

# Formalizing Tag-Based Metadata with the Brick Ontology

Gabe Fierro<sup>1,\*</sup>, Jason Koh<sup>2</sup> Shreyas Nagare<sup>3</sup> Xiaolin Zang<sup>3</sup> Yuvraj Agarwal<sup>3</sup>  
Rajesh K. Gupta<sup>2</sup> and David E. Culler<sup>1</sup>

<sup>1</sup>Computer Science, UC Berkeley, Berkeley, CA, USA

<sup>2</sup>Computer Science, UC San Diego, San Diego, CA, USA

<sup>3</sup>Computer Science, Carnegie Mellon University, Pittsburgh, PA, USA

Correspondence\*:

Gabe Fierro

gtfierro@cs.berkeley.edu

## 2 ABSTRACT

3 Current efforts establishing semantic metadata standards for the built environment span  
4 academia, industry and standards bodies. For these standards to be effective, they must be  
5 clearly defined and easily extensible, encourage consistency in their usage, and integrate cleanly  
6 with existing industrial standards, such as BACnet. There is a natural tension between informal  
7 tag-based systems that rely upon idiom and convention for meaning, and formal ontologies  
8 amenable to automated tooling.

9 We present a qualitative analysis of Project Haystack, a popular tagging system for building  
10 metadata, and identify a family of inherent interpretability and consistency issues in the tagging  
11 model that stem from its lack of a formal definition. To address these issues, we present the  
12 design and implementation of the Brick+ ontology, a drop-in replacement for Brick with clear  
13 formal semantics that enables the inference of a valid Brick model from an informal Haystack  
14 model, and demonstrate this inference across five Haystack models.

15 **Keywords:** Smart Buildings, Building Management, Metadata, Ontologies, OWL, RDF, Brick, Haystack

## 1 INTRODUCTION

16 Smart buildings have long been a target of efforts aiming to reduce energy consumption, improve occupant  
17 comfort, and increase operational efficiency. Although a substantial body of work advances the state-of-the-  
18 art — including automated control [39, 10, 31], modeling [32] and analysis [38, 20] — such approaches  
19 do not see widespread use due to the prohibitive cost of configuring their instantiation to each building. A  
20 major factor in this cost is due to lack of interoperability standards; without such standards, the rollout of  
21 energy efficiency measures involves customizing implementations to the one-off combinations of hardware  
22 and software configurations that are unique to each building. Limited deployment of energy efficiency  
23 applications constrains the ability to evaluate potential savings [26]. Recent studies by the US Department  
24 of Energy [27, 18] have established that a lack of interoperability standards for buildings reduces the  
25 cost-effectiveness and scalability of energy efficiency techniques and analyses.

26 Semantic metadata standards present a promising path to enabling interoperability by offering uniform  
27 descriptions of building resources to application developers and building operators. Today, semantic  
28 metadata standardization efforts for buildings span academia [4], industry [1, 37] and standards bodies [35,  
29 3]. As applications developed for the built environment have become increasingly data-focused, recent  
30 metadata standard efforts have shifted from supporting the initial construction and commissioning phases  
31 of operation to enabling robust descriptions of the provenance and context of collected data.

### 32 1.1 Brick and Haystack Metadata Systems

33 Emerging data-oriented metadata standards differ in their support for *consistent* and *extensible* use.  
34 De-facto industrial metadata practices have embraced unstructured vendor- and building-specific idioms  
35 intended for human consumption rather than programmatic manipulation. Several standardization efforts  
36 have arisen to address the ad-hoc nature of building metadata. Of these, Brick [4] and Project Haystack [1]  
37 have seen adoption and investment from academic and industrial sources, and are involved in the ASHRAE  
38 223P effort to standardize semantic tagging for building data [3].

39 Project Haystack is a commonly-used open building metadata standard that replaces unstructured labels  
40 with semi-structured sets of tags<sup>1</sup>. However, the informal and ad-hoc composition of these tags precludes  
41 consistent usage; this leaves interpretation of tags up to the tacit knowledge of domain experts.

42 Brick is a recently introduced metadata standard designed for completeness (describing all of the relevant  
43 concepts required for applications), expressiveness (capturing the explicit and implicit relationships required  
44 for applications) and usability (fulfilling the needs of domain experts and application developers). Although  
45 evaluations of Brick demonstrate its ability to robustly capture a wide variety of application requirements,  
46 the story of how Brick integrates with existing tooling and industrial practices, such as Haystack, has been  
47 less clear.

48 Put simply, Brick and Haystack serve different goals. Brick is designed for the complete and consistent  
49 modeling of concepts required for developing portable software that can be deployed at scale. Haystack is  
50 designed for building managers and engineers who need *familiar idioms* for developing and using software  
51 designed to function on a small number of buildings. However, these informal practices are not sufficient for  
52 the large-scale standardization of *consistent* semantic metadata necessary for the widespread deployment  
53 of energy efficiency applications. Consistent metadata requires a set of rules formalizing how metadata can  
54 be *defined, structured, composed* and *extended*.

