

Geoff: A Linked Data Vocabulary for Describing the Form and Function of Spatial Objects

Kris McGlinn^[1], Declan O’Sullivan^[1], Christophe Debruyne^[1], Eamonn Clinton^[2]
and Rob Brennan^[3]

¹ ADAPT, Trinity College Dublin, Dublin 2, Ireland

² Ordnance Survey Ireland, Dublin 8, Ireland

³ ADAPT, Dublin City University, Dublin 9, Ireland
`kris.mcglinn@adaptcentre.ie`

Abstract. Geoff, the geospatial form and function vocabulary, is a comprehensive RDF-based spatial object classification scheme based on a separation of the concepts of form and function. Geoff is based on our analysis of the extensive (over 50 million spatial object instances) Digital Landscape Model (DLM) Core model maintained by Ordnance Survey Ireland (OSi). We propose Geoff as there are currently no open geospatial form and function classification systems that cover the full range of geospatial objects (from buildings and roads to lakes and other natural features) modelled as Linked Data or in any other formalism. Geoff is a generalization of the DLM Core schema and adopts the GeoSPARQL ontology. Geoff was initially developed to make these classifications available for OSi’s geospatial Linked Data as they facilitate the publications of more expressive models of spatial features. For example, to state that a church building (form) is now used as apartments (function). Geoff is now presented to the wider community for reuse and extension to meet their own needs. Geoff supports geospatial queries based on form and function and interlinking of geo-information datasets using different form and function code lists. The Geoff ontology follows Linked Data publishing best practice in terms of available metadata, documentation, and quality assurance.

Keywords: Ontology, Form and Function, Geospatial Feature.

1 Introduction

Publishing geospatial Linked Data is growing in popularity with a wide range of national initiatives across Europe, e.g., in Switzerland¹, Spain², France³, and Ireland⁴. The initial geospatial concepts modelled have focused on boundaries and placenames [1], but as geospatial Linked Data becomes established, we expect a wider array of

¹ <https://www.geo.admin.ch/en/geo-services/geo-services/linkedata.html>

² <http://vocab.linkedata.es/datosabiertos/def/sector-publico/territorio>

³ <http://data.ign.fr/def/geofla/20140822.htm>

⁴ <http://data.geohive.ie/>

geospatial information will be published. New vocabularies are thus needed to express the different types of features and diverse foci of the next generation of geospatial datasets. For example, building on our previous boundary datasets [2], the ADAPT Centre and Ordnance Survey Ireland (OSi) initiated a project to publish Linked Data identifiers and data for the 50 million spatial entities in Ireland currently tracked by OSi in 2017. This Linked Data will include the built environment, e.g., buildings, roads, and monuments, and the natural environment such as rivers, lakes, and mountains.

One important geospatial information system (GIS) classification taxonomy for both the built and natural environment [3], [4] is the specification of a spatial entity's form (which characterizes its physical shape and appearance) and the related concept of a spatial entity's function (which specifies the type of use of a spatial feature). For example, a building may have the form "church" and the function "apartments" in the case of an older church that was converted into apartments. Standardization agencies such as ISO [5] are active in GIS domain and have published standardized taxonomies for form and function. Unfortunately, to the authors' knowledge, none of these classification systems has yet been made available as RDF, RDFS, or OWL to facilitate geospatial Linked Data publication that includes these fundamental GIS concepts in an interoperable way.

In this paper, we present and describe Geoff⁵, the geospatial form and function vocabulary based on over a decade of data model development within OSi on form and function modelling within their Prime2 [7] national spatial data infrastructure that has assessed both existing public spatial form and function standards. Geoff also draws upon the OSi's experience in classifying 50 million spatial features in DLM Core. In line with the principles of reuse and minimal extension, Geoff is designed to be used to sub-classify any individuals of the class GeoSPARQL Feature [6] (which is an abstract class for representing any geospatial entity such as buildings, rivers, and cities) and to link them with a set of function types defined in Geoff based on existing industry practice (the Irish national spatial infrastructure OSi DLM Core) where that is more comprehensive than existing standards. Geoff form and function types will be easy to extend as new use cases emerge. These extensions can be performed locally or as proposed revisions to Geoff itself. We demonstrate the vocabulary's suitability as a solution by elaborating on its deployment within OSi as part of the Irish national spatial object identifier publication initiative, which has resulted in over 200 thousand buildings in Galway published with geolocation and form and function⁶.

