






OPEN

DATA DESCRIPTOR

CODE-ACCORD: A Corpus of building regulatory data for rule generation towards automatic compliance checking

Hansi Hettiarachchi^{1,8}  , Amna Dridi^{2,8}, Mohamed Medhat Gaber², Pouyan Parsafard², Nicoleta Bocaneala², Katja Breitenfelder³, Gonçal Costa⁴, Maria Hedblom⁵, Mihaela Juganaru-Mathieu⁶, Thamer Mecharnia⁷ , Sumeer Park³, He Tan⁵, Abdel-Rahman H. Tawil² & Edlira Vakaj²

Automatic Compliance Checking (ACC) within the Architecture, Engineering, and Construction (AEC) sector necessitates automating the interpretation of building regulations to achieve its full potential. Converting textual rules into machine-readable formats is challenging due to the complexities of natural language and the scarcity of resources for advanced Machine Learning (ML). Addressing these challenges, we introduce CODE-ACCORD, a dataset of 862 sentences from the building regulations of England and Finland. Only the self-contained sentences, which express complete rules without needing additional context, were considered as they are essential for ACC. Each sentence was manually annotated with entities and relations by a team of 12 annotators to facilitate machine-readable rule generation, followed by careful curation to ensure accuracy. The final dataset comprises 4,297 entities and 4,329 relations across various categories, serving as a robust ground truth. CODE-ACCORD supports a range of ML and Natural Language Processing (NLP) tasks, including text classification, entity recognition, and relation extraction. It enables applying recent trends, such as deep neural networks and large language models, to ACC.

Background & Summary

Building codes establish regulations and standards dictating the minimum safety and welfare requirements for buildings and structures. Compliance with these codes throughout a building's lifecycle, including design, construction, and renovation/demolition, is essential to guarantee the stability, safety, usability and reliability of building designs. Although compliance checks have traditionally been conducted manually, there is now a growing motivation to automate this process due to the significant time and human resources demanded by the manual approach^{1,2}. The advancement of more semantically rich Building Information Models (BIMs) has further encouraged this tendency, making Automated Compliance Checking (ACC) achievable³. However, since building codes are written in text, as the first step, their underlying information needs to be extracted and converted into machine-readable formats to enable ACC via an intelligent approach⁴. However, extracting information from text has been a challenge due to the complexities associated with natural language, especially due to the unstructured nature and the human-centred design^{5,6}.

¹Faculty of Science and Technology, Lancaster University, Lancaster, LA1 4WA, UK. ²Faculty of Computing, Engineering and Built Environment, Birmingham City University, Birmingham, B4 7XG, UK. ³Fraunhofer Institute for Building Physics IBP, Department Indoor Climate and Climatic Impacts, Fraunhofer Str. 10, 83626, Valley, Germany. ⁴Human Environment Research (HER), La Salle, Ramon Llull University, Barcelona, Catalonia, Spain. ⁵Department of Computing, School of Engineering, Jönköping University, Box 1026, 551 11, Jönköping, Sweden. ⁶Mines Saint-Etienne, Institut Henri Fayol, Département ISI, F - 42023, Saint-Etienne, France. ⁷Université de Lorraine, CNRS, LORIA, 54506, Vandœuvre-lès-Nancy, France. ⁸These authors contributed equally: Hansi Hettiarachchi, Amna Dridi. ✉e-mail: h.hettiarachchi@lancaster.ac.uk

Early efforts in machine-readable rule formalisation focused on developing software tools and predefined procedures that domain experts could use to format textual information into rules manually. One such approach involved RASE tagging⁷, which annotates text blocks into requirements (R), applicabilities (A), selections (S), and exceptions (E), together with metadata to help formulate rules⁸. Another approach used an object-oriented procedure to filter declarative (or computer-interpretable) clauses and extract their physical entities and attributes, creating an elemental view of building regulations to facilitate easier information processing³. Similarly, a logical rule-based mechanism using noun and verb phrases was introduced for Korean building codes⁹. However, the manual conversion process was time-consuming and required the constant involvement of domain experts, prompting researchers to seek automated solutions.

Building on this trend, rule-based approaches were commonly used to extract information from regulatory text to generate machine-readable rules^{10–13}. Both semantic and syntactic features have been used to design the rules, considering the complexities associated with the text. These methods were also integrated with varied natural language processing (NLP) techniques, such as syntactic parsing that analyses text structure using domain-specific context-free grammars (CFGs) to improve accuracy^{4,14}. However, rule-based methods generally lack adaptability across different domains since most rules are specific to a particular domain¹⁵. Also, designing rules requires intensive manual efforts, which are time-consuming and error-prone and demand extensive domain expertise. To mitigate these difficulties, there has been a shift towards supervised learning-based approaches, recently. Deep learning techniques, in particular, have shown great promise for extracting information from regulatory texts, aligning with the recent trends in NLP^{16–19}. Various architectures such as Bidirectional Long Short-Term Memory (Bi-LSTM), Convolutional Neural Networks (CNN) and transformers have been adopted to extract information effectively^{4,15,20–22}. More recently, this task has also been approached as a machine translation problem, converting text into formats like LegalRuleML²³ using advanced text generation models such as T5 and BART²⁴.

While supervised learning approaches overcome some challenges of earlier methods, such as reducing the need for extensive involvement of human or domain experts and saving time, they rely heavily on high-quality annotated data for training. Various strategies have been employed to annotate regulatory data and capture key information. For instance, one study introduced a dataset in JSON format aligned with an ontology for Chinese building codes¹⁴. Another applied a similar approach to German building codes, using data-specific attributes defined by domain experts²⁵. However, these datasets are often limited to specific domains (e.g. fire protection) or locations (e.g. Vienna), as they depend heavily on predefined information structures like ontologies. Such dependencies make the annotation strategies difficult to adapt to other domains or regions, requiring significant revisions to the underlying structures. Additionally, there is a noticeable lack of openly available datasets.

To address these limitations, this study aims to develop a simple annotation strategy that can be generalised across building codes from different sub-domains and regions. Our approach focuses specifically on extracting information from natural language, which poses challenges due to its unstructured nature. On closer examination, text typically consists of two fundamental types of information: (1) entities (also referred to as named entities) and (2) relations, which are essential for understanding the conveyed ideas²⁶. An entity is a specific piece of information or a concept that can be categorised, or simply, anything that can be referred to using a proper name²⁶. For example, given the sentence “A fire door must be self-closing.”, “fire door” and “self-closing” describe entities. A relation is a semantic connection/association between an entity pair²⁶. For example, there is a “necessity” relation between the entities “fire door” and “self-closing” in the above sentence. Altogether, entities and relations form a network/knowledge graph that captures the rule(s) expressed in the text. Building on this concept, we propose an annotation strategy based on a generic set of entities and relations to convert unstructured text into machine-readable formats. Using this strategy, we annotated a dataset named CODE-ACCORD and made it publicly available to encourage the adoption of modern NLP and supervised learning techniques to advance the ACC processes.