55 In this paper, we present the design and implementation of *Brick+*, a drop-in replacement for the Brick  
56 ontology with clear formal semantics designed for the sensible composition of concepts required for  
57 portable building applications. The key design principle of *Brick+* is the choice to *model concepts in terms*  
58 *of the formal composition of their properties*. This is more expressive than the original Brick class hierarchy  
59 which captures *specialization*, but not behavior. *Brick+* enables the inference of properties beyond what can  
60 be captured by tag-based metadata schemes or the original Brick schema, including modeling the behavior  
61 of equipment and points and formalizing Haystack models.

62 Ultimately, the formal representation of metadata enables a greater degree of consistency and consistency  
63 on behalf of the model, while enabling a family of supporting tooling that facilitate the production of *Brick+*  
64 models from existing tag-based metadata, the systematic validation of those models, and the migration of  
65 existing Brick models to the proposed *Brick+*.

---

<sup>1</sup> Referred to as “Haystack” in the rest of the paper

## 66 1.2 Overview

67 §3 presents an analysis of the systemic interpretability and consistency issues endemic to the Haystack  
68 metadata system, motivating the need for formal rules for composition. This is one of the first systematic  
69 evaluations of Project Haystack’s approach to metadata: how it impacts consistency and to what extent  
70 it enables or inhibits semantic interoperability. §4 presents the design of Brick+, a drop-in replacement  
71 for Brick with clear formal semantics. Brick+ defines a class lattice that structures the composition of  
72 concepts. This lattice enables Brick+ to define inference from Haystack’s informal tags to formal Brick  
73 classes. §5 presents the implementation of Brick+ using the OWL-DL ontology language and defines the  
74 process by which a Brick model can be inferred from a set of tagged Haystack entities. §8 evaluates the  
75 Brick+ ontology and inference methodology by observing the accuracy of classifying entities from five  
76 Haystack models to Brick+, and examining the additional properties that can be inferred by Brick+ over  
77 104 existing Brick models. §9 summarizes ongoing and future efforts to integrate the Brick and Haystack  
78 metadata standards and concludes.

79 Since publication of [13], Brick+ has been adopted into the release of Brick v1.1. This paper extends  
80 [13] with:

- 81 1. a deeper discussion of the challenges in formalizing metadata tags, and how the implementation of  
82 Brick+ resolves these issues (§5.2)
- 83 2. the design and implementation of a tool for validating usage of the Brick ontology using SHACL, and  
84 techniques for defining templates and idioms that assist in Brick usability (§6)
- 85 3. the design and implementation of a tool for migrating Brick models from older versions of the Brick  
86 ontology to newer versions (§7)

87 The production and evaluation of the above Brick+ tooling validates the choice of a formalized semantic  
88 metadata model. Not only is the tooling straightforward to construct given the Python-based implementation  
89 and formal construction, it also presents an opportunity to unify the Brick and Haystack metadata standards  
90 beyond a fragile “house of sticks” constructed from idiom and convention.

## 2 BACKGROUND

91 We define a set of concepts for later use, provide an overview of the Brick and Haystack metadata models,  
92 and discuss how Brick+ fits into the existing body of literature.

### 93 2.1 Definitions

94 We refer to the following terms throughout the paper:

- 95 • A **tag** is an atomic fact or attribute; tags may or may not be associated with a value.
- 96 • A **tag set** is an unordered collection of tags associated with an entity.
- 97 • A **valid tag set** is a tag set with a clear, real-world definition.
- 98 • An **entity** is an abstraction of a physical, logical or virtual item.
- 99 • A **class** is a category of entities defined by a particular shared purpose and properties.

100 In Brick and Brick+, classes are organized by the subclass and superclass relationships between classes.  
101 This approach organizes classes naturally in terms of more specific or more general concepts. For example,  
102 the class of “sensors” is more general than the class of “temperature sensors” (sensors that measure the  
103 temperature property of some substance) and the class of “air sensors” (sensors that measure properties  
104 of air), which are both more general than the class of “air temperature sensors” (sensors that measure

105 the temperature property of air). In Brick, `Air Temperature Sensor` is the class of all entities that  
106 measure the temperature of air.

## 107 2.2 Haystack

108 Haystack defines entities as a set of *value* tags (representing key-value pairs) and *marker* tags (singular  
109 annotations). Value tags define attributes of entities such as name, timezone, units and data type. Ref tags  
110 are a special kind of value tag that refer to other Haystack entities. These hint at relationships, but are  
111 entirely generic; the relationship is understood by convention. Haystack provides a dictionary of defined  
112 tags on its website [1]. The set of marker tags for an entity constitute the “tag set” for that entity and  
113 construe the concept of which the entity is an example (its “type”).