The rest of this paper is structured as follows: Section 2 describes the current state of the art in form and function classification, section 3 provides an overview of the Geoff vocabulary, Section 4 describes an example application of Geoff in Ireland's new geospatial Linked Data spatial entity publishing initiative, Section 5 evaluates Geoff under the headings; impact, reusability and accessibility, design and implementation. Finally, Section 6 provides conclusions and discusses future directions for Geoff.

⁵ <http://ontologies.geohive.ie/geoff/index-en.html>

⁶ <http://data.geohive.ie/downloadAndQuery.html>

2 Form and Function Classifications

A fundamental principle in ontology engineering is the reuse of vocabularies to avoid semantic heterogeneity. It was envisaged that existing form and function standards could be used as a basis for semantically annotating the OSi DLM core data. However, none were found in the geospatial domain. We, therefore, explored other non-geospatial standards. Form and function are concepts that are well understood in architecture.

The main classification systems for structural entities in the AEC domain are the US OmniClass [3] and UK UniClass2 [4]. Both are based on the ISO 12006-2 [5] classification for construction. OmniClass has a rich set of types for classifying entities, spaces, elements, products, etc. with a focus on the construction phase of buildings, landscape, and infrastructure. Both form and function can classify entities. UniClass2 also provides classifications for the built environment, but with a different range, including regions, complexes, spaces, and activities, the definition of which are too fine-grained at this point for the OSi, which classifies at a building level. However, it may be useful in future iterations as existing work is exploring the conversion of footprint polygons into wall products. UniClass2 does not classify entities by form. Instead, entities appear to be defined using terminology that reflects their form in a table called “Entities”. UniClass has 452 entities (not called forms) defined and 78 functions (defined for elements, which are components of a structure), OmniClass has been described above.

Due to the different approach to classification in UniClass2 (i.e., entities without any description of form, and function described only for elements), direct comparisons between it and the Geoff form and function concepts become difficult. There is a more closely aligned mapping between OmniClass and OSi’s classification of form and function in this respect. In either case, there is a lack of an open, freely available OWL-based spatial object classification scheme that clearly defines form and function for geospatial features, including buildings and infrastructure. This work seeks to address this through the development of OSi form and function types into a single OWL vocabulary. Also considering OSi has classified the form and function values for each of their (> 3.5 million) buildings and is in the process of publishing a subset of their building data as RDF (currently > 200 thousand buildings published), the classification of the form and function of their buildings may be of interest to others who wish to classify buildings. We describe Geoff, which provides this classification, in the next section.

3 Overview of Geoff

Geoff is designed to provide a Linked Data vocabulary to separate the concepts of form and function for geospatial features, as is standard practice in geospatial information systems. Additionally, it provides the most comprehensive set of reusable identifiers for forms and functions that have been published to date in any encoding scheme by spanning both the GIS domain (OSi DLM Core) and the AEC domain (OmniClass), through the definition of a comprehensive hierarchy of standard GIS forms. Functions are additionally sub-classified according to their likely forms. This scheme provides support for an extensible classification of features when a feature’s function has been

identified and links functions with their typical forms. This provides support for bi-directional inference (as far as is possible within the limitations of OWL 2's EL profile).

Finally, we have manually identified links between the related GIS forms and AEC forms to support interworking between GIS-AEC applications, e.g., so a construction company could load a national spatial data representation of buildings and use that as the basis for further construction design. Knowledge engineers with a background in Building Information Modelling within ADAPT were responsible for driving this knowledge engineering activity. The vocabulary and links were validated both by subject matter experts in the OSi and by demonstrating the artefacts before deployment.

- geoff:Form describes “the physical composition of an object; what it actually is”
- geoff:Function describes “what an object does, or is used for”

Each form is a spatial entity and, as such, is modelled as a sub-class of the Open Geospatial Consortium standard GeoSPARQL [6] class geo:Feature. To minimize ontological commitment, geoff:Function is not linked to any specific external vocabulary or ontology of location functions as we are unaware of any single authoritative source of such knowledge.

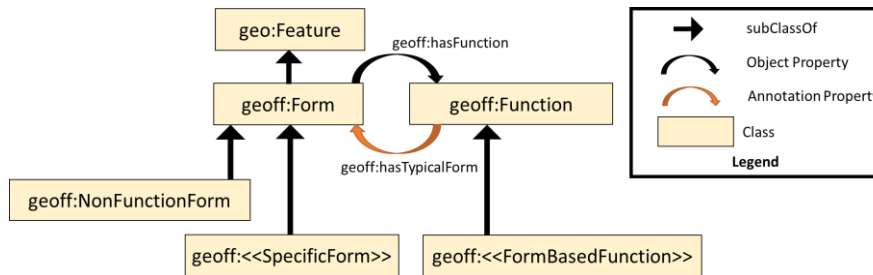


Fig. 1 illustrates the overall structure of the geospatial form and function vocabulary (Geoff). It defines two top-level classes geoff:Form and geoff:Function.

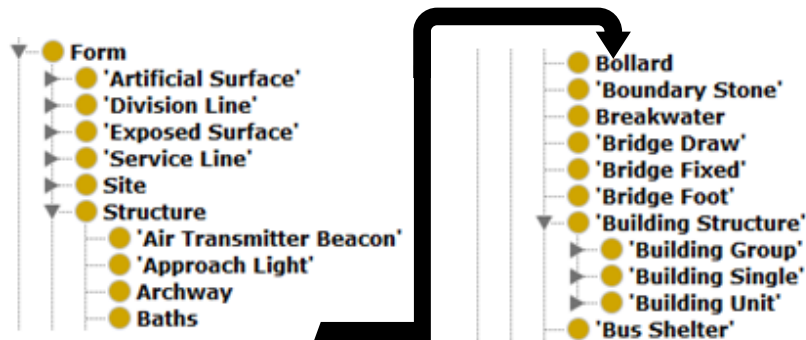


Fig. 2 Overview of some of Geoff's Class Form as shown in Protégé [12]

An object property geoff:hasFunction that links a specific geospatial feature to its function (or functions) is provided. By examination of DLM Core, we have identified some cases where no specific function is (or can) be associated with a feature or form

(e.g., a ruin). For these cases, we propose a sub-class of `geoff:Form` called `geoff:NonFunctionForm` (see **Fig. 1**). The second property supplied is `geoff:hasTypicalForm`, and since this is an association between an individual (subject) and a Form class (object), it is outside OWL EL. However, this is valuable information to record, we model the property as an annotation property. This makes the property at least available for querying through SPARQL, even if it is outside the formal logical model of Geoff.

The final component of Geoff is a pattern for specifying specific sub-classes of form based on the identified function of a specific spatial feature. To enable this, `geoff:Function` has 281 named individuals, each describing a specific type of function that can be assigned to a geospatial feature. A comprehensive set of functions appropriate to each form is specified in Geoff as defined classes. For example, the `geoff:PowerlineOverheadElectricityFunction` class has individual members such as 110kV power transfer function. Following this scheme, we introduced 337 specific (i.e., lowest level) form classes that can be assigned to a geospatial feature (**Fig. 2**). Given that we propose an ontology, users can easily extend these definitions for their purposes. Following Linked Data best practices, Geoff has the following metadata related to provenance and licensing included; `dc:title`, `owl:versionIRI`, `owl:versionInfo`, `dc:date`, `dc:creator`, `dc:description`, `dc:rights` and `cc:license`. Classes defined according to Live OWL Documentation Environment (LODE) requirements and each have `rdfs:label` values (`@en`) and `rdfs:comment` to classes. This is the same for properties.

In **Listing 1**, the class for form “Church” is shown. To assign the `formID`, it is necessary to create an individual. It is also at the individual level we assign the different function to a form. In **Listing 2**, the function “Catholic Church” is given. It should be noted that the numbering for the forms and functions is a direct result of the numbering within the DLM core data set. Each form and function is stored in a table, and has a unique integer identifier localized to that table. Therefore, the number provides nothing more than an idea of the sequence of the form or function in that table. As within the OSi Prime2 governance model these numbers are the authoritative identifiers and therefore are unlikely to change, so these predicates are considered authoritative.

```
:Form84 rdf:type owl:Class ;
  rdfs:subClassOf :BuildingSingle ,
                :Site ;
  rdfs:comment "A spatial feature with the shape of Church" ;
  rdfs:label "Church"@en , "Church" .
```

Listing 1 The form class “Church” in Geoff (turtle serialization)

```
:function81
  rdf:type owl:NamedIndividual , :BuildingSingleFunction , :SiteFunction ;
  hasTypicalForm :Form84 ;
  rdfs:comment "The function Catholic Church" ;
  rdfs:label "Catholic Church"@en , "Catholic Church" .
```

Listing 2 The function “Catholic Church” in Geoff (turtle serialization)

4 Use Case: OSi Linked Data Platform and Publication of Irish Buildings Data as Linked Data

The OSi's geospatial data is being converted and made available as RDF. This began with the publication of boundary data in 2017 [2] and is now proceeding with building data [15], with the aforementioned buildings published for Galway county. Making OSi building data open and available as Linked Data has the potential to become an authoritative, central catalogue of Irish buildings available for interlinking and enriching Irish building data (e.g., with form and function). Geoff plays a central role in this new larger tranche of Linked Data being published by OSi. This section briefly describes some of the components of the OSi Linked Data Platform to provide some context to the use case on building data publication.

The central component of the platform is an Oracle Spatial and Graph⁷ version 12.1 server which hosts the OSi Prime2 data and is also a native triplestore with geospatial inferencing available. It should be noted that this work is conducted on the DLM Core distribution of Prime2, which is a non-normalised subset of Prime2 provided to OSi customers. Most of the concepts carry over and it is expected that only small changes will have to be made to Geoff to represent Prime2 directly. Prime2 has therefore been the main driver of the ontology development. The process of generating the original boundary vocabulary was presented previously in [2], and, more recently, the analysis and publication of the building data was discussed in [15].

The analysis resulted in both a vocabulary for the geospatial objects required, as well as R2RML mappings to generate the RDF data from the master tabular data stored in Oracle. The vocabularies and data dumps (generated in RDF using these R2RML mapping) are made available on the geohive national spatial data website (data.geohive.ie). The generated RDF data is also published on an instance of Fuseki running on Oracle, exposing an endpoint to a Pubby⁸ instance adapted to also display geospatial data based upon OSi base maps⁹ (see [Fig. 3](#) & [Fig. 4](#)), making the RDF data available to browsers in a human-readable form. The endpoint also exposes the authoritative geospatial data, providing a catalogue of URIs, which can be interlinked with other Linked Data datasets. Currently, the focus of OSi is to make their building data more open.

The Prime2 buildings descriptions contain information about the buildings geospatial coordinate (the centroid of the building's footprint), a polygon footprint, names of buildings (in English and Irish), a GUID, addresses, way segments (i.e., where the entrance joins the road network), the status of the building (in use, derelict, etc.), and form and function. For each of these values, the OSi also maintains provenance information detailing who made changes, who authorized those changes, and when those changes took place [15]. Of this data, the OSi is currently only making openly available a subset, which includes its name (`rdfs:label`), geospatial coordinate (`geo:wktLiteral`), form and

⁷ <https://www.oracle.com/technetwork/database/options/spatialandgraph/overview/index.html>

⁸ <https://github.com/chrdebru/pubby> (forked from <https://github.com/pubby> to include maps)

⁹ <https://www.osi.ie/services/mapgenie/>

function (geoff:Form and geoff:Function), and most recent change to one of those values (a sub-property of the data property prov:wasGeneratedAt called osi:lastUpdate).

Galway (Ceannt) at OSi Geohive
<http://data.geohive.ie/resource/building/e1b361b0-11c8-495d-8803-e7771b618a38>

Property	Value
geoff:hasFunction	▪ geoff:Function632
geo:hasGeometry	▪ <geometrical representation>
rdfs:label	▪ Galway (Ceannt) (en) ▪ Stáisiún Ceannt (ga)
geohive:lastUpdate	▪ 2017-03-27T00:00:00 (xsd:dateTime)
rdf:type	▪ geoff:Form313 ▪ geohive:Building ▪ geo:Feature

Fig. 3 A sample of building data visualized using pubby¹⁰

Galway (Ceannt) >>> geo:hasGeometry at OSi Geohive

Property	Value
geo:asWKT	▪ POINT (530165.079999761 725308.870393976) (geo:wktLiteral)
is:geo:hasGeometry of	▪ <http://data.geohive.ie/resource/building/e1b361b0-11c8-495d-8803-e7771b618a38>
rdf:type	▪ geo:Geometry

As Turtle | As RDE/XML

© 2020 Ordnance Survey Ireland. Linked Data Frontend powered by Pubby

OSi Ordnance Survey | GeoHive | Trinity College Dublin | DCU

¹⁰ <http://data.geohive.ie/page/building/e1b361b0-11c8-495d-8803-e7771b618a38>

Fig. 4 A sample of building geometry data visualized using pubby

This subset provides an important basis for interlinking building data with the authoritative data the OSi collects and more recent work has specifically looked at converting the geospatial data into ifcOWL [16], to further semantically enrich the OSi building data. For each OSi building resource (>3.5 million), a human-readable URI is created to provide unique identification along with geospatial data and form and function. In **Listing 2** the result of this process is shown in turtle format. Here a building is defined with an Irish and English name, a geospatial point (defined using geo:asWKT), a lastUpdate and a geoff:hasForm and geoff:hasFunction relation assigning it a form and function value (like that in **Listing 1**). A full list of form and function values can be found in the ontology. Using this approach, form and function values are assigned to each building.

```
<http://data.geohive.ie/resource/building/ffa0e12b0b0f48719b72cb6f860ee96f>
  rdfs:label "Carrickemond Church"@en ;
  geoff:hasFunction geoff:function81;
  osi:lastUpdate "2015-06-19T16:37:55"^^xsd:dateTime ;
  geo:hasGeometry
    <urn:osi:build:geom:pnt:e487c4baa50941508de62bc6aaba762e1998-12-
25T00:00:00> ;
  a geo:Feature, osi:Building, geoff:Form84 .
```

Listing 2 Example Building Resource for the form Church (1 of 2355 Church form individuals)

```
PREFIX geo: <http://www.opengis.net/ont/geosparql#>
PREFIX geoff: <http://ontologies.geohive.ie/geoff#>
PREFIX qudt: <http://qudt.org/1.1/schema/qudt#>
SELECT * WHERE {
  ?building geo:hasGeometry ?feature .
  ?building geoff:hasFunction geoff:function81 .
  ?feature
    geo:nearby(53.3442497253418 -6.240039825439453 2 qudt:Kilometer) .
}
```

Listing 3 Example GeoSPARQL query for all building resources with function “Catholic Church” (geoff:function81) within 2 kilometers of a specific geospatial point (using a property function)

4.1 Methodology for data uplift

Data uplift is the conversion of structured or semi-structured data into Linked Data based upon semantic-web technologies. Our process for supporting geospatial semantic uplift is based on a standard methodology for ontology development¹¹, which consists of defining the scope, reuse of existing ontologies and vocabularies, enumeration of terms, definition of classes, properties and constraints, and finally the creation of in-

¹¹ https://protege.stanford.edu/publications/ontology_development/ontology101.pdf

stances. Ontology development is required where analysis determines no existing vocabulary can be found to satisfy the data exchange requirements defined within the scoping stage, or to support the interlinking process where multiple ontologies have been found. For Geoff, ontology development consisted of an analysis of the DLM Core data schema [15], as well as an iterative development life cycle using the Protégé tool [12]. This process involved periodic reviews by members of the OSi to ensure proper alignment with the Prime2 data schema. For data uplift, mappings must then be generated, either directly or through the use of mapping tools, such as R2RML, to support data conversion. We do not address the development of these mappings here, as these will be presented in a parallel publication, but we do provide an analysis of the RDF generated in the next section along with an evaluation of the Geoff ontology.

5 Evaluation

This section first presents an assessment of the Geoff vocabulary according to (i) potential impact, (ii) reuse and availability, and (iii) design & implementation. The sections below summarize these evaluations under a set of statements. Secondly, we provide an analysis of the distribution of instances in Geoff classes from a real-world dataset to gain insight into Geoff's ability to model the domain.

i) Impact: Geospatial information systems use form and function as important semantic attributes of features, for example, to enable function-based search across a range of otherwise unrelated feature types. GIS professionals expect to have form and function support in geospatial Linked Data. Geoff is the first geospatial feature and function vocabulary to be published in RDF. Before Geoff, there were no existing geospatial entity form and function classification schemes in RDF, except as embedded in the entity type classification, e.g., the DBpedia¹² subclasses of `dbo:Building`, which does not meet the needs of the GIS domain or give the same expressivity. Combined with the OSi Linked Data building registry, the additional semantics provided by Geoff enable new levels of interoperability between GIS data and AEC domain data which is necessary as future advances in remote drones, smart construction, self-driving vehicles and integrated energy grids leveraging the massive sensing power of national spatial data infrastructure.

ii) Reusability and Availability: Geoff is currently used by OSi for their publication of the National Irish buildings Linked Data (which is ongoing). However, building open data is being published by a growing number of national and regional agencies such as the recent publication by the city of Zurich of its 50000 detailed 3D buildings as open data¹³. Unfortunately, none of this open data is Linked Data due in part to a lack of vocabularies. This lack is the problem that Geoff addresses. The W3C Linked Building Data community group¹⁴ is starting to address Linked Data for building models and could adopt Geoff as it has no work to date on form and function models.

¹² <https://wiki.dbpedia.org/>

¹³ <https://80.lv/articles/zurichs-3d-buildings-available-as-open-data/>, <https://80.lv/articles/zurichs-3d-buildings-available-as-open-data/>

¹⁴ <https://w3c-lbd-cg.github.io/lbd/>

To facilitate reuse, Geoff is published as a LOD-comformant self-describing vocabulary, so metadata is provided and all entities have RDFS labels and comments sufficient to generate HTML documentation for the vocabulary in the W3C style. This paper also describes Geoff and provides additional background for users. Geoff can be (imported and) extended in two ways: (i) by adding new individuals for new forms and functions and (ii) by defining new inverse functional ID properties to represent different standard codes for individual forms or functions (both OSi and OmniClass codes are already supplied). For users of this resource to create their extensions, there would be required a process of integration and publication of future versions of Geoff. Geoff (vocabulary, metadata, and documentation) is published at a persistent w3id URI here¹⁵ and is licensed as CC 4.0 BY, and thus commercial reuse and sharing are permitted.

iii) Design and Implementation: Geoff has been validated using the OOPS! [13] ontology validation service, it supports LOD metadata and self-documentation, the W3C Data on the Web Best Practices¹⁶ such as versioning, metadata, vocabularies licensing, identifiers and quality having been validated by the Luzzu data quality framework¹⁷ for the quality dimensions of Understandability, Consistency, Syntactic Validity, and Licensing. In developing Geoff, we aimed at maximizing the reuse of existing standards. Geoff reuses OGC’s GeoSPARQL ontology to model geospatial features and the Dublin Core Element Set version 1.1 and OWL for metadata. To demonstrate its applicability, Geoff has been deployed and validated internally for OSi’s dataset of circa 3.5 million buildings. Currently, a subset of these have been published (over 200 thousand) with the intention of extending to include most if not all of the 50 million spatial entities OSi currently tracks in Ireland.

5.1 Instance Distribution by Applying Geoff to OSi DLM Core

In this section, we demonstrate the application of Geoff to classify individuals from a real GIS dataset (OSi’s DLM Core) to gain insight into Geoff’s fitness for distinguishing between individuals. This is a vocabulary quality indicator. Table 2 below presents a breakdown and analysis of the data generated through an R2RML-based conversion process from DLM Core to Geoff. It shows the numbers of different forms identified in the vocabulary and their relationship to the OSi DLM Core. The “DLM Core Table Name” is the name of a table in DLM Core which corresponds with the Class. The “Form” column gives the number of distinct form classes associated with that table. Tables for structures, water, and ways are merged into a single class. The “Function” column gives the number of distinct functions associated with those forms. The “Individuals” column gives the number of individuals which belong to that class in DLM Core, ranging from >22 million individuals for “Division Line” (which has 9 subclasses which are not illustrated here for brevity) to just 2 individuals for “Building Group”. The forms “Building_Group”, “Building_Unit” and “Way_Point” are modelled in the database as having the function value “Not_Applicable”. “Not_Applicable” asserts that

¹⁵ <https://github.com/perma-id/w3id.org/tree/master/geoff>

¹⁶ <https://www.w3.org/TR/dwbp/>

¹⁷ <https://eis-bonn.github.io/Luzzu/>

a function does not apply to that form, rather than simply unknown, it is modelled in Geoff as a member of the NonFunctionForm class. Units and groups are collections of buildings, therefore a single function is not applicable.

It can be seen from this table that Geoff is a more generalized representation of form and function than is represented within DLM Core. Nonetheless, Geoff presents the first vocabulary of its kind which geospatial agencies can now represent their internal geospatial data sets form and function. Except for “Division Line”, whose 9 subclasses are not studied in detail here, each Geoff class represents around 1-10% of the total instances which is an appropriate granularity, especially given the hierarchy. As noted previously, even though this is a GIS-oriented dataset, there is a rich representation of form and function values for buildings. This provides evidence for the opportunities presented by integrating form and function within the AEC domain.

Table 1. Break down of high-level form classes, number of distinct form classes and functions related to each class, and number of individual instances of forms for each class. A form with no functions associated with it (geoff:NonFunctionForm) is listed as ‘n/a’.

Geoff Class	DLM Core Table Name	Form	Function	Individuals	%
Artificial Surface	ARTIFICIAL_POLY	3	70	5,834,758	14.39
Division Line	DIVISION_LINE	10	10	22,173,787	54.67
Exposed Surface	EXPOSED_POLY	10	12	114,761	0.28
Service Line	SERVICE_LINE	2	6	329	0
Site	SITE_POLY	79	62	13,945	0.03
Structure	STRUCTURE_POLY	48	66	336,246	0.83
	STRUCTURE_PNT	29	30	962,590	2.37
	STRUCTURE_LINE	15	28	936,798	2.31
Building Single	BUILDING_POLY/PNT	52	145	3503541	8.64
Building Unit	BUILDING_UNIT_PNT_LT	5	n/a	573849	1.41
Building Group	BUILDING_GROUP_PNT	2	n/a	2	0
Water	WATER_PNT	11	5	663,984	1.64
	WATER_LINE	13	3	251,620	0.62
Way	WAY_GDF1	14	9	348,404	0.86
	WAY_GDF2	8	7	336,773	0.83
	RAIL_PNT	5	2	9,795	0.02
Vegetation	VEGETATION_POLY	12	40	4,498,521	11.09

6 Conclusions and Future Work

A Semantic Web resource for describing forms and functions of features did not exist, even though these notions are prevalent within Geospatial datasets such as the OSi. Another problem in this domain is that non-RDF standards have different perspectives

on form and function, which need to be aligned. We address this gap and challenge with the Geoff geospatial form and function vocabulary, which sets out to provide an open, freely available OWL-based spatial object classification scheme which clearly defines form and function as distinct concepts for geospatial features like manmade structures, e.g., buildings. The ontology is being used to classify form and function for the Ordnance Survey Ireland’s over 50 million spatial objects, with current development efforts focused on publishing c. 3.5 million building descriptions as Linked Data. Geoff has been developed to make these classifications available to the OSi. The Geoff vocabulary follows Linked Data publishing best practices in terms of metadata, documentation and quality assurance. Future work will explore in more detail the linking of Geoff to other standards for form and function, such as OmniClass, as well as the definition of rules (SHACL) to support data set validation.

Acknowledgment. This research has received funding from Ordnance Survey Ireland and the ADAPT Centre for Digital Content Technology, funded under the SFI Research Centres Programme (Grant 13/RC/2106) and co-funded by the European Regional Development Fund.

References

- [1] C. Stadler, J. Lehmann, K. Höffner, and S. Auer, “LinkedGeoData: A core for a web of spatial open data,” *Semant. Web*, vol. 3, no. 4, pp. 333–354, 2012.
- [2] C. Debruyne *et al.*, “Ireland’s Authoritative Geospatial Linked Data,” in *In Proceedings of The Semantic Web – ISWC – 16th International Semantic Web Conference*, 2017.
- [3] “OmniClass,” 2019. [Online]. <http://www.omniclass.org/>. [Accessed:22-Jan-2019].
- [4] “Uniclass2,” 2015. [Online]. <http://www.cpic.org.uk/uniclass/>. [Accessed:22-Jan-2019].
- [5] “ISO 12006-2:2015,” 2019. [Online]. Available: <https://www.iso.org/standard/61753.html>. [Accessed: 22-Jan-2019].
- [6] O. G. Consortium and others, “OGC GeoSPARQL-A geographic query language for RDF data,” *OGC Candidate Implement. Stand.*, vol. 2, 2012.
- [7] Ordnance Survey Ireland, “Prime2 - Data Concepts & Model Overview,” p. 17, 2014.
- [8] A. Borrmann, M. König, C. Koch, and J. Beetz, *Building Information Modeling: Technology Foundations and Industry Practice*, 10.1007/97. Springer, 2018.
- [9] L. Oyedele *et al.*, “Big Data in the construction industry: A review of present status, opportunities, and future trends,” *Adv. Eng. Informatics*, vol.30, pp.500–521, Aug. 2016.
- [10] “buildingSMART Data Dictionary Browser,” 2018. [Online]. Available: <http://bsdd.buildingsmart.org/>. [Accessed: 22-Jan-2019].
- [11] W. Terkaj and A. Šojić, “Ontology-based representation of IFC EXPRESS rules: An enhancement of ifcOWL ontology,” *Autom. Constr.*, vol. 57, pp. 188–201, Sep. 2015.
- [12] “The Protégé Ontology Editor and Knowledge Acquisition System,” 2019. [Online]. Available: <http://protege.stanford.edu/>. [Accessed: 22-Jan-2019].
- [13] S. Peroni, D. Shotton, and F. Vitali, “The Live OWL Documentation Environment: A Tool for the Automatic Generation of Ontology Documentation,” in *Knowledge Engineering and Knowledge Management*, 2012, pp. 398–412.

- [15] K. McGlinn, *et al.*, “Integrating Ireland’s Geospatial Information to Provide Authoritative Building Information Models,” in *In 13th International Conference on Semantic Systems*, 2017.
- [16] J. O’Donovan *et al.*, “A method for converting IFC geometric data into GeoSPARQL,” in *CEUR Workshop Proceedings*, 2019, vol. 2389, pp. 7–20.