In summary, the CODE-ACCORD corpus comprises 862 self-contained sentences extracted from the building regulations of England and Finland. A self-contained sentence is defined as a regulatory sentence that expresses a rule and contains all the details itself without any linguistical co-references that are unresolvable within the sentence, references to external sources or incomplete/ambiguous concepts. Such sentences are essential for ACC as they express rules that can be directly extracted and interpreted without extensive cross-referencing or additional context. Each sentence was manually annotated for entities and relations, with subsequent rounds of curation to ensure accuracy. Overall, CODE-ACCORD contains 4,297 annotated entities distributed across four categories and 4,329 relations distributed across ten categories.

Methods

The development of the CODE-ACCORD corpus involved two main stages: (1) data collection and (2) data annotation. Initially, sentences were carefully extracted, and then, they underwent a thorough annotation process, resulting in the final dataset.

Data collection methodology. Our data collection approach mainly focused on extracting sentences that describe rules from building regulatory data to support our ultimate goal of creating a dataset that enables the automatic generation of machine-readable rules.

Data sources. CODE-ACCORD utilised the published building regulations of England and Finland, which are openly available, as its primary/raw data sources. England’s building codes were collected from the official website of the UK Department for Levelling Up, Housing and Communities and the Ministry of Housing, Communities and Local Government (<https://www.gov.uk/government/collections/approved-documents>).

Regulations	Approved Document/Decree	#Volumes	#Pages
England	A: Structure	1	54
	B: Fire Safety	2	384
	C: Site preparation and resistance to contaminants and moisture	1	52
	D: Toxic Substances	1	10
	E: Resistance to Sound	1	86
	F: Ventilation	2	110
	G: sanitation, hot water safety and water efficiency	1	55
	H: drainage and waste disposal	1	64
	J: Combustion appliances and fuel storage systems	1	89
	K: Protection from falling, collision and impact	1	68
	L: Conservation of fuel and power	2	220
	M: Access to and use of buildings	2	143
	O: Overheating	1	44
	P: Electrical safety	1	22
	Q: Security in dwellings	1	20
	R: Infrastructure for electronic communications	2	56
	S: Infrastructure for charging electric vehicles	1	47
Material and workmanship: Approved Document 7	1	24	
Finland	Accessibility	1	6
	Fire Safety	1	25
	Energy Efficiency	1	18
	Planning and Supervision	1	7
	Strength and Stability of Structures	1	55
	Safety of Use	1	9
	Health (Indoor Climate; Water and sewerage; and Humidity)	3	16
	Acoustic Environment	1	4
Total		33	1688

Table 1. Description and Statistics of CODE-ACCORD data sources.

Finland's building codes were collected from the official website of the National Building Code of Finland from the Ministry of Environment (<https://ym.fi/en/the-national-building-code-of-finland>).

The English translation of the Finnish National Building Code was used as this work aimed to build a corpus in English. England's building regulations included the guidelines and standards that dictate the construction and maintenance of buildings in England. These regulations were designed to ensure the safety, health, and welfare of people in and around buildings, as well as to conserve fuel and power in these structures. They cover various aspects of building construction, including structural integrity, fire safety, accessibility, energy efficiency, and more. These regulations are organised into different chapters or sections, each of which addresses a specific domain or aspect of construction, such as Part A: Structure, Part B: Fire Safety, Part K: Protection from Falling and Part L: Conservation of Fuel and Power. Finland's building regulations are similar to England's regulations in terms of their purpose but may have variations in specific requirements to suit the local context. They are issued as official government decrees organised into sections or chapters, such as Accessibility, Fire Safety and Energy Efficiency, each addressing specific domains or functional requirements. These regulations mainly guide the planning, design, construction, and maintenance of buildings to meet the country's standards for safety and environmental considerations.

In both England and Finland, building regulations are published in PDF documents which are available online to the public. There is a combined total of 33 documents, consisting of 23 documents for English regulations, which span 1548 pages, and 10 documents for Finnish regulations, covering 140 pages. Table 1 presents the statistical overview of the data sources encompassing both English and Finnish regulations.

Sentence collection and processing. We processed the data from the sources mentioned above to create our initial regulatory sentence collection, following the methodology illustrated in Fig. 1. We excluded England's documents C, D, H and J from further processing due to the complexities associated with their format. Initially, all PDF documents sourced from both English and Finnish data repositories were converted into plain text. Following this, a semi-automated process was applied to filter regulatory sentences. During the filtering, we particularly focused on data that encompasses quantitative, subjective, and deontic requirements, essential for rule identification. Afterwards, the extracted sentences were manually filtered to select self-contained sentences that clearly expressed rules without relying on preceding or subsequent sentences. Each of the steps is further described below.

Document processing. This step consists of (i) extracting the textual data from the original PDF format using the PDFMiner library (<https://pypi.org/project/pdfminer/>), (ii) parsing the digital regulatory

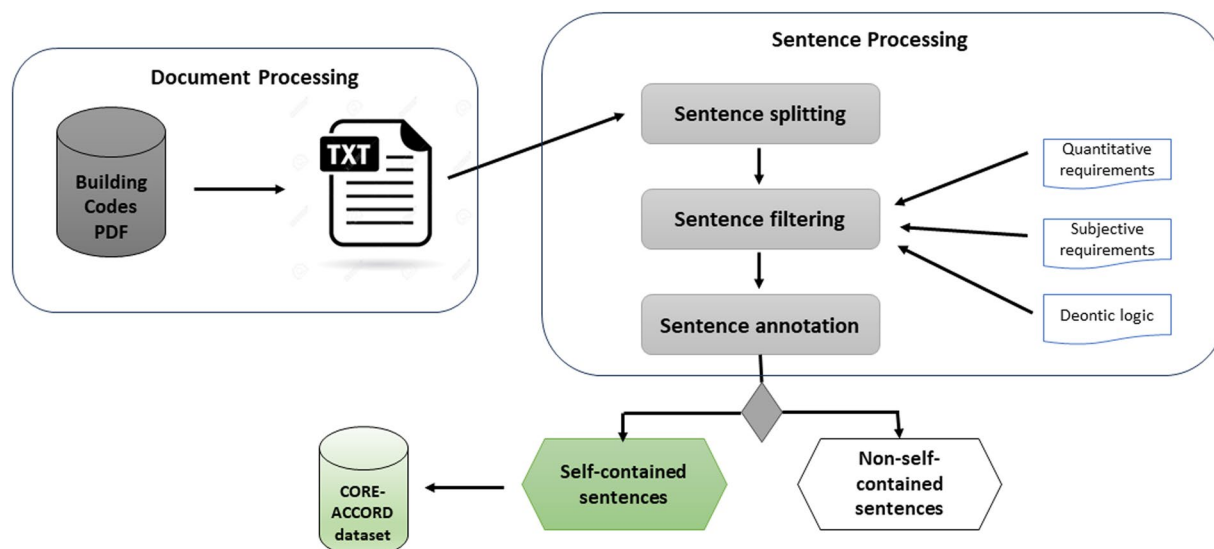


Fig. 1 The Semi-Automatic CODE-ACCORD Data Preparation Methodology.

documents by performing actions like de-hyphenation, removing line breaks and footnotes, and (iii) removing the non-convertible tables and figures, and unnecessary sections.