## 114 2.3 Brick

115 The Brick ontology has two components: an extensible *class hierarchy* representing the physical and  
116 logical entities in buildings, and a minimal set of *relationships* that capture the connections between  
117 entities. A Brick model of a building is a labeled, directed graph in which the nodes are entities and  
118 the edges are relationships. Brick is defined using the Resource Description Framework (RDF) data  
119 model [24], which represents graph-based knowledge as tuples of (subject, predicate, object)  
120 termed *triples*. A triple states that a subject entity has a relationship (predicate) to an object entity. Line 2  
121 of Figure 2 is a triple for which the subject is `:sensor1`, the predicate (relationship) is `a`, and the object  
122 is `brick:Temperature_Sensor`.

123 Brick and Brick+ are both defined with the RDFS [17] and OWL [7] knowledge representation languages.  
124 These languages allow the expression of rules and constraints for authoring ontologies, which can be  
125 interpreted by a *semantic reasoner* such as Hermit [16] to materialize inferred triples from a set of input  
126 triples. Brick+’s use of a semantic reasoner is covered in §5.

## 127 2.4 Prior Metadata Construction Efforts

128 Several other works conduct inference and classification to extract structure from unstructured building  
129 metadata, including leveraging semi-supervised learning approaches to learn parsing rules for unstructured  
130 labels [9, 22], classifying sensors by examining historical timeseries data [19, 15], or by combining  
131 timeseries analysis with label clustering [6]. These efforts are largely complementary to Brick+. Brick+  
132 defines formal methods for inference-based classification of tagged entities, but requires external support  
133 for extracting tagged entities from unstructured metadata.

134 Beyond the building domain, there is a family of work [28, 29] using ontologies to provide structure  
135 to tag-based folksonomies [25]. The approaches developed in these works, along with work on formal  
136 concept analysis [43] and concept lattices [42], form the theoretical basis for Brick+.

## 137 2.5 Relation to Building Information Modeling

138 Building information modeling (BIM) relates to the data exchanged for the design, construction and  
139 commissioning of a building. BIM models contain extensive lists of building assets in addition to 3D  
140 geometry, and there is active research into extending the use of BIM for operation and maintenance [40, 44].  
141 However, BIM models lack direct representation of the contextual metadata described by Brick and  
142 Haystack [8, 23]. For example, a BIM model can represent a fan and describe its physical properties such as  
143 the shape of the blades, but does not explicitly label whether the fan is installed on the supply or return side  
144 of an HVAC system. Deriving that information requires traversing the complex objects and relationships  
145 that describe the ducts, connectors and other components of the HVAC system [12], which is difficult  
146 to do in an automated manner. Despite these difficulties in retrieving the contextual metadata required  
147 to run data-driven applications, BIM is largely complementary to Brick and Haystack. Recent work in

148 representing BIM models using an OWL-based ontology [30] will enable well-defined mappings between  
149 the Brick+ and ifcOWL ontologies.

### 3 SYSTEMIC TAG ISSUES IN HAYSTACK

150 Although Haystack models have seen increasing adoption, the design of the Haystack data model has several  
151 intrinsic issues that limit its consistency and interpretability. Here, we present one of the first analyses of  
152 the Haystack data model and how its tag-based implementation impacts consistency, interpretability and  
153 interoperability.

#### 154 3.1 Lack of Formal Class Hierarchy

155 A well-formed class hierarchy organizes concepts by their specificity. This is essential for the creation of  
156 consistent metadata models because it facilitates automated discovery of classes by way of traversing the  
157 hierarchy for more general or more specific concepts. In the process of identifying an appropriate class  
158 for an entity, a user can browse the hierarchy from the most general classes (equipment, location, sensor,  
159 setpoint, substance) to the specific class whose definition best describes the entity.

160 A well-formed class hierarchy is extensible. Users can create new, more specific classes that subclass  
161 existing, but more general, superclasses. Even in the absence of a textual definition for this new class, the  
162 subclass relationship provides an immediate contextual scoping for how the class is meant to be used.

163 Haystack lacks an explicit class hierarchy, and the informal construction of Haystack complicates the  
164 automated generation of one. Recall that the type of each entity in a Haystack model is defined by a set of  
165 tags. We can formalize this as:

166 **Definition 3.1.** The set of tags for a class or entity  $x$  is given by  $T(x)$ . The definition of a class  $C$  is any  
167 entity that has the tags defined by  $T(C)$ . An entity  $e$  is a member of class  $C$  if  $T(e) \supseteq T(C)$

168 This definition embeds the key assumption of tag-based metadata: smaller tag sets convey more generic  
169 concepts. As an example, the pseudo-class identified by the Haystack tags `discharge air temp`  
170 `setpoint` is a subclass of the class identified by the tags `air temp setpoint`. However, the use of  
171 tag sets to specify the subclass relationship is insufficient for a well-formed class hierarchy because it is  
172 possible to construct two class definitions  $C_i$  and  $C_j$  such that  $T(C_i) \supseteq T(C_j)$  but the definition of  $C_i$  is  
173 not semantically more specific than  $C_j$ .

174 Consider the counter example of two concepts: `Air Flow Setpoint` (the desired cubic feet per  
175 minute of air flow) and `Max Air Flow Setpoint` (the maximum allowed air flow setpoint). In  
176 Haystack, `Air Flow Setpoