Cultural heritage, Manuscripts

1. Introduction

The Virtual Record Treasury of Ireland (VRTI) is an all-island and international legacy project developed as part of Ireland's Decade of Centenaries. Hosted by Trinity College Dublin, it represents the culmination of a seven-year research programme funded by the Government of Ireland under Project Ireland 2040, through the Department of Tourism, Culture, Arts, Gaeltacht, Sport and Media. The VRTI reconstructs digitally the archival collections which were destroyed in June 1922, during the opening days of the Irish Civil War, when the Public Record Office of Ireland (PROI) was consumed in a catastrophic fire and explosion. The PROI, situated within the Four Courts complex in Dublin, was established in 1867. As Ireland's state archive, it housed centuries of official records and state papers—some dating back over 700 years. In Easter 1922, escalating tensions between the National Army of the new Irish Free State and Anti-Treaty forces led to the occupation of the Four Courts. Open conflict erupted on 28 June, and on 30 June, an explosion destroyed the eastern wall of the Record Treasury. The resulting fire consumed nearly all the contents of the PROI. A century later, the Beyond 2022 project was launched to begin the virtual

*Corresponding author. E-mail: beyza.yaman@adaptcentre.ie.

reconstruction of this lost archive. This work culminated in the launch of the Virtual Record Treasury of Ireland as a digital resource on 30 June 2022, www.virtualtreasury.ie.

At its core, VRTI brings together historical resources from numerous archives to recover the content of the destroyed records. A relational database captures information about these documents—including digital images, searchable text, and associated metadata. A central component of the VRTI digital platform is the Knowledge Graph for Irish History (VRTI-KG). A Knowledge Graph (KG) was chosen as the most effective way to represent historical information, allowing data integration, interlinking entities, and reasoning across datasets. VRTI-KG models information contained in reconstructed sources—such as books, parchments, and other documents—rather than metadata about the documents contained in the virtual treasury. The KG currently captures knowledge about people and places from Ireland's past, drawn from both within and beyond the Virtual Treasury's reconstructed records. VRTI-KG also provides authoritative identifiers for people, places, events, offices, organisations, and other entities in the Irish historical record. Information in the KG has been sourced both from subject-matter experts (e.g. historians examining artefacts) and from pre-existing collections of structured data which have been transformed into RDF.

In [3], we previously outlined the foundational design and implementation work for VRTI-KG that took place in the period 2018 to 2022, leading to the successful launch of the VRTI website in June 2022. The first phase of KG development took place in parallel and focused on creating a structured yet flexible resource for subject-matter experts to record and curate historical data prior to transformation and ingestion by the project's knowledge engineers. To ensure that historians can contribute effectively, the project adopted a familiar and accessible technology: spreadsheets. The spreadsheet-based approach was designed to allow historians to contribute data without needing to interact directly with the technical underpinnings of the KG. This effort resulted in a suite of tailored spreadsheet schemas developed to capture structured data about entities such as people, offices, organisations, and places. For example, the People sheet enables historians to record biographical details, relationships, tenure, rank, and status. The schemas were developed through close collaboration between VRTI's computer scientists and humanities scholars. Designed with future extensibility in mind, the schemas were intended to support a wide variety of use cases and accommodate the complexities and diversities of Ireland's historical record. While designed for use within VRTI-KG, they are also intended to be adaptable so that other cultural heritage projects and users can both use and contribute data to the KG. In addition to generating new data for the Knowledge Graph, a key requirement was enabling historians to explore, review, and edit its contents.

Prior to the second phase of KG development (2022-2025), all aspects of phase one were reviewed, from the historian/computer scientist collaboration process through to the initial implementation of the KG and out of the box user interface related tools. As a result of this review, three main areas of new functional requirements emerged and how they have been addressed are outlined in this article. First, update to the VRTI ontology (expanding properties of people entities including linking them to key life events), and the introduction of VRTI vocabulary to allow for eras, and geospatial entities. See Section 2. The second new area of functional requirements relate to the need for improved search, query and visualisation of the entities in the KG (see section 4). It was apparent that interacting with the KG posed challenges for users without technical expertise. In addition, the spreadsheet-based data entry process, while accessible for historians, represented a significant workload and made it difficult to track changes and updates, particularly when correcting or refining existing information. Maintaining consistent editorial standards across contributions was an additional concern. Addressing these limitations became a key priority. As a result, a new requirement was agreed for two new interfaces to support non-knowledge engineer engagement with the KG. The first is an exploration interface designed for pubic end users, allowing seamless navigation between the VRTI's document collections and the KG. The second is an editing interface aimed at VRTI historians, enabling them to add, amend, and manage data directly within the KG. The third area for new functional requirements emerged from the desire to enrich the modeling of Places in the KG, to both support deeper historical analysis and to provide more engaging user geospatial based interaction through the new public user interface (see section 5). From the review a set of non-functional requirements also emerged. This resulted in actions related to making the infrastructure more evolvable and sustainable in the longer term, flexible interconnection of the KG with the VRTI manuscripts database and website, as well as URIs redesign. See section 3.

Finally, Section 6 provides a use-case from VRTI-KG; Section 7 discusses related work focusing on cultural heritage domain; and Section 8 concludes the paper by the lessons learnt.

2. Updates to the VRTI Ontology

The VRTI Knowledge Graph uses the VRTI Ontology and the VRTI Vocabulary (see footnotes below for links), both of which are grounded in the *CIDOC Conceptual Reference Model* (CIDOC-CRM). CIDOC-CRM is an international standard developed for modeling complex historical entities, events, and their interrelations [2]. As a digital humanities initiative, the VRTI project leverages CIDOC-CRM's event-centric approach to accurately represent the dynamic and layered nature of Irish history. Table 1 presents a summary of the total instances in VRTI-KG. In the last period we have made a number of simple house-keeping changes, but the more substantial changes are outlined in this section.

Table 1
Total number of instances in KG

Type	Count
People	10,653
Places	66,771
Triples	2.9M

2.1. Persistent ID for the VRTI Ontology and Vocabulary

To ensure the long-term stability and resolvability of its ontology, the VRTI project now employs `w3id.org` URIs as persistent identifiers [18]. `w3id.org` is a community-driven initiative that provides stable, HTTP-based identifiers that can be resolved by web clients and are interoperable with RDF, OWL, and other Semantic Web technologies. By registering a W3ID, the VRTI project guarantees that references to its vocabulary will remain valid and resolvable even if underlying hosting infrastructure changes, aligning with Linked Data and FAIR principles. The VRTI Ontology[1] and the VRTI Vocabulary[2] are defined within the dedicated namespaces:

```
@prefix vrti: <https://www.w3id.org/virtual-treasury/ontology#>
@prefix vrtivoc: <https://www.w3id.org/virtual-treasury/vocabulary#>
```

The compliance of the ontology with FAIR principles was checked through the use of the FOOPS tool[3] [4]. FOOPS is an online tool used to assess an ontology for design pitfalls such as missing annotations, incorrect relations, logical inconsistencies, or bad modeling practices. It defines different types of pitfalls, their properties, and how they relate to ontological elements, supporting automated evaluation and quality assurance of ontologies. The VRTI ontology was iteratively assessed using FOOPS and the resulting quality report was used to guide the resolution of detected quality issues. The latest version of the ontology was assessed using FOOPS and received a score of 81% (Figure 1), with it not achieving full compliance primarily due to the ontology not being listed in some public registries. It is planned that the ontology will be submitted to the relevant public registries to improve its compliance in future assessments.

Fig. 1. Results of FOOPS applied to VRTI Ontology

[1] https://www.w3id.org/virtual-treasury/ontology
[2] https://www.w3id.org/virtual-treasury/vocabulary
[3] https://github.com/oeg-upm/fair_ontologies

2.2. *Person Schema Expansion*

The VRTI Person Schema is a structured metadata model designed to standardise the representation of individuals within the KG. The old VRTI Person Schema consisted of 43 fields, while the updated schema expands this to 53 fields. This increase reflects a broader and more nuanced capacity to support:

- richer name variant handling
- more granular career information
- richer descriptions of identity (e.g. nationality, religion) and social status
- formal links to external authority datasets
- improved tracking of contributor and provenance information

Fig. 2. VRTI Person Schema

The 53 fields are organised into categories such as names and aliases, biographical dates and places, familial relationships, career and references to external data sources. The schema supports a rich variety of name forms

including normalised and variant spellings of forenames/surnames, aliases, patronymics/matronymics, honorifics, and sobriquets. This design facilitates name disambiguation and cross-referencing in historical records, which often contain inconsistent or multilingual naming conventions.

Beyond names, the schema also accommodates key life events (birth, death, marriage), relationships to people, places, offices and organisations, and indicators for gender, religion, nationality, and social status—allowing for nuanced prosopographical analysis. The inclusion of links to external authorities such as Wikidata, FAST[4] (Faceted Application of Subject Terminology), and the Dictionary of Irish Biography demonstrates a commitment to Linked Data principles, supporting interoperability and data enrichment. Optional fields for contributor attributions, citations and images further enhance transparency and context.

Figure 2 provides a visual overview of the structure of the updated VRTI Person Schema. The diagram groups fields into thematic categories to reflect how person records are constructed and curated. Fields highlighted in yellow indicate those that were added or substantively extended in the current phase of the project, including more detailed descriptions for aliases, enhanced temporal modelling (i.e. era dating), expanded use of controlled vocabularies for career data, and additional links to external authority datasets.

2.3. Introduction of the VRTI Vocabulary

To accommodate the specific requirements of the VRTI, new concepts were introduced as an extension to the core CIDOC-CRM ontology. This VRTI Vocabulary is used to classify entities such as places, organisations, offices, and roles, as well as to control a range of descriptive attributes. It is implemented using a combination of `E55 Type` from CIDOC-CRM and the Simple Knowledge Organisation System (SKOS) [23] concepts, which are treated as functionally equivalent mechanisms for classification and controlled description within the project. This approach supports both CRM-aligned semantic modelling and SKOS-based hierarchical organisation and navigation.

The vocabulary is currently used to classify entities such as people (gender), places (place types), and dates (era). The vocabulary is designed to be extensible, and future phases of the project will apply the same combined E55/SKOS pattern to additional domains, including the modelling of event types, and for descriptive attributes such as occupation, rank, religion and national identity.

Era

The VRTI Vocabulary models historical *eras* as SKOS concepts and E55 Types. This allows for the ability to group and search for entities by era (see Figure 3). These controlled era terms are also used as part of the URI for places, offices and organisations, (see section 5.3) and include:

- `present-day`
  Refers to the contemporary context, representing entities as they exist today.
- `medieval-400-1499`
  Encompasses the early and high medieval periods, approximately from the 5th to the late 15th century.
- `early-modern-1500-1749`
  Covers the period from the early 16th century to the mid-18th century.
- `modern-1750-1921`
  Represents the late modern period leading up to Irish independence.
- `modern-post-1922`
  Applies to entities operating in the context of the modern Irish state, post-establishment of the Free State in 1922.

[4] https://www.oclc.org/research/areas/data-science/fast.html

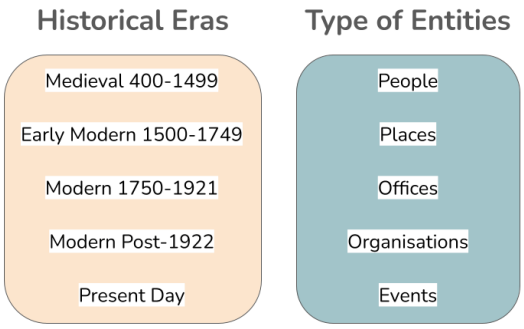

Fig. 3. Data Organisation in VRTI-KG

Place Types

The vocabulary also defines a comprehensive typology of geographic entities across time, including `province`, `county`, `barony`, `parish`, `townland`, `city` and `town`. The vocabulary also includes historical variations such as `present-day-county`, `modern-1750-1921-townland`, `early-modern-1500-1749-barony`, and other temporally-scoped place types. This enables the graph to capture both spatial and historical nuance, ensuring that places are accurately contextualised within their relevant time-period.

3. Updates to Technical Infrastructure

3.1. KG Technical Updates

The original (prior to July 2025) SPARQL endpoint of VRTI-KG resided in the Blazegraph triplestore. However, due to several reasons [24] we decided to migrate to Virtuoso, namely: i) Following its acquisition by Amazon Neptune, Blazegraph is no longer maintained as an open-source tool. ii) Blazegraph lacks full support for the GeoSPARQL extension, which is essential in support of the project's visualisation requirements that depend on executing GeoSPARQL queries; iii) we found that Blazegraph frequently encounters timeouts for certain query types and experiences unexpected crashes, leading to system instability. Virtuoso[5] outperforms many other options in handling complex SPARQL queries demonstrating good scalability, and effectively managing datasets exceeding 25 billion triples, which is crucial for accommodating VRTI's rapid growth. Its support for high-frequency read and write operations ensures that the query service remains responsive and up-to-date with the latest data. Additionally, Virtuoso offers comprehensive SPARQL 1.1 capabilities with GeoSPARQL, which are essential for complex data retrieval across diverse datasets. Thus, Virtuoso's performance, scalability, and feature set made it the most suitable choice for VRTI KG updated requirements.

In preparation for July 2025 launch of the new website, two different virtual machines are being used to provide a secure environment for the production and deployment of VRTI systems, as well as providing a seamless infrastructure for public users. The first virtual machine, known as the production server, hosts web applications utilizing Apache, providing access to the Virtuoso triplestore. Users can explore and visualise data through various applications, including the Virtuoso query interface, VRTI-KG explorer and Ontodia (see [24]). Figure 4 illustrates the technical structure of the KG in the VRTI project. The developer deploys the applications and triplestore on the web server, and users can access KG exploration tools via hyperlinks on the VRTI website. Once VRTI KG is relaunched in July 2025, users will interact with visualisation tools directly through the website.

The second virtual machine, referred to as the development server (Figure 5), is designated for hosting tools currently under development. To ensure data security on the production server, this separate server has restricted

[5] https://virtuoso.virtualtreasury.ie/sparql/

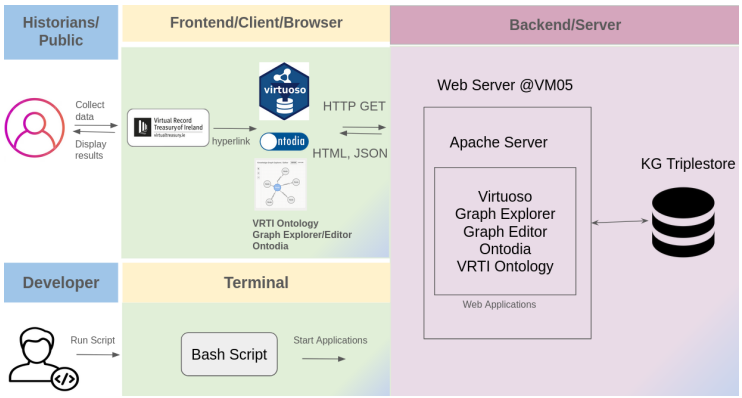

Fig. 4. Production Server Diagram

access, limited to VRTI team members. It hosts versions of VRTI-KG explorer and editor, with functionality that is being actively developed. Authenticated users can access, modify, and add data to the development server's triplestore as soon as the developer deploys updates for testing new features.

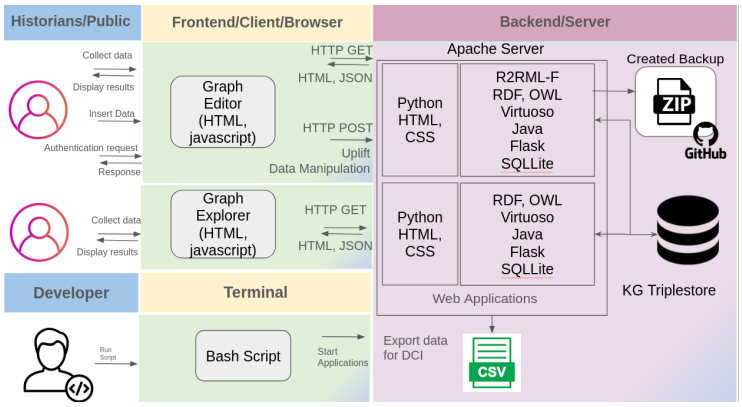

Fig. 5. Development Server Diagram

3.2. URI Redesign

The URI scheme for entities in the VRTI KG was redesigned to improve clarity, scalability, and semantic expressiveness. Earlier URIs followed a simpler structure, such as https://kb.virtualtreasury.ie/person/{person-name}_{century} and https://kb.virtualtreasury.ie/place/{place-type}/{place-name}, where the `kb` subdomain denoted the knowledge base. In the updated model, the namespace transitioned to `kg.virtualtreasury.ie` to reflect the shift from knowledge base to knowledge graph. The new URI pattern introduces two key changes: (1) the inclusion of a structured `era` segment to support temporal context, and (2) the addition of a betanumeric identifier to ensure uniqueness and avoid collisions among similarly named entities.

Temporal Context

A distinctive feature of the VRTI URI re-design is its inclusion of a temporal context for entities that evolve or are situated historically. The `era` segment in the URI provides explicit chronological anchoring for entities such as places, offices, and organisations. The following values are used for the `era` component:

- `present-day`

- `medieval-400-1499`
- `early-modern-1500-1749`
- `modern-1750-1921`
- `modern-post-1922`

By integrating these standardised temporal labels into URIs, the VRTI KG provides a framework for representing the historical evolution of institutions and places across time, and allows for the filtering of search results by historical period.

Betanumeric Identifiers

Each URI in the VRTI KG includes a final unique identifier to distinguish entities, such as people, places, offices or organisations, with identical or similar names. These identifiers are designed following principles inspired by the Archival Resource Key (ARK) identifier scheme [1]. The identifiers use a controlled betanumeric subset composed only of digits and consonants, excluding the letter `l` (to avoid confusion with the digit `1`) and all vowels (to prevent the accidental generation of words). In addition, all identifiers are prefixed with a `v1` *shoulder*, which serves as a short identifier extension. This structure facilitates the division of identifiers across future autonomous projects or divisions associated with the VRTI project, thus enabling independent identifier assignment without the risk of duplicates.

The VRTI KG adopts a mixed URI design strategy that combines human-readable components with these betanumeric identifiers. This approach reflects the dual requirements of the project: supporting scholarly usability while ensuring long-term identifier stability and disambiguation.

Descriptive URI components are included primarily to support human readability, as URIs are surfaced directly within the Virtual Record Treasury interface and are routinely handled by historians during data curation. Readable URI fragments aid recognition and reduce errors during manual reconciliation tasks. At the same time, the opaque betanumeric identifiers are employed to guarantee uniqueness, enable differentiation between entities sharing identical or near-identical names, and prevent excessive URI length where descriptive strings would become unwieldy.

Updated URI Design Examples

The base namespace for all VRTI KG URIs is:

```
https://kg.virtualtreasury.ie/
```

Following the base, the URI path reflects the **entity type** (e.g., `person`, `place`, `office`, `organisation`) and incorporates relevant contextual elements such as temporal scope (era), name, type, and a final betanumeric identifier to ensure uniqueness.

Person URIs

Person URIs encode the individual's name and the century in which they were active, followed by a unique identifier:

```
https://kg.virtualtreasury.ie/person/{person-name}_{century}/{unique-id}
Example: https://kg.virtualtreasury.ie/person/Butler_James_c18/v1j4w1m
```

Place URIs

Place URIs are organised hierarchically and include the historical era, the type of place e.g., `town`, `city`, `townland`, `parish`, `barony`, `county`, the name of the place, and a unique identifier:

```
https://kg.virtualtreasury.ie/place/{era}/{place-type}/{place-name}/{unique-id}
Example: https://kg.virtualtreasury.ie/place/present-day/town/Ennis/v15qqn1
```

In the case of place entities, the VRTI ontology supports language-tagged literals (e.g. @en, @ga) for labels and descriptive fields in order to accommodate multilingual expression and historical name variations. Plain string literals are retained where language tagging is not meaningful or where the source material does not warrant linguistic interpretation.

Office URIs

Office URIs reflect administrative, institutional or governmental offices tied to a specific historical context. The path includes the era, the country in which the office was located, the office name, and a unique identifier:

```
https://kg.virtualtreasury.ie/office/{era}/{country}{office-name}/{unique-id}
```

Organisation URIs

Organisation URIs follow a structure similar to office URIs and represent historical organisations such as religious institutions, civic bodies, or government departments:

```
https://kg.virtualtreasury.ie/organisation/{era}/{country}{organisation-name}/{unique-id}
```

The inclusion of country-level identifiers in the URIs of offices and organisations is intended to ensure unambiguous distinction between entities with identical or closely related titles operating in different jurisdictions. Many Irish offices and institutions share names with contemporaneous British counterparts, particularly in administrative and legal contexts. Embedding a country-level qualifier in these URIs provides a pragmatic disambiguation mechanism that supports uniqueness, interoperability, and reuse of identifiers beyond the immediate project context.

In contrast, person URIs do not encode country or place information, as personal identity within the VRTI KG is not necessarily defined by a single national or geographic attribute. Although the majority of individuals represented are Irish, place of birth, residence, and national affiliation are explicitly modelled as properties within the person schema rather than being embedded in the identifier itself. This design choice avoids over-specifying identity at the URI level and better accommodates historical complexity, including migration, multiple residences, and changing political boundaries.

Similarly, place URIs do not include a country-level component because the VRTI KG is intentionally scoped to Irish places only. Within this bounded geographic domain, country-level qualification would be redundant. Place disambiguation is instead achieved through hierarchical spatial modelling (e.g. townland, parish, barony, county).

Overall, the URI design for entities in the VRTI KG balances machine-processability with human readability, and reflects temporal and contextual dimensions where appropriate. Each URI serves as a persistent, dereferenceable identifier, ensuring that entities within the graph - such as people, places, offices and organisations - are uniquely addressable over the web.

Redirect Policy

A redirect policy was designed to handle requests to each URI in VRTI-KG, which redirects to a relevant summary page or restricted search results page in VRTI-KG explorer (see Section 4), depending on the whether the URI is applicable to only one resource (e.g. a person URI) or multiple resources (e.g. a surname name of a person). The redirection facilitates the discovery of related information in the KG using the interface of the explorer. For instance, a URI (e.g. https://kg.virtualtreasury.ie/birth/Talbot_Richard_c17/v1xf6p1) for the birth event of a person is redirected to the summary page of the selected person, which details the birth and death information, among other things. Similar redirection occurs when a URI for a place, such as a centroid or boundary is selected which redirects to the

summary page of the place. Restricted search result redirects occur when a URI, which can be related to multiple resources is selected. For instance, a URI representing the surname (e.g. https://kg.virtualtreasury.ie/normalized-appellation-surname/Talbot) or forename is selected. The redirection results in a search results page being displayed with all of the people in the KG with the selected name shown.

3.3. DB and KG Connection

The VRTI-DB contains digitised versions of the historical records and their metadata in a relational database, such as manuscripts and census records from Irish history. Entities such as people and places identified in VRTI-KG that appear in a record, are referenced within the metadata of the record stored in the DB. To facilitate this, a process was provided that allows historians to link entities from the KG with related historical records in the DB. For instance, the person who sent a historic letter, or where a person is mentioned in a legal document or a will. The process involves the historian searching for the URI of the specified entity using the KG explorer and then identifying for the unique reference code of the associated record. Then, the historian selects the role (e.g. creator, editor, sender, publisher, mentioned in) of the KG entity with respect to the record. Once the KG URIs (and the identified roles) are ingested into the metadata record in the database, the links to the records are made available via VRTI-KG explorer through a restful API endpoint that allows new links to be shown on the page of the KG entity (whether person or place etc), thus enabling straightforward traversing of related historical records. Similarly, when a user is browsing a historical record on the database, they are provided the opportunity to click on an included KG reference and explore the entity through the KG explorer.

As shown in Figure 6, the system connects the VRTI-DB to VRTI-KG through a structured ingestion and URI management process. This process facilitates the creation of links between historical records and their associated KG URIs, enabling seamless exploration by a user of both manuscripts and relevant entities in the KG and vice versa.

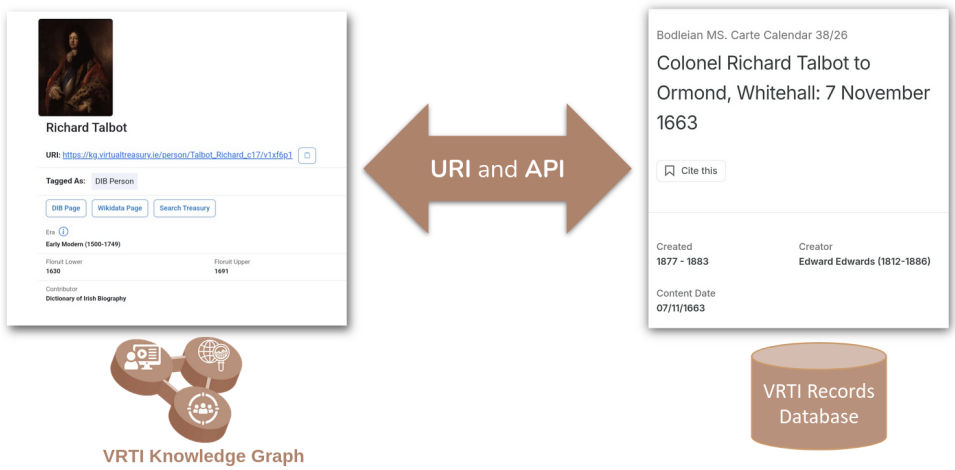

Fig. 6. Database to Knowledge Graph connection architecture.

4. New Requirement: Search, Query and Visualise Cultural Heritage Data on the Semantic Web

The tools, discussion and experimental results presented in this section have been previously presented in previous papers [19, 21], nevertheless, they are included here for completeness of the narrative. When initially launched in 2022, users could only interact with VRTI-KG through the use of out of the box customised tools such as Ontodia, LodView etc. In response to user feedback, it was decided to develop specific tools for both end users and

historians to interact with VRTI-KG. The KG Explorer [21] was developed to enable non-technical users to search and interact with resources in VRTI-KG. While its design draws inspiration from Sampo-UI's multi-faceted search approach [10], it introduces several enhancements, including natural language information and a data-driven interface configuration. These improvements were informed by input from a focus group of computer science and digital humanities researchers within the VRTI. The Explorer is configured through a set of JSON files, which streamline synchronisation between the KG's underlying data model and the user interface components. These files use key-value pairs to define various features, such as SPARQL queries triggered upon search execution and advanced filtering options to refine search results.

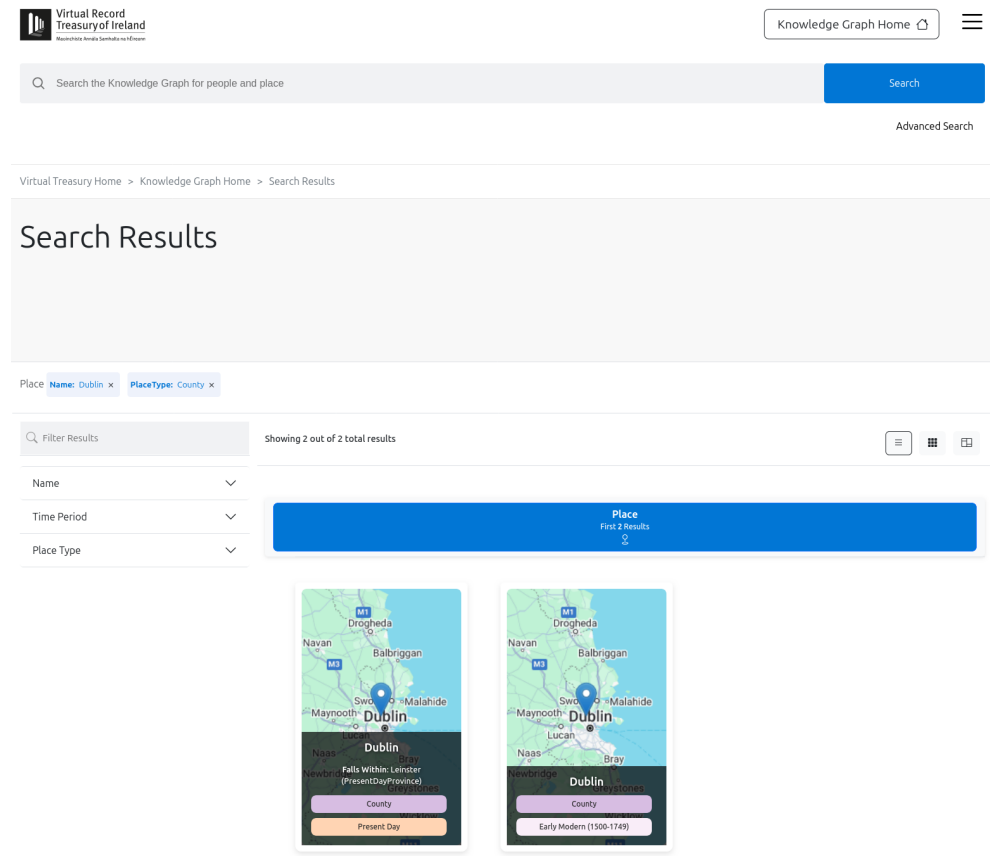

Fig. 7. Sample search results presented on the KG explorer for places with name "Dublin"

Figure 7 presents search results on the explorer for the county of Dublin (see Listing 2). The screenshot presents the main search bar which suggests resource labels from the KG which match the user's input. Advanced search filters can be applied to the initial search term to narrow down the search results. The type of place and time period related to each place in the search results are shown to allow users to easily differentiate between each result. The user is redirected to a summary page[6] (Figure 8) describing the place, including geospatial visualisations and listings of people who were born or died at the respective place. In addition, black markers are shown on the map which represent nearby places that are the same type (e..g. townland, counties) and can be clicked to explore their respective summary page.

Similar summary pages[7] (see Figure 9) exist for people, which describe their attributes in tabular and natural language format, related links (such as Wikidata page), other people in the KG who share the same surname and a

[6]https://kg.virtualtreasury.ie/place/present-day/county/Dublin/v1s4kv2
[7]https://kg.virtualtreasury.ie/person/Talbot_Richard_c17/v1xf6p1

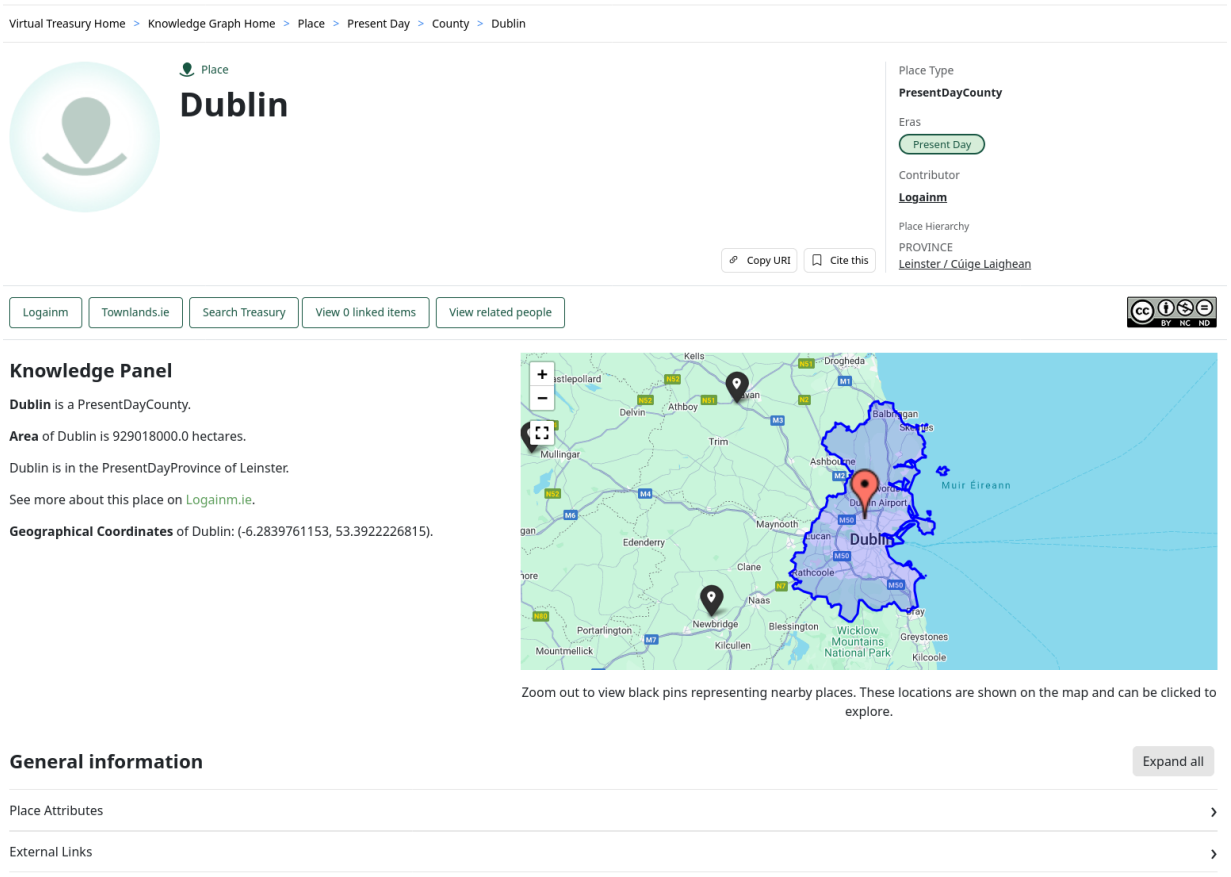

Fig. 8. Summary page for present day county Dublin

listing of records in the treasury related to the person. In addition, an extract of the *Dictionary of Irish Biography* (DIB)[8] entry for the person is shown as part of the summary.

The explorer enables users to search for places in VRTI-KG using intuitive geospatial visualisations. The interactive map allows navigation through places, displaying tooltips with key information upon hovering. Clicking on a place redirects users to a detailed summary page, providing further information about the selected place.

A user evaluation [21] was conducted on the explorer with 20 participants of unspecified background knowledge with varying levels of semantic web experience. A task sheet was provided to the participants with a list of tasks which mimic expected user interaction, which included finding relevant information on people and places and navigating different visualisations. The participants completed the tasks asynchronously using a provided online link to the explorer. Three key metrics were measured in the experiment: satisfaction of the user interaction, understanding of information provided on the interface and efficiency of task completion. Satisfaction was measured using the Post Study Usability Questionnaire (PSSUQ) [16] which contains a Likert scale 1 (best) - 7 (worst) to rate various aspects of a software system, such as interface and information quality. Understanding was measured using a bespoke questionnaire which asked participants to provide information displayed on the interface, such as the birth date and place of a person. All of the PSSUQ sub-scales scored between 5% and 35% better than their acceptable research threshold [16]. Efficiency was measured by asking participants to self-report time of completion of the tasks. 19 of the 20 participants (95%) answered at least 9 out of 11 questions in the understanding questionnaire correctly. The

[8] https://www.dib.ie

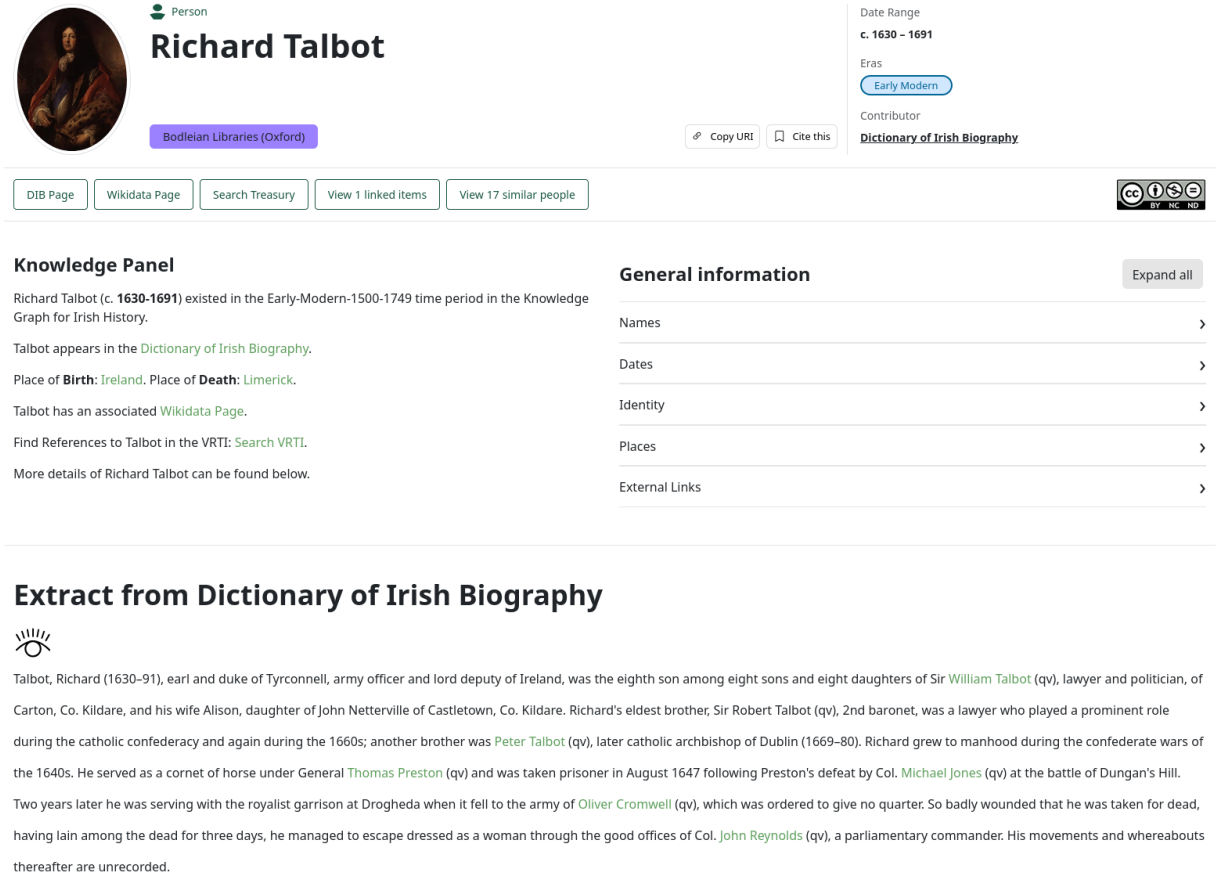

Fig. 9. Summary page for a person entity

efficiency results showed that not all participants completed the tasks equally efficiency with a standard deviation of 6.8 minutes, which could have been as a result of some participants having previous semantic web knowledge.

The KG explorer was released to the public at the end of June 2025. The usage of the users was captured by Google analytics. Since being released, the explorer has received over 11,000 unique visits from 80 different countries with over 60% of visitors based outside of Ireland. Over 60,000 pages have been viewed and over 12,000 keyword searches have been conducted. These numbers indicate active and meaningful engagement from end users.

The KG Editor [19] is a tool used internally in the project by the historians in the VRTI team to edit existing resources in the KG, such as people and places, without the need for them to be familiar with RDF, SPARQL, or SHACL. In addition, the editor facilitates the creation of new resources in the KG, which can then be linked with existing resources using the editor. The use of a tool to mediate engagement with VRTI-KG not only provides direct editing of VRTI-KG for the historians, but also ensures that any edits are undertaken in a valid and consistent manner, ensuring quality of data in the KG is maintained. The interface provides bespoke forms to change existing data or create new data by suggesting through dropdowns relevant resources in the KG, such as a place for the birth place of a person. The data input into the forms is then automatically translated into SPARQL [6] queries that are executed on the KG. SHACL [11] is used to validate the quality of changes before committing to ensure that specified constraints are satisfied and overall data quality did not decrease. The editor has a similar configuration framework to the KG Explorer, allowing it to be configured and synchronized easily with changes in the KG data model.

An initial user evaluation [19] was also conducted on the editor, however, the participants consisted of 9 historians. The three same metrics were used, in addition, to accuracy which was used to compare each participants resulting graph pattern with the expected pattern for each edit. Similar scores for satisfaction, understanding and efficiency were observed. The participants completed 8 edit tasks, resulting in a mean score of 7 (77%) correct edits. Several of the incorrect edits were a result of the participants committing the new resource to the KG without including all of the requested attributes. The subsequent commit was blocked as the URI had already been created. This issue was resolved by allowing subsequent commits under exceptional circumstances, after informing the users of potential duplicates that exist in the KG. Other incorrect edits were caused as a result of the participant not finding the correct tab that contains the specified attribute. This issue was resolved by including tooltip information on the tabs to inform users on the attributes available to edit in each of them.

Similar refinements were suggested by participants in both Explorer and Editor usability experiments, such as including additional tooltip information to describe technical terms and softening technical terms. The refinements identified from the findings of the user evaluation were implemented into the latest version of the tools and are hoped to improve user experience. It is planned to conduct a subsequent evaluation on the refined version of the editor.

5. New Requirement: Geospatial Linked Data for Irish History

As outlined in [3] one of the achievements in the first period of VRTI-KG development was the modeling of Place and provision of a process by which historians could connect Persons, Organisations and Offices with their associated Places. At the time, the representation of Place and creation of instances within the KG were sufficient for proof of concept. However as more data sources started to emerge for incorporation into VRTI, and the need for a more natural interaction for end users with VRTI-KG (and place being one of the natural jumping off points for user interaction), the geospatial model and dataset was revisited in order to provide for a more comprehensive and authoritative geospatial data stack for places across historical eras. In addition, we want to make this geospatial data available as Linked Data to support other historical and diverse domain applications. It is important to follow the best practices for data quality when doing that as it is essential and ensures that information is accurate, consistent, and fit for its intended use, which in turn enables reliable integration and effective data governance across heterogeneous data sources [25].

5.1. Underpinning Geospatial Datasets

Geospatial data uplift refers to enhancing geographic datasets by transforming them into rich and more detailed semantic geospatial resources. In the scope of this work, we leveraged the geospatial data including information about the whole island of Ireland (North and South) downloaded from townlands.ie. Townlands.ie is a website that provides a detailed resource for exploring Irish townlands, the smallest administrative land divisions in Ireland. The data in the platform is derived from OpenStreetMap (OSM) data and aims to document and showcase the geographical boundaries of Ireland's townlands, baronies, civil parishes, electoral divisions, counties, and provinces in a hierarchical way. In townlands.ie, the data is available as GeoJSON, Shapefile, and CSV formats.

However townlands.ie does not provide town and city-level information, and we needed to identify other potentially useful attributes to add to the OSM-derived resources. To do this we also used the Logainm.ie dataset in order to add some other types of land divisions. Logainm is the Placenames Database of Ireland. It is an official resource that provides authoritative information on the names of places across Ireland. It is developed and maintained by Fiontar & Scoil na Gaeilge at Dublin City University (DCU) in collaboration with the Placenames Branch of the Department of Tourism, Culture, Arts, Gaeltacht, Sport and Media, Logainm serves as a comprehensive database for Irish-language and English-language place names.

5.2. Geospatial Data Uplift

Using the retrieved datasets, first, we accessed the data in a structured format such as GeoJSON and CSV. Next, we examined the key attributes like the townland name, geographic coordinates, county, and boundaries. We used

these attributes to create RDF triples by mapping them to relevant ontologies, including GeoSPARQL, CIDOC-CRM and VRTI Ontology. We used RML (RDF Mapping Language) tool called Morph-KGC[9] to transform the data programmatically. Each townland becomes a unique RDF subject (URI), while attributes (e.g., geo:lat, geo:long) and relationships define predicates and objects. We used Turtle serialisation ensuring compatibility for use in Linked Data and SPARQL queries. An extract of an RML mapping used to uplift townlands is shown in Listing 1.

The dataset connects people to places by linking individuals to specific geographic locations, such as counties or townlands, represented within the dataset. These connections are established through relationships that capture associations like birthplace, and death place which shows significant life events tied to the places. For example, a person might be associated with Dublin as their birthplace. By integrating these connections, the dataset provides a rich contextual network that intertwines geographic and personal historical data, enabling insights into the movement, influence, and distribution of people across regions. This enhances the understanding of social, cultural, and historical dynamics within the geographical framework.

Listing 1: Extract of RML mapping to uplift townland centroid coordinates

```
<#TownlandCentroid>
        a rr:TriplesMap ;

        rml:logicalSource [
                rml:source "data_townland.csv" ;
                rml:referenceFormulation ql:CSV
        ];

        rr:subjectMap [
                rr:template "https://kg.virtualtreasury.ie/place/present-day/townland/centroid/{NAME_COMBINED
                    }/{VRTI_ID}" ;
                rr:class cidoc:E54_Dimension , geo:Geometry ;
        ] ;

        rr:predicateObjectMap [
                rr:predicate geo:asWKT ;
                rr:objectMap [
                        rr:template "POINT␣({LONGITUDE}␣{LATITUDE})" ;
                        rr:datatype geo:wktLiteral ;
                ] ;
        ] ;
.
```

To illustrate, Listing 2 describes a geographic entity, Dublin County, in terms of its attributes, relationships, and contextual data, using a combination of ontologies, namespaces, and specific properties. Each geographic entity (e.g., townland, county) is assigned a unique IRI. The Dublin County resource is identified as both a **cidoc:E53_Place** (a general place in CIDOC-CRM) and a **geo:Feature** (a geographic feature). It has a preferred identifier through **cidoc:P48_has_preferred_identifier** referencing the appellation resource. This is especially needed where there are more than one place name for the places. The type of the entity is **vrti:PresentDayCounty**, identifying it as a present-day administrative county. Since we will have counties for different periods, it is crucial to have types for specific eras. **cidoc:P43_has_dimension** links to resources defining the county's boundary and centroid. **geo:hasBoundingBox** and **geo:hasCentroid** reiterate these spatial attributes. Spatial properties reference to points and polygons. Townlands.ie provides geographic data derived from Open Street Maps data with present day Irish administrative boundaries. **geo:hasMetricArea** provides the area in square meters as 929017721.782945. **cidoc:P89_falls_within** places Dublin County within the province of Leinster and **cidoc:P71i_is_listed_in** links to external resources. **owl:sameAs** asserts equivalence to <http://data.logainm.ie/ place/100002>, establishing interoperability with Logainm's data. **vrti:OSM_ID** provides the OpenStreetMap (OSM) identifier for Dublin. **vrti:VRTI_KG_ID** assigns a unique internal reference ID for Dublin within the Virtual Treasury Knowledge Graph.

[9]https://github.com/oeg-upm/morph-kgc

Listing 2: A code snippet for the county Dublin for different eras (turtle syntax)

```
<https://kg.virtualtreasury.ie/place/present-day/county/Dublin/v1s4kv2> a cidoc:E53_Place,
        geo:Feature ;
    rdfs:label "Dublin"@en ;
    cidoc:P1_is_identified_by <https://kg.virtualtreasury.ie/place/present-day/county/appellation/Dublin/
        v1s4kv2> ;
    cidoc:P2_has_type vrti:County, vrti:PresentDayCounty ;
    cidoc:P43_has_dimension <https://kg.virtualtreasury.ie/place/present-day/county/boundary/Dublin/v1s4kv2>,
        <https://kg.virtualtreasury.ie/place/present-day/county/centroid/Dublin/v1s4kv2> ;
    cidoc:P48_has_preferred_identifier <https://kg.virtualtreasury.ie/place/present-day/county/appellation/
        Dublin/v1s4kv2> ;
    cidoc:P71i_is_listed_in <http://www.townlands.ie/dublin>,
        <https://www.logainm.ie/en/100002> ;
    cidoc:P89_falls_within <https://kg.virtualtreasury.ie/place/present-day/province/Leinster/v14rt5h> ;
    geo:hasBoundingBox <https://kg.virtualtreasury.ie/place/present-day/county/boundary/Dublin/v1s4kv2> ;
    geo:hasCentroid <https://kg.virtualtreasury.ie/place/present-day/county/centroid/Dublin/v1s4kv2> ;
    geo:hasMetricArea "929017721.782945"^^xsd:float ;
    owl:sameAs <http://data.logainm.ie/place/100002> ;
    vrti:OSM_ID -282800 ;
    vrti:VRTI_ERA <https://www.w3id.org/virtual-treasury/vocabulary#Present-Day> ;
    vrti:VRTI_KG_ID "v1s4kv2" .

<https://kg.virtualtreasury.ie/place/modern-1750-1922/county/Dublin/v1k1j4p>
        a                       crm:E53_Place ;
        rdfs:label              "Dublin"@en ;
        crm:P1_is_identified_by <https://kg.virtualtreasury.ie/place-appellation/Dublin> , <https://kg.
            virtualtreasury.ie/vrti-identifier/v1k1j4p> ;
        crm:P2_has_type         vrti:ModernCounty , vrti:County ;
        crm:P48_has_preferred_identifier
                <https://kg.virtualtreasury.ie/place-appellation/Dublin> ;
        crm:P189_approximates   <https://kg.virtualtreasury.ie/place/present-day/county/Dublin/v1s4kv2> ;
        crm:P89_falls_within    <https://kg.virtualtreasury.ie/place/modern-1750-1922/province/Leinster/
            v1y2d2b> ;
        vrti:VRTI_ERA           <https://www.w3id.org/virtual-treasury/vocabulary#Modern-1750-1922> ;
        vrti:VRTI_KG_ID "v1k1j4p" .

<https://kg.virtualtreasury.ie/place/early-modern-1500-1749/county/Dublin/v16ww1c> a cidoc:E53_Place,
        geo:Feature ;
    rdfs:label "Dublin"@en ;
    cidoc:P189_approximates <https://kg.virtualtreasury.ie/place/present-day/county/Dublin/v1s4kv2> ;
    cidoc:P1_is_identified_by <https://kg.virtualtreasury.ie/place-appellation/Dublin> ;
    cidoc:P2_has_type vrti:County,
        vrti:EarlyModernCounty ;
    vrti:VRTI_ERA <https://www.w3id.org/virtual-treasury/vocabulary#Early-Modern-1500-1749> ;
    vrti:VRTI_KG_ID "v16ww1c" .
```

5.3. Geospatial Data for Historical Time Periods

As mentioned previously, a number of geospatial "stacks" have been created in VRTI-KG for different historical eras enabling a structured collection of spatial and temporal data that captures how locations, entities, and relationships evolve over time. Thus, VRTI-KG geospatial dataset spans multiple time periods (Fig. 3, enabling historical comparisons by connecting diverse datasets (e.g., historical maps, census records) to build a rich, interlinked dataset. The identifiers for different time periods of Dublin and the connections between them can be seen in Listing 2.

The Present-Day geospatial dataset contains data with geographic references, such as: i) Coordinates (latitude, longitude) (Listing 3) ii) Spatial features (points, polygons) (Listing 3) iii) Temporal aspects (VRTI-Eras). On the other hand, it is not as trivial for other eras to discover that information, since the borders of places were not as certain or precise in the previous era as of now. Thus we connected previous eras to Present-Day and use the interlinks to approximate the previous eras spatial features using present day ones. As more details of historic boundaries become known (e.g. through digitisation of old maps), these will replace the approximations calculated over time. Having these geospatial dataset connections across different eras for people, events, and objects enables the integration of historical and modern data while maintaining relationships among entities. This approach enhances spatiotemporal analysis, linking historical narratives with contemporary present day geographical contexts. If the entity is not a living entity (e.g. an organisation or office) then historical context can help to track entities over different eras.

For example, Early Modern Places is a dataset that includes a unique identifier for Early Modern places. The idea is to separate the Present Day places from the previous versions. Because most of the time neither name of the places or the boundaries were the same with the Present-Day data. Instead most of the time there is a 1-to-n relationship between places. This means a townland in the Early Modern dataset can be connected to more than 1 Present-Day townlands. For example, the early modern townland Abbey Land of Sligo was later divided into two present-day townlands: Abbeyquarter North and Abbeyquarter South. This shows us that the lands and organisational areas are divided to smaller lands as of now. Early Modern places are connected to Present-Day Dataset using **cidoc:P189_approximates** property. This property associates an instance of E53 Place with another instance of E53 Place, which is defined in the same reference space to approximate the former[10].

Listing 3: Point and polygon info for the county Dublin (turtle syntax)

```
<https://kg.virtualtreasury.ie/place/present-day/county/centroid/Dublin/v1s4kv2> a cidoc:E54_Dimension,
        geo:Geometry ;
    geo:asWKT "POINT (-6.2839761153 53.3922226815)"^^geo:wktLiteral .
<https://kg.virtualtreasury.ie/place/present-day/county/boundary/Dublin/v1s4kv2> a cidoc:E54_Dimension,
        geo:Geometry ;
    geo:asWKT "MULTIPOLYGON (((-6.5468918 53.2794837, -6.5464238 53.2797712, -6.5451069 53.2810147,
        -6.5438859 53.282084, -6.5434005 53.2822258, -6.5418904 53.2825208, -6.5404141 53.2828031,
        -6.5395815 53.2830361, -6.538967 53.2832576, -6.5373046 53.2839006, -6.5367397 53.2841049,
        -6.5361939 53.2842907, ... -6.0884458 53.2749605)))"^^geo:wktLiteral .
```

6. VRTI Case Study for Cultural Heritage

In order to illustrate the user experience in seamless navigation across VRTI-KG and VRTI-DB, this section presents an example case study. The case study focuses on the "State Papers Ireland, 1660–1715 Gold Seam"[11] in the VRTI, which presents the correspondence that passed between the governments at Dublin and Whitehall in the sixteenth, seventeenth, and eighteenth centuries.

A search for Roger Boyle in the VRTI Knowledge Graph returns a KG entity card as seen in Figure 10 which is the KG entity card for him, in this case populated using information from the KG and from the Royal Irish Academy's *Dictionary of Irish Biography* entry for Boyle. The information was written by a professional historian and is understood to be as precise as possible and verifiable. Roger Boyle was a prominent politician in seventeenth century Ireland, the twelfth child and third son of Richard Boyle, earl of Cork, the largest landowner in Ireland. Roger was designated as lord Broghill until 1660, when the monarch Charles II granted him the earldom of Orrery in the Irish peerage. Boyle had a distinguished, if controversial, career. A soldier and enemy of the crown in the wars of the 1640s and 1650s, he thereafter reconciled with the monarch Charles II in 1660 and was appointed to a commission as lord justice of Ireland (an interim governor). He continued in positions of authority until his death. In the course of his duties he generated significant quantities of paperwork, which have survived in UK The National Archives, where they form part of the State Papers Ireland series.

Boyle's primary correspondents were the secretaries of state at Whitehall, who acted as a clearing house for all royal correspondence intended for the monarch. These links are shown Treasury Links in Figure 10. For example one of the Treasury links for Boyle when clicked brings the user to that correspondence as stored in the VRTI-DB as shown in Figure 11.

The role of lord justice was undertaken as a commission, and Boyle served with the lord chancellor of Ireland, Sir Maurice Eustace. As a result, they co-wrote, or co-signed, numerous pieces of correspondence (as seen in Figure 12) and they were mentioned in their respective letters to London. If the user clicks on Eustace in Figure 11, they are brought to his KG entity card (as seen in Figure 13) and in his duties he generated similar volumes of correspondence.

[10] https://site2024.cidoc-crm.org/Property/p189-approximates/version-7.1.1

[11] https://virtualtreasury.ie/gold-seams/state-papers-ireland-1660-1715

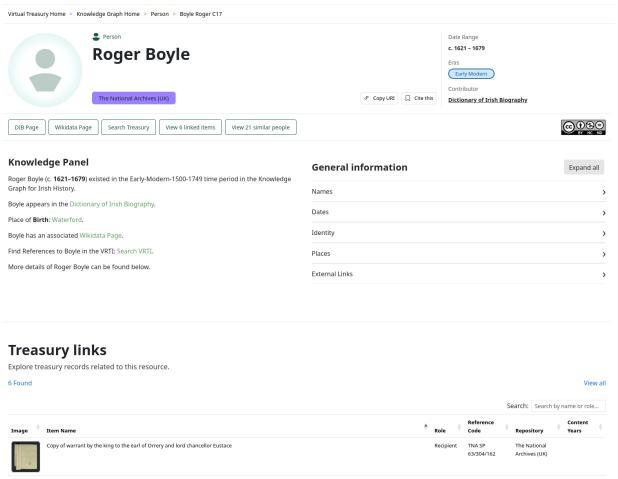

Fig. 10. Screenshot for URI https://kg.virtualtreasury.ie/person/Boyle_Roger_c17/v1md59f

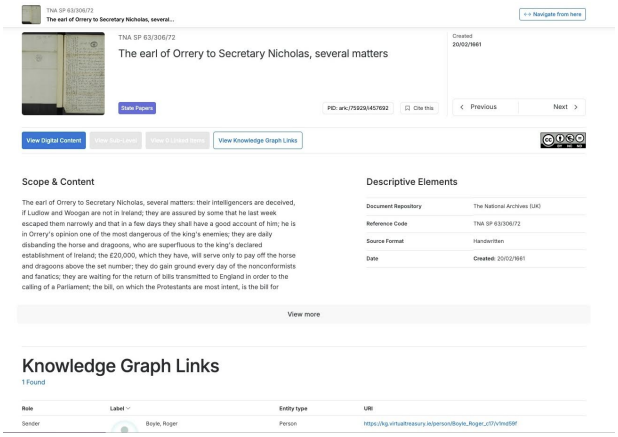

Fig. 11. Screenshot for https://virtualtreasury.ie/item/TNA-SP-63-306-72

The metadata that sits behind the correspondence has been curated to specifically identify the sender and recipient of each piece of correspondence, and this data is then matched to determine if a VRTI-KG entry exists. As the correspondents appear across the series, their letters are automatically linked to the KG-entry for the sender or recipient, for presentation to the user.

7. Related Work

As more historical original sources become digitised and more available for historians to analyse, increasingly Digital Humanities historical projects are emerging that incorporate the use of semantic web and Linked Data approaches.

A notable foundation for some of these is the SampoUI framework created and maintained by the Semantic Computing Research Group (SeCo) in Aalto University, Finland. There are a large number of example successful projects arising. "AcademySampo" [14, 15] is a semantic portal for biographical and prosopographical research, and the article presents how the Finnish registries "Ylioppilasmatrikkeli" 1640–1852 and 1853–1899, containing comprehensive biographical information on academics in Finland, have been created into a Linked Open Data service adhering to the FAIR principles . Taking a similar approach "BiographySampo" [7, 22] is designed to provide

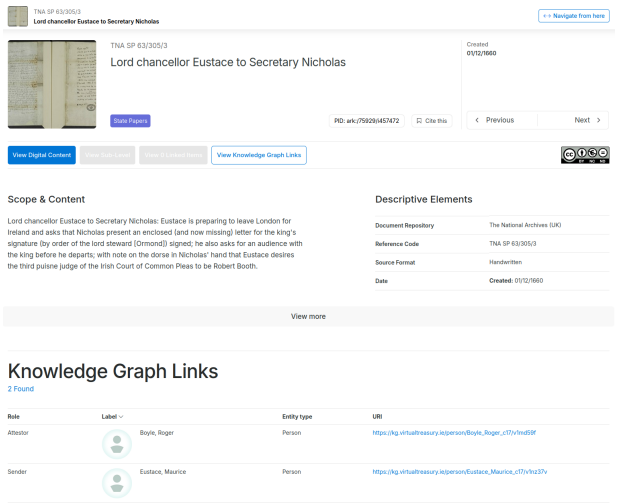

Fig. 12. Screenshot for https://virtualtreasury.ie/item/TNA-SP-63-305-3

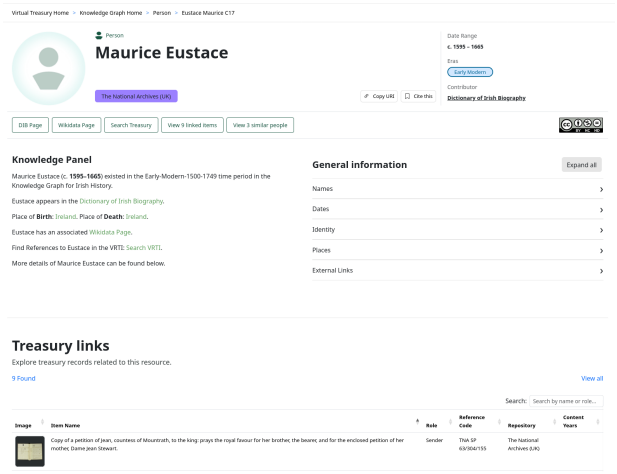

Fig. 13. Screenshot for URI https://kg.virtualtreasury.ie/person/Eustace_Maurice_c17/v1nz37v

a semantic portal to Finnish Biographies. This is based on an automatically derived knowledge graph from a set of 13,100 textual biographies, enriched with information linking to sixteen additional data sources that have been obtained through data harvesting from external collections found in archives, museums, and libraries. "WarSampo" Knowledge Graph [9, 12], is a Linked Open Data resource with detailed metadata on over 100,000 individuals related to Finland in World War II, through a Linked Data approach aggregating open data distributed across various local military cemeteries where the dead were buried. It also provides tools for data analysis and support for digital humanities studies. A more recent example of the use of the Sampo framework that being being implemented is where Epistolar data (1809–1917) were uplifted to analyze their metadata and social structure [8]. The History of Vienna [13] uses Semantic MediaWiki for collaborative editing and KG curation in the Vienna History Wiki, linking entities such as people, events, and places to external sources such as Wikidata and Schema.org.

Elsewhere, using different approaches, "The People of Medieval Scotland" project[12] [17] incorporates information on individuals involved in actions in or related to Scotland between the death of Malcolm III in 1093 and

[12] www.poms.ac.uk

Robert I's parliament in 1314. The data is structured to reflect social interactions and relationships. It uses a general "factoid-oriented" structure model that links people to the information about them that appears in the primary source manuscript. The "factoid model" of digital prosopography has been applied to major databases dealing with medieval England, Scotland, France, and Byzantium. "The Digital Prosopography of the Roman Republic" [5] aims to enhance prosopographical research on the Roman Republic's elite by creating a searchable database with a terminological approach. Similiar to VRTI avoiding it provides information on careers, office holdings, family relationships, and personal status, with the intention of enabling analysis through statistical and quantitative methods.

8. Lessons Learnt

In this project, we combined multiple data sources to provide comprehensive coverage of geospatial data. Transforming geographic datasets into rich semantic resources improves data usability and interoperability. Using ontologies such as GeoSPARQL, CIDOC-CRM, and VRTI Ontology enables structured and meaningful geospatial data representation. Townlands.ie provides detailed hierarchical Irish administrative divisions but lacks city-level and town-level data. Logainm complements this by offering authoritative Irish place names and additional land divisions. Integrating data from multiple sources ensures a more complete and accurate dataset. On the other hand, both datasets have different numbers of entities (e.g. townlands, parishes) which makes it difficult to identify the accuracy of the source dataset.

Historical places often have different names and boundaries across time periods. The 1-to-n relationship between historical and present-day places reflects territorial changes over time. The use of **cidoc:P189_approximates** property allows linking present-day places with their historical counterparts. Linking people, events, and objects to places enables spatiotemporal analysis across different eras. Unlike modern datasets, historical locations often have uncertain or changing borders. Approximation techniques are needed to map past locations to present-day equivalents for consistency.

We have used a user-friendly interface to enable user interaction. Natural language presentation and data-driven configurations improve accessibility for non-technical users. Providing intuitive search filters and dropdown suggestions enhances the search and data entry experience. Interactive maps help users navigate and understand spatial relationships better. Geospatial visualisations and summary pages provide clearer insights into people and places. Expanding data sources and enhancing visualisation tools can support better historical analysis. Using JSON configurations ensures seamless updates between the interface and the underlying KG. SHACL validation[13] helps maintain data integrity when modifying or adding resources.

User evaluations showed high satisfaction and understanding, with improvements in efficiency for those familiar with semantic web concepts. Historians using the editor demonstrated strong accuracy in editing tasks (77% correct edits). Some users struggled with efficiency, likely due to varying levels of prior knowledge. Participants requested additional tooltips and simplification of technical terms for better clarity. Based on user feedback, refinements were implemented to improve tool usability. Further evaluations are planned to assess the impact of these refinements.

References

[1] A. Alliance, ARK Alliance.*Home of the Archival Resource Key (ARK)*, 2025. https://arks.org/.
[2] CIDOC CRM Special Interest Group, Definition of the CIDOC Conceptual Reference Model, Version 7.1.3, 2024, Accessed April 2025.
[3] C. Debruyne, G. Munnelly, L. Kilgallon, D. O'Sullivan and P. Crooks, Creating a Knowledge Graph for Ireland's Lost History: Knowledge Engineering and Curation in the Beyond 2022 Projec, *ACM Journal on Computing and Cultural Heritage* **15**(2) (2022), 1–25.
[4] D. Garijo, O. Corcho and M. Poveda-Villalón, FOOPS!: An Ontology Pitfall Scanner for the FAIR Principles, in: *International Semantic Web Conference (ISWC) 2021: Posters, Demos, and Industry Tracks*, CEUR Workshop Proceedings, Vol. 2980, CEUR-WS.org, 2021. http://ceur-ws.org/Vol-2980/paper321.pdf.
[5] M. Hammond, From Digital Prosopography to Social Network Analysis, *Medieval People: Social Bonds, Kinship, and Networks* **36** (2021), 235–262.

[13] https://github.com/alex-randles/SWJ-Supplementary-Data/blob/main/shapes.ttl

[6] S. Harris, A. Seaborne and E. Prud'hommeaux, SPARQL 1.1 Query Language, Vol. 21, 2013, p. 778. https://www.w3.org/TR/sparql11-query/.
[7] E. Hyvönen, P. Leskinen, M. Tamper, H. Rantala, E. Ikkala, J. Tuominen and K. Keravuori, BiographySampo–publishing and enriching biographies on the semantic web for digital humanities research, in: *European Semantic Web Conference*, Springer, 2019, pp. 574–589.
[8] E. Hyvönen, P. Leskinen, H. Poikkimäki, H. Rantala, J. Tuominen, S. Drobac, O. Koho, I. Pikkanen and H.-L. Paloposki, Searching, exploring, and analyzing historical letters and the underlying networks: LetterSampo Finland (1809–1917) data service and semantic portal, 2025.
[9] E. Ikkala, M. Koho, E. Heino, P. Leskinen, E. Hyvönen, T. Ahoranta et al., Prosopographical Views to Finnish WW2 Casualties Through Cemeteries and Linked Open Data., in: *WHiSe@ ISWC*, 2017, pp. 45–56.
[10] E. Ikkala, E. Hyvönen, H. Rantala and M. Koho, Sampo-UI: A full stack JavaScript framework for developing semantic portal user interfaces, *Semantic Web* **13**(1) (2022), 69–84.
[11] H. Knublauch and D. Kontokostas, Shapes Constraint Language (SHACL), 2017. https://www.w3.org/TR/shacl/.
[12] M. Koho, H. Rantala and E. Hyvönen, Digital Humanities and Military History: Analyzing Casualties of the WarSampo Knowledge Graph, in: *Proceedings of the 6th Digital Humanities in the Nordic and Baltic Countries Conference (DHNB 2022)*, CEUR-WS. org, 2022.
[13] B. Krabina, Building a knowledge graph for the history of vienna with semantic mediaWiki, *Journal of Web Semantics* **76** (2023), 100771.
[14] P. Leskinen and E. Hyvönen, Linked open data service about historical Finnish academic people in 1640–1899, in: *Proceedings of the Digital Humanities in the Nordic Countries 5th Conference*, CEUR-WS. org, 2020.
[15] P. Leskinen, H. Rantala and E. Hyvönen, Analyzing the lives of Finnish academic people 1640–1899 in nordic and Baltic countries: AcademySampo data service and portal, in: *Proceedings of the 6th Digital Humanities in the Nordic and Baltic Countries Conference (DHNB 2022)*, CEUR-WS. org, 2022.
[16] J.R. Lewis, Psychometric Evaluation of the PSSUQ Using Data from Five Years of Usability Studies, *International Journal of Human-Computer Interaction* **14** (2002), 463–488. doi:10.1080/10447318.2002.9669130.
[17] M. Pasin and J. Bradley, Factoid-based prosopography and computer ontologies: towards an integrated approach, *Digital Scholarship in the Humanities* **30**(1) (2015), 86–97.
[18] Permanent Identifiers Community, W3ID: Permanent Identifiers for the Web, 2025, Accessed April 2025.
[19] A. Randles, L. McKenna, L. Kilgallon, B. Yaman, P. Crooks and D. O'Sullivan, Evaluating the Knowledge Graph Editor of the Virtual Record Treasury of Ireland, in: *Proceedings of the 9th International Workshop on the Visualization and Interaction for Ontologies, Linked Data and Knowledge Graphs co-located with the 23rd International Semantic Web Conference (ISWC 2024), Baltimore, USA, November 12, 2024*, B. Fu, P. Lambrix, H. Li, S. Nunes and C. Pesquita, eds, CEUR Workshop Proceedings, Vol. 3773, CEUR-WS.org, 2024, pp. 62–77. https://ceur-ws.org/Vol-3773/paper5.pdf.
[20] A. Randles, D. O'Sullivan, L. McKenna, B. Yaman, L. Kilgallon and P. Crooks, Facilitating Search of the Virtual Record Treasury of Ireland Knowledge Graph using ChatGPT, in: *Joint Proceedings of Posters, Demos, Workshops, and Tutorials of the 20th International Conference on Semantic Systems co-located with 20th International Conference on Semantic Systems (SEMANTiCS 2024), Amsterdam, The Netherlands, September 17-19, 2024*, D. Garijo, A.L. Gentile, A. Kurteva, A. Mannocci, F. Osborne and S. Vahdati, eds, CEUR Workshop Proceedings, Vol. 3759, CEUR-WS.org, 2024. https://ceur-ws.org/Vol-3759/paper15.pdf.
[21] A. Randles, L. McKenna, L. Kilgallon, B. Yaman, P. Crooks and D. O'Sullivan, The Knowledge Graph Explorer for the Virtual Record Treasury of Ireland, in: *Proceedings of VOILA 2024: The 9th International Workshop on the Visualization and Interaction for Ontologies, Linked Data and Knowledge Graphs co-located with the 23rd International Semantic Web Conference (ISWC 2024)*, CEUR-WS. org, 2024.
[22] H. Rantala, E. Hyvönen and J. Tuominen, Finding and explaining relations in a biographical knowledge graph based on life events: Case BiographySampo, in: *CEUR Workshop Proceedings*, Vol. 3443, RWTH Aachen University, 2023.
[23] W3C, SKOS Simple Knowledge Organization System Reference, 2009, W3C Recommendation, accessed April 2025.
[24] B. Yaman, L. McKenna, A. Randles, L. Kilgallon, P. Crooks and D. O'Sullivan, Digital Prosopographical Information in the Virtual Record Treasury of Ireland's Knowledge Graph for Irish History, in: *Proceedings of the 1st International Workshop of Semantic Digital Humanities (SemDH) Co-Located with the 21st Extended Semantic Web Conference*, 2024.
[25] B. Yaman, K. Thompson, F. Fahey and R. Brennan, LinkedDataOps:quality oriented end-to-end geospatial linked data production governance, *Semantic Web* **15**(2) (2024), 555–581. doi:10.3233/SW-233293.