Sentence splitting. This step consists of splitting the regulatory text into individual sentences by analysing word sequences and punctuation marks to determine sentence boundaries. To accomplish this, we used the sentence tokeniser provided by the Natural Language Toolkit (NLTK) library (<https://www.nltk.org/>).

Sentence filtering. This step initially involved an automatic filtering process to select sentences that pertain to regulations based on three distinct features described below. The filtering was performed using a keyword-matching approach.

- **Quantitative requirements** refer to specific stipulations that can be expressed numerically or with quantitative terms. These requirements often specify precise values, measurements, or numerical criteria that must be met to ensure compliance with the regulations. Examples of quantitative requirements may include keywords such as “less than”, “greater than”, “equal”, “at least”, “higher than”, “more than”, “lower than”, followed by numerical values or thresholds. The quantitative requirements are considered since they are mostly used in building codes for describing requirements⁴.
- **Subjective requirements** are stipulations that involve the use of subjective language or expressions. These requirements are not defined by precise numerical values or measurements, but rather by language that conveys recommendations, preferences, or suggestions. Subjective requirements often include terms like “should be”, “recommended”, “preferred” or “advisable”. While subjective in nature, these requirements are important in building regulations as they allow for flexibility and adaptation to different situations while still providing a framework for best practices and quality standards. To the best of our knowledge, existing research in the field of applying NLP for the automation of building regulations has not addressed subjective requirements in their analyses, methodologies, or datasets⁴.
- **Deontic logic** pertains to the logic that deals with the expression of permissions, obligations, prohibitions, and other normative statements. It is used to represent rules and requirements that are binding or mandatory, such as rules that specify what “must”, “shall”, “could” or “prohibit” within building regulations. Deontic logic plays a crucial role in modelling the normative aspects of these regulations, providing a formal framework to represent and reason about mandatory and discretionary requirements. Similarly to the subjective requirements, deontic logic has not been extensively considered in previous research efforts. This is primarily due to the focus of most research on quantitative requirements, given their higher frequency within building regulations⁴.

Following the automatic filtering, the next step involved manual curation to ensure the accuracy of the filtered sentences. One of the authors manually reviewed and removed false positives from the automatically filtered (auto-filtered) sentences, resulting in the semi-filtered sentence set, which was then used in subsequent steps.

Sentence annotation. The final step included manual annotation of filtered sentences to identify self-contained sentences, which are defined as sentences that encompass all necessary details without any linguistic co-references that cannot be resolved within the sentence itself. Moreover, self-contained sentences should not include references to external sources, such as sections, chapters, or documents, and should avoid the inclusion of incomplete

Regulations	#Sentences	#Auto-filtered Sentences	#Semi-filtered Sentences	#Self-contained Sentences
England	19201	4219	3695	963
Finland	1473	824	764	283
Total	20674	5043	4459	1246

Table 2. Statistics of the outputs of CODE-ACCORD semi-automatic data preparation methodology.

or ambiguous concepts. Our work specifically focused on self-contained sentences because they convey rules that can be directly extracted without requiring cross-referencing or additional context that needs to be resolved by a human. This allows their information extraction and machine-readable rule formalisation to be fully automated. Non-self-contained sentences were excluded from the final regulatory sentence collection.

Sentence statistics. After applying the semi-automatic data collection methodology to the selected data sources, we obtained some noteworthy statistics summarised in Table 2. The total number of sentences was 20,674, out of which 5,043 were subjected to auto-filtering for capturing regulatory sentences, representing 24% of the total sentences in the building regulations textual content. However, during the curation process of auto-filtered sentences, an unintended consequence of the sentence-splitting approach was noticed in a few documents due to conversion errors caused by PDF formatting. Some sentences were found to be grouped together with section and subsection headers, and introductory sentences that often contain colons and semi-colons, resulting in incorrect sentence structures. To address this, manual tweaking was carried out during sentence filtering to extract the sentences carefully. The combined sentences were cross-checked against the original documents, and additional elements were removed. Following this step, the false positives in auto-filtered sentences were also removed resulting in 4,459 semi-filtered sentences. From these, 1,246 self-contained sentences were manually extracted for the annotation task described in the following section.

Data annotation methodology. In data annotation, we primarily focused on extracting information from text to facilitate automatic rule generation. There are two key types of information found in the text: named entities and relations, which are essential for comprehending the ideas conveyed in natural language²⁶. Hence, our primary focus in this work was on annotating entities and relations.

For CODE-ACCORD manual annotations, we used a group of 12 annotators with either a computer science or a civil engineering/construction background. Since this work targets the automation of compliance checking using machine learning concepts, we believe it is important to involve experts from both areas in the annotation process. All the annotators are from the ACCORD project (<https://accordproject.eu/>) and received allocated compensation for the work packages involved in this task. We used the LightTag text annotation platform²⁷ to collect human annotations, considering its coverage of different text annotations, including entities and relations, project management support and user-friendly interfaces. We provide an overview of our annotation methodology, highlighting the key concepts below. More detailed information is available in our comprehensive annotation manual (https://github.com/Accord-Project/CODE-ACCORD/blob/main/annotated_data/Annotation_Strategy_V1.0.0.pdf). Further details on annotation quality are provided in the Technical Validation section.

Entity annotation. By named entity/entity, we refer to a specific piece of information or a concept that can be categorised. Simply, named entities are anything that can be referred to using a proper name²⁶.

Entity categories. Following the idea proposed in^{4,15}, we picked four named entity categories, described in Table 3, for entity annotation. However, deviating from previous work, we adopted a simple category structure, mainly aiming at the generalisability of our annotation approach across different subdomains, such as structure, fire safety and accessibility, when defining the named entities. Also, we considered the coverage of all information in a regulatory sentence.

Annotation process. The entity annotation process mainly consisted of two steps: (1) mark entity text spans and (2) assign entity categories. A few annotated samples are shown in Table 4. As can be seen, the selected categories are versatile enough to capture all entities in different sentence structures. Also, these samples are from Accessibility and Fire Safety regulations to indicate the general applicability of our annotation strategy in different subdomains.

Relation annotation. Relations are semantic connections/associations among entities in the text²⁶. Extraction of relations together with entities is a crucial process to transform information embedded in unstructured texts into structured data formats such as knowledge graphs.

Relation categories. Due to the lack of generic, non-domain-specific relation annotation approaches in compliance checking, after carefully analysing the possible relations in the regulatory text, we identified ten relation categories, detailed in Table 5. Similar to our approach in defining entity labels, we mainly focused on the generalisability across different subdomains and coverage of semantic information when identifying the relation categories. The final category, ‘none’, is added considering the potential model requirements for identifying instances without relation between entity pairs.

Named Entity	Description
object	An ontological concept which represents a thing that is subject to a particular requirement (e.g. window, fire door)
property	Property of an object (e.g., width, height)
quality	Quality or uncountable characteristic of an object/property (e.g. horizontal, self-closing)
value	A standard or a numerical value that defines a quantity (e.g. 1,500 millimetres, five per cent)

Table 3. Entity categories.

Sentence	Annotated Sentence
The gradient of the passageway located in an outdoor space may not exceed five per cent.	The <property>gradient</property> of the <object>passageway</object> located in an <object>outdoor space</object> may not exceed <value>five per cent</value>.
There shall be a horizontal landing with a length of at least 1,500 millimetres at the lower and upper end of the ramp.	There shall be a <quality>horizontal</quality> <object>landing</object> with a <property>length</property> of at least <value>1,500 millimetres</value> at the <property>lower and upper end</property> of the <object>ramp</object>.
A fire door must be self-closing and self-bolting.	A <object>fire door</object> must be <quality>self-closing</quality> and <quality>self-bolting</quality>.

Table 4. Sample named entity annotations.

Relation	Description
selection	A limit to the scope of an object/property based on another object or a quality
necessity	A qualitative/subjective/existential necessity of an object/property (e.g., should, should have, shall be, etc.)
part-of	Being a part of an object/property
not-part-of	Not being a part of an object/property
greater	A value that should be greater than to
greater-equal	A value that should be greater than or equal to
equal	A value that should be equal to
less-equal	A value that should be less than or equal to
less	A value that should be less than to
none	No relation

Table 5. Relation categories.

Annotation process. Overall, the relation annotation process consisted of four steps: (1) mark entity text spans, (2) assign entity categories, (3) identify entity pairs which form relations, and (4) assign relation categories. This task has proven more challenging than entity annotations, primarily due to its multiple intricate steps and the potential for error propagation. However, we instructed the annotators to adhere to the entire flow and extracted entity and relation annotations upon completing the second and fourth steps, allowing them the flexibility to highlight all relevant content simultaneously. Furthermore, this approach enabled them to review the provided annotations by examining the final entity-relation representation. We only applied the manual annotation process targeting the first nine entity categories without the ‘none’ category because once all the available relations are known, the remaining possible entity pairs form the no relations. Table 6 shows a few annotated samples following the complete process.

Data Records

The CODE-ACCORD dataset is available at <https://doi.org/10.5281/zenodo.10210022>²⁸. The data repository contains three main folders: *English_Regulations*, *Finnish_Regulations* and *Annotated_Data*. The first two folders include the textual data utilised for sentence collection, and the latter contains the annotated data, as further described below.

Sentence collection. There are separate data folders, named *English_Regulations* and *Finnish_Regulations* for England’s and Finland’s regulatory text, respectively. Each folder’s hierarchy is structured as follows. Within the main data folder, there are two primary sub-folders. The first sub-folder, *PDF*, contains the original PDF files

Sentence	Entity Pair	Relation
The gradient of the passageway located in an outdoor space may not exceed five per cent.	The <property>gradient</property> of the <object>passageway</object> located in an outdoor space may not exceed five per cent.	part-of
	The gradient of the <object>passageway</object> located in an <object>outdoor space</object> may not exceed five per cent.	part-of
	The <property>gradient</property> of the passageway located in an outdoor space may not exceed <value>five per cent</value>.	less-equal
Entity-relation representation		
There shall be a horizontal landing with a length of at least 1,500 millimetres at the lower and upper end of the ramp.	There shall be a <quality>horizontal</quality> <object>landing</object> with a length of at least 1,500 millimetres at the lower and upper end of the ramp.	selection
	There shall be a horizontal <object>landing</object> with a <property>length</property> of at least 1,500 millimetres at the lower and upper end of the ramp.	part-of
	There shall be a horizontal landing with a <property>length</property> of at least <value>1,500 millimetres</value> at the lower and upper end of the ramp.	greater-equal
	There shall be a horizontal <object>landing</object> with a length of at least 1,500 millimetres at the <property>lower and upper end</property> of the ramp.	necessity
	There shall be a horizontal landing with a length</property> of at least 1,500 millimetres at the <property>lower and upper end</property> of the <object>ramp</object>.	part-of
Entity-relation representation		

Table 6. Sample relation annotations.

of the regulations documents. The second sub-folder, *Text and CSV*, is where TXT and CSV versions of the PDF files are stored after undergoing various pre-processing stages. This *Text and CSV* sub-folder consists of eight sub-folders, each corresponding to a specific pre-processing step. They are meticulously organised in sequential order to facilitate systematic data handling. The first sub-folder, *RawTextData*, contains the initial raw text data obtained through PDF-to-text conversion. The subsequent sub-folder, *CleanedData-RawText*, holds the cleaned data derived from the initial raw text. *AllSentences* sub-folder contains all sentences extracted from the cleaned text. Next, *AutoFilteredSentences* comprises sentences that have been automatically filtered, following the specific criteria described in the Data Collection Methodology above. *ManuallyFilteredSentences* contains manually curated sentences to ensure consistency and remove unnecessary content. The *FinalFilteredSentences* sub-folder

Attribute	Description
example_id	Unique ID assigned for each sentence
content	Original textual content of the sentence
processed_content	Tokenised (using NLTK's wordtokenize package) textual content of the sentence
label	Entity labelled sequence in IOB format
metadata	Additional information of sentence (i.e. original approved document from which the sentence is extracted)

Table 7. Format of entity data file.

Attribute	Description
example_id	Unique ID assigned for each sentence
content	Original textual content of the sentence
metadata	Additional information of sentence (i.e. original approved document from which the sentence is extracted)
tagged_sentence	Sentence with tagged entity pair
relation_type	Category of the relation in between the tagged entity pair

Table 8. Format of relation data file.

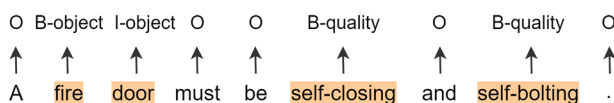


Fig. 2 Sample of entity labels in BIO format.

stores the ultimate raw text of the semi-automatically filtered sentences after eliminating empty lines and redundant information. Moving on, *CSV-FinalFilteredSentences* presents the sentences from the *FinalFilteredSentences* folder in CSV format, preparing the data for the final sub-folder, *Classification* which categorises the sentences as either “self-contained” or “others”, where only the self-contained sentences were considered for final data annotation. This structured hierarchy streamlines the data processing and analysis of regulations documents, ensuring an organised and efficient workflow.

Annotated data. A randomly selected 862 sentences out of 1246 self-contained sentences extracted from the building regulations of England and Finland underwent our comprehensive data annotation process, targeting both entities and relations, which are essential for extracting information from text. All annotated data are available in *Annotated_Data* folder. There are two sub-folders named *entities* and *relations* within it, which hold entity-annotated data and relation-annotated data, respectively. Each folder has three CSV files named *all.csv*, *train.csv* and *test.csv*. File *all.csv* contains the full dataset. *train.csv* has 80% of the full dataset, which can be used to train machine learning models, and *test.csv* has the remaining 20%, which can be used for models’ performance testing. All three files within one folder follow the same format. The entity and relation data file formats are further described below.

Entity data format. The format of an entity data file available within the *entities* folder is summarised in Table 7. The entity annotations are available in the BIO (Beginning, Inside, Outside) format, which is considered the standard for information extraction tasks²⁹, as shown in the sample in Fig. 2. *object* and *quality* tags shown in the figure are two categories of the selected entity categories, which are further described in the Data Annotation Methodology above.

Relation data format. The format of a relation data file available within the *relations* folder is summarised in Table 8. We adopted the following format to tag the entity pair denoting the relationship within a data sample, in accordance with formats utilised in recent studies^{18,30}:

The <e1>gradient</e1> of the <e2>passageway</e2> located in an outdoor space may not exceed five per cent.

The special tags <e1> and </e1> represent the start and end of the first entity that appeared in the sentence. Similarly, <e2> and </e2> represent the second entity.

Descriptive statistics. The statistical analysis of the final annotated dataset is vital for future efforts that will use CODE-ACCORD as a resource for automating the conversion of textual rules into machine-readable formats, facilitating Automated Compliance Checking (ACC).

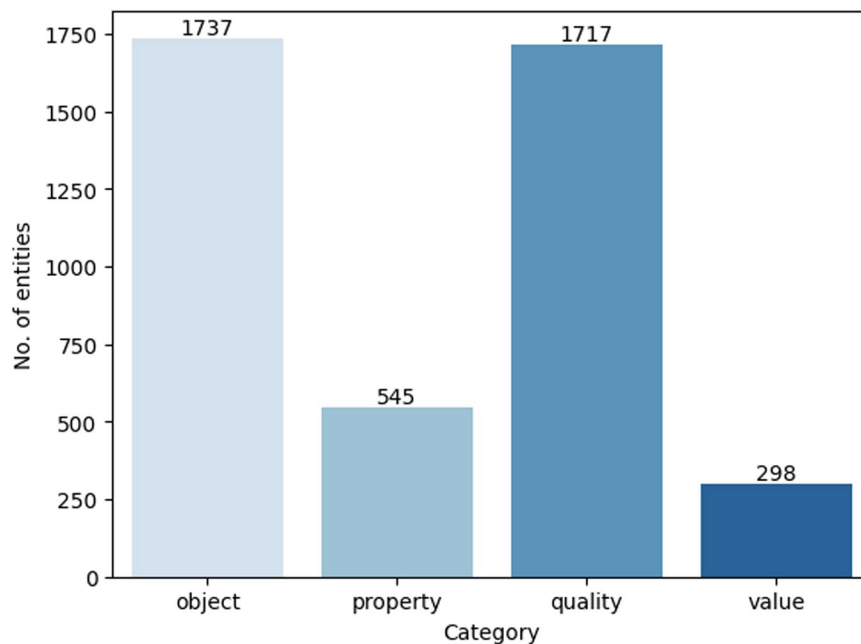


Fig. 3 Distribution of entity categories.

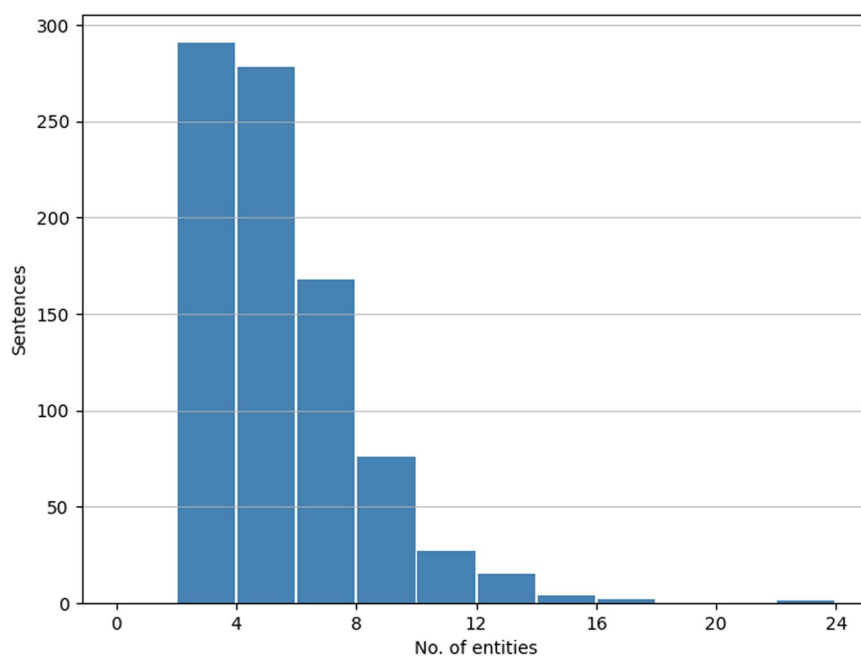


Fig. 4 Distribution of the number of entities per sentence.

Entity statistics. Our final entity annotated dataset contains 862 sentences. It has 4,297 entities distributed over four categories, as shown in Fig. 3. The illustrated distribution of the number of entities per sentence in Fig. 4 provides a detailed insight into the annotated data. As can be seen, most sentences contain up to five entities. We further analysed the sequence lengths of text spans from each category, and the resulting histograms are presented in Fig. 5. Most of the entity spans are composed of one or two words/tokens. However, overall, there are more lengthy text spans under *quality* than in the other categories.

Relation statistics. Altogether, we annotated 862 sentences, resulting in 3,329 relations. We automatically identified the unannotated entity pairs within sentences as unrelated entity pairs which belong to the ‘*none*’ category. Out of 8,104 samples categorised as ‘*none*’, we included a random subset of 1,000 in our final dataset to ensure a balanced distribution with other relations. The breakdown of a total of 4,329 relations across ten categories is

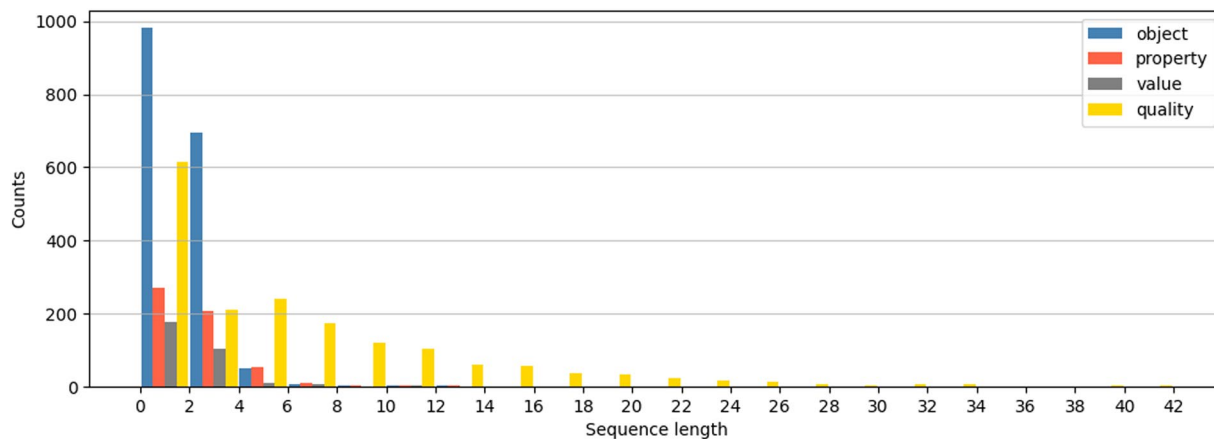


Fig. 5 Sequence length distribution of annotated text spans as entities.

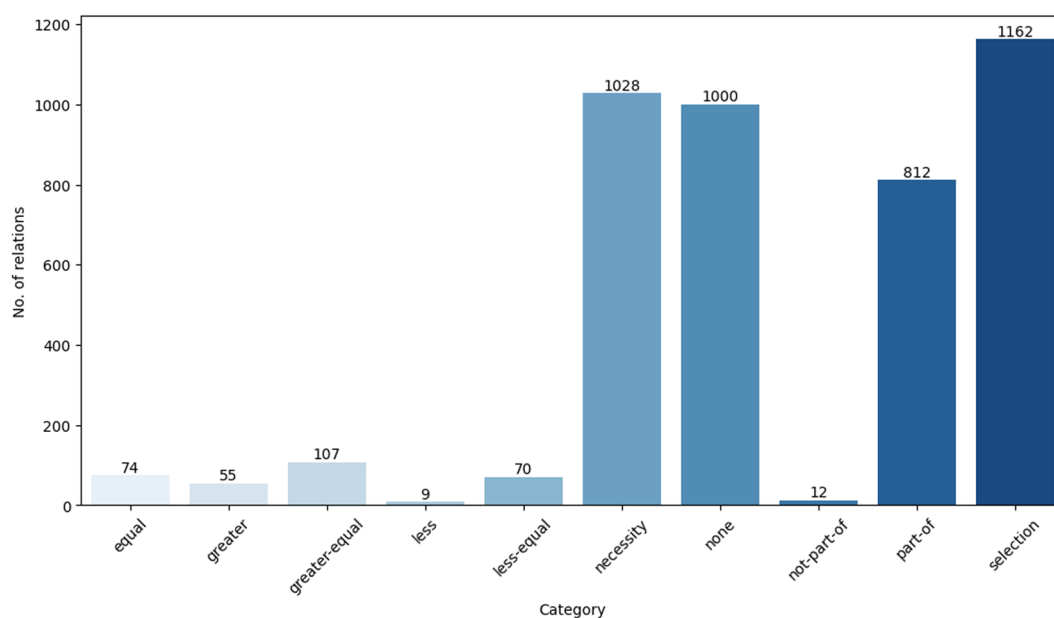


Fig. 6 Distribution of relation categories.

depicted in Fig. 6. Additionally, Fig. 7 illustrates the distribution of the number of relations per sentence. Most sentences contained two or three relations, although a minority had over ten relations.

Technical Validation

Annotation quality. We employed several methods to ensure the quality of the annotations. The annotation strategy was initially designed based on two fundamental types of information in the text (i.e. entities and relations) along with available methods. It was then refined through multiple iterations, incorporating feedback from several domain experts. Additionally, we conducted two rounds of test annotations with a team of two annotators familiar with the strategy to validate its feasibility and coverage before finalising it.

We conducted the annotations for both entities and relations in multiple rounds. The first round served as a training session for the annotators. They were given sufficient time to review the annotation strategy and examples, with additional clarifications provided upon request before beginning the round. After this round, a question-and-answer session was held to address any queries with the help of several domain experts. The actual annotation rounds then began, with a total of seven rounds. Each sentence was independently annotated by two or three annotators through these rounds.

To measure the quality of annotations, we calculated Inter Annotator Agreements (IAAs) throughout our rounds. As the entity IAA, we used the pairwise relative agreement of entities. The annotator A's agreement with Annotator B is calculated using Eq. 1²⁷. An entity annotated by one annotator is considered a match to an entity by another annotator only if the marked text span and assigned label are equal. Figure 8 summarises the distribution of entity IAA values obtained throughout the annotation rounds. The mean IAA is 0.37, with a maximum

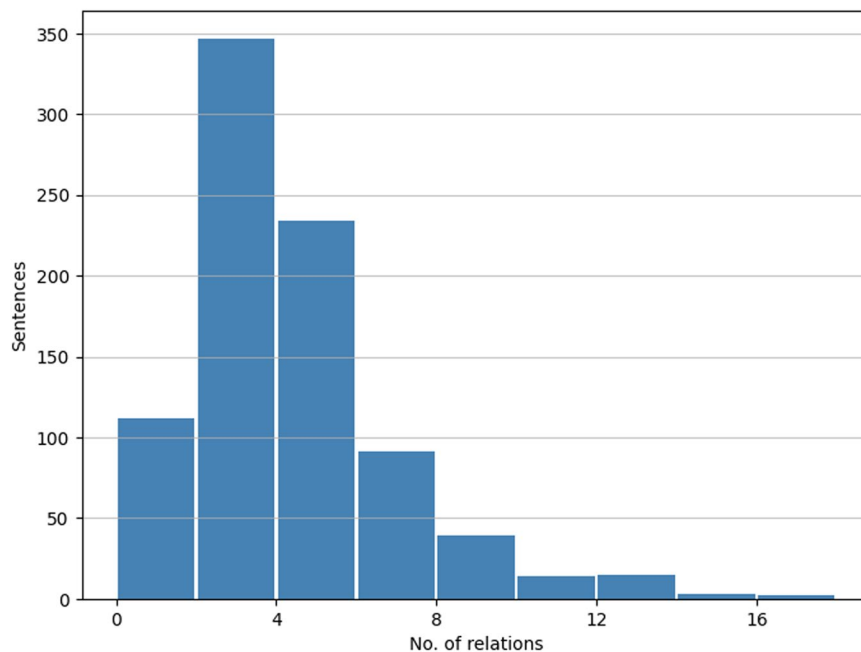


Fig. 7 Distribution of the number of relations per sentence.

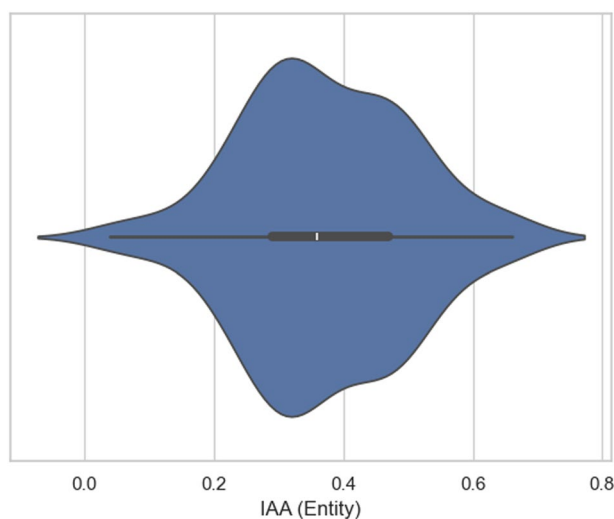


Fig. 8 Distribution of entity IAA values.

of 0.66. The task's complexities (i.e. two-step annotation process and domain-specific knowledge requirements) and strict matching criteria used during the agreement calculation can be identified as the primary factors contributing to this distribution. However, given the further complexities associated with relation annotations following its four-step process, we limited our IAA calculations only to entities.

$$\text{A's agreement with B} = \frac{\text{Number of B's entities matched with A's entities}}{\text{Total entitie sannoted by B}} \quad (1)$$

To enhance the accuracy of the annotations, each annotation round was followed by a curation round to determine the final annotations. Three members from the annotation group, who are domain experts, served as data curators. Their curation jobs were carefully assigned without any overlaps with the annotation jobs. During curation, the final annotations for all entities and relations with disagreements between annotators have been decided by the curator, considering the proposed annotations and the overall entity-relation representation of each sentence. Additionally, the content of each sentence was cross-checked against the final entity-relation representation or knowledge graph to confirm its completeness and to verify that all relevant information in the text was accurately captured.

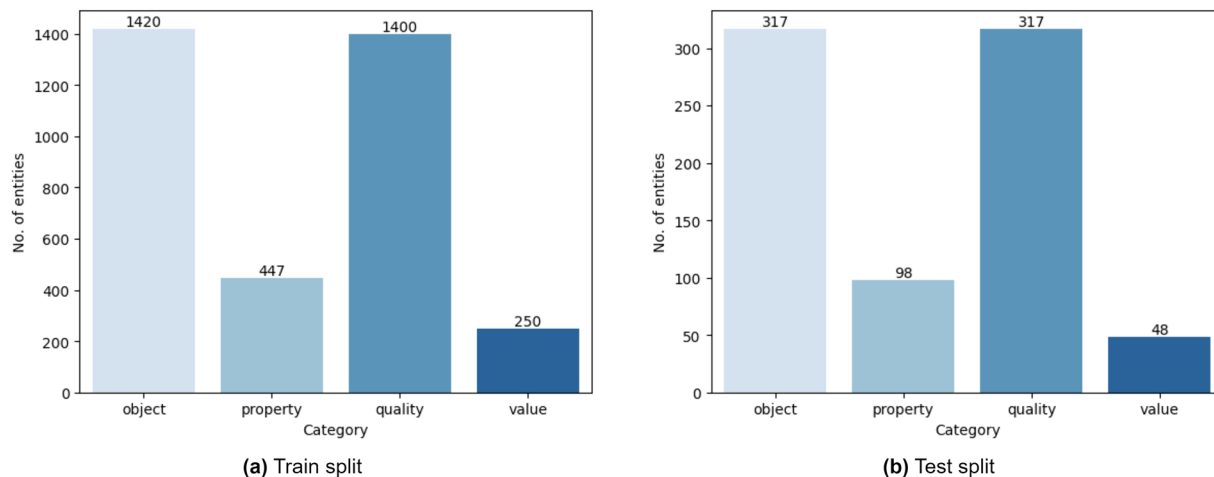


Fig. 9 Distribution of the entity categories in train and test data splits.

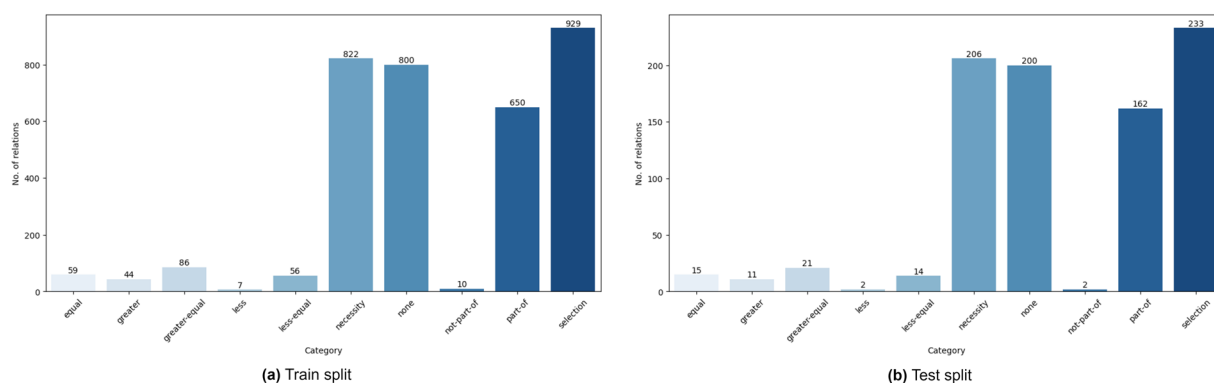


Fig. 10 Distribution of the relation categories in train and test data splits.

Data splits quality. Alongside the complete annotated dataset, CODE-ACCORD offers two data splits: train and test, for each annotation type (i.e. entities and relations). Each train split comprises 80% of the corresponding full dataset and is intended for training machine learning models. The remaining 20% forms the test splits, designed for evaluating model performance. It is essential to ensure an equivalent class distribution across data splits, as it directly affects the model training and evaluation procedures and the overall accuracy of the final model. Thus, we used stratified sampling to create these splits, ensuring that the class distribution in each split mirrors that of the original data. We further validated this fact using exploratory data analysis resulting in Figs. 9, 10, which illustrate the distribution of entity and relation categories in the train and test data splits, respectively, demonstrating the consistency of distributions across these splits.

Code availability

CODE-ACCORD GitHub repository is available at <https://github.com/Accord-Project/CODE-ACCORD/>.

Received: 31 July 2024; Accepted: 17 December 2024;

Published online: 29 January 2025

References

- Fuchs, S. & Amor, R. Natural language processing for building code interpretation: A systematic literature review. *Proceedings of the Conference CIB W78* **2021**, 11–15 (2021).
- Zhang, Z., Ma, L. & Nisbet, N. Unpacking ambiguity in building requirements to support automated compliance checking. *Journal of Management in Engineering* **39**, 04023033, <https://doi.org/10.1061/JMENEA.MEENG-5359> (2023).
- Malsane, S., Matthews, J., Lockley, S., Love, P. E. & Greenwood, D. Development of an object model for automated compliance checking. *Automation in Construction* **49**, 51–58, <https://doi.org/10.1016/j.autcon.2014.10.004> (2015).
- Zhou, Y.-C., Zheng, Z., Lin, J.-R. & Lu, X.-Z. Integrating nlp and context-free grammar for complex rule interpretation towards automated compliance checking. *Comput. Ind.* **142**, <https://doi.org/10.1016/j.compind.2022.103746> (2022).
- Zhang, J. & El-Gohary, N. Extraction of construction regulatory requirements from textual documents using natural language processing techniques. In *Congress on Computing in Civil Engineering, Proceedings*, 453–460, <https://doi.org/10.1061/9780784412343.0057> (2012).
- Locatelli, M., Seghezzi, E., Pellegrini, L., Tagliabue, L. C. & Di Giuda, G. M. Exploring natural language processing in construction and integration with building information modeling: A scientometric analysis. *Buildings* **11**, <https://doi.org/10.3390/buildings1120583> (2021).

7. Hjelseth, E. & Nisbet, N. Capturing normative constraints by use of the semantic mark-up rase methodology. In *Proceedings of CIB W78-W102 Conference*, 1–10 (2011).
8. Beach, T., Rezgui, Y., Li, H. & Kasim, T. A rule-based semantic approach for automated regulatory compliance in the construction sector. *Expert Syst. Appl.* **42**, 5219–5231, <https://doi.org/10.1016/j.eswa.2015.02.029> (2015).
9. Lee, H., Lee, J.-K., Park, S. & Kim, I. Translating building legislation into a computer-executable format for evaluating building permit requirements. *Automation in Construction* **71**, 49–61, <https://doi.org/10.1016/j.autcon.2016.04.008>. The Special Issue of 32nd International Symposium on Automation and Robotics in Construction (2016).
10. Zhang, J. & El-Gohary, N. M. Semantic nlp-based information extraction from construction regulatory documents for automated compliance checking. *Journal of Computing in Civil Engineering* **30**, 04015014, [https://doi.org/10.1061/\(ASCE\)CP.1943-5487.0000346](https://doi.org/10.1061/(ASCE)CP.1943-5487.0000346) (2016).
11. Zhang, J. & El-Gohary, N. M. Integrating semantic nlp and logic reasoning into a unified system for fully-automated code checking. *Automation in Construction* **73**, 45–57, <https://doi.org/10.1016/j.autcon.2016.08.027> (2017).
12. Zhou, P. & El-Gohary, N. Ontology-based automated information extraction from building energy conservation codes. *Automation in Construction* **74**, 103–117, <https://doi.org/10.1016/j.autcon.2016.09.004> (2017).
13. Xue, X. & Zhang, J. Building codes part-of-speech tagging performance improvement by error-driven transformational rules. *Journal of Computing in Civil Engineering* **34**, 04020035, [https://doi.org/10.1061/\(ASCE\)CP.1943-5487.0000917](https://doi.org/10.1061/(ASCE)CP.1943-5487.0000917) (2020).
14. Zheng, Z., Zhou, Y.-C., Lu, X.-Z. & Lin, J.-R. Knowledge-informed semantic alignment and rule interpretation for automated compliance checking. *Automation in Construction* **142**, 104524, <https://doi.org/10.1016/j.autcon.2022.104524> (2022).
15. Zhang, R. & El-Gohary, N. A deep neural network-based method for deep information extraction using transfer learning strategies to support automated compliance checking. *Automation in Construction* **132**, 103834, <https://doi.org/10.1016/j.autcon.2021.103834> (2021).
16. Shen, Y. *et al.* Parallel instance query network for named entity recognition. In Muresan, S., Nakov, P. & Villavicencio, A. (eds.) *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 947–961, <https://doi.org/10.18653/v1/2022.acl-long.67> (Association for Computational Linguistics, Dublin, Ireland, 2022).
17. Shen, Y. *et al.* DiffusionNER: Boundary diffusion for named entity recognition. In Rogers, A., Boyd-Graber, J. & Okazaki, N. (eds.) *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 3875–3890, <https://doi.org/10.18653/v1/2023.acl-long.215> (Association for Computational Linguistics, Toronto, Canada, 2023).
18. Plum, A., Ranasinghe, T., Jones, S., Orasan, C. & Mitkov, R. Biographical semi-supervised relation extraction dataset. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '22*, 3121–3130, <https://doi.org/10.1145/3477495.3531742> (Association for Computing Machinery, New York, NY, USA, 2022).
19. Yang, S., Choi, M., Cho, Y. & Choo, J. HistRED: A historical document-level relation extraction dataset. In Rogers, A., Boyd-Graber, J. & Okazaki, N. (eds.) *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 3207–3224, <https://doi.org/10.18653/v1/2023.acl-long.180> (Association for Computational Linguistics, Toronto, Canada, 2023).
20. Wang, X. & El-Gohary, N. *Deep Learning-Based Named Entity Recognition from Construction Safety Regulations for Automated Field Compliance Checking*, 164–171 <https://ascelibrary.org/doi/pdf/10.1061/9780784483893.021> (2021).
21. Zheng, Z., Lu, X.-Z., Chen, K.-Y., Zhou, Y.-C. & Lin, J.-R. Pretrained domain-specific language model for natural language processing tasks in the aec domain. *Computers in Industry* **142**, 103733, <https://doi.org/10.1016/j.compind.2022.103733> (2022).
22. Wang, X. & El-Gohary, N. Deep learning-based relation extraction and knowledge graph-based representation of construction safety requirements. *Automation in Construction* **147**, 104696, <https://doi.org/10.1016/j.autcon.2022.104696> (2023).
23. Athan, T. *et al.* OASIS LegalRuleML. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Law, ICAIL '13*, 3–12, <https://doi.org/10.1145/2514601.2514603> (Association for Computing Machinery, New York, NY, USA, 2013).
24. Fuchs, S., Witbrock, M., Dimyadi, J. & Amor, R. Neural semantic parsing of building regulations for compliance checking. In *IOP Conference Series: Earth and Environmental Science*, vol. **1101**, 092022, <https://doi.org/10.1088/1755-1315/1101/9/092022> (IOP Publishing, 2022).
25. Recski, G., Iklódi, E., Lellmann, B., Kovács, Á. & Hanbury, A. BRISE-Plandok: A german legal corpus of building regulations. *Language Resources and Evaluation* 1–40, <https://doi.org/10.1007/s10579-024-09747-7> (2024).
26. Jurafsky, D. & Martin, J. H. *Speech and Language Processing (2Nd Edition)* (Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 2009).
27. Perry, T. LightTag: Text annotation platform. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 20–27 (Association for Computational Linguistics, Online and Punta Cana, Dominican Republic, 2021).
28. Hettiarachchi, H. *et al.* Accord-project/CODE-ACCORD: v1.0.0. *Zenodo* <https://doi.org/10.5281/zenodo.10210022> (2023).
29. Ramshaw, L. & Marcus, M. Text chunking using transformation-based learning. In *Third Workshop on Very Large Corpora* (1995).
30. Bastos, A. *et al.* RECON: Relation extraction using knowledge graph context in a graph neural network. In *Proceedings of the Web Conference 2021, WWW '21*, 1673–1685, <https://doi.org/10.1145/3442381.3449917> (Association for Computing Machinery, New York, NY, USA, 2021).

Acknowledgements

This work is funded by the European Union's Horizon Europe research and innovation programme under grant agreement no 101056973 (ACCORD). UK Participants in Horizon Europe Project [ACCORD] are supported by UKRI grant numbers [10040207] (Cardiff University), [10038999] (Birmingham City University) and [10049977] (Building Smart International).

Author contributions

H.H.: Conceptualisation, Data preparation, Data annotation methodology, Data annotation, Data curation, Validation, Visualisation, Project administration, Writing - Original Draft, Writing - Review & Editing. A.D.: Conceptualisation, Data preparation methodology, Data preparation, Formal analysis, Data annotation methodology, Data annotation, Data curation, Writing - Original Draft. M.M.G.: Conceptualisation, Data annotation methodology, Supervision, Writing - Review & Editing. P.P.: Data annotation, Data curation, Writing - Review & Editing. N.B.: Data annotation, Writing - Review & Editing. K.B.: Data annotation, Writing - Review & Editing. G.C.: Data annotation, Writing - Review & Editing. M.H.: Data annotation, Writing - Review & Editing. M.J.-M.: Data annotation, Writing - Review & Editing. T.M.: Data annotation, Writing - Review & Editing. S.P.: Data annotation, Writing - Review & Editing. H.T.: Data annotation, Writing - Review & Editing. A.-R.H.T.: Data annotation, Writing - Review & Editing. E.V.: Conceptualisation, Data annotation methodology, Project administration, Funding acquisition, Writing - Review & Editing.

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to H.H.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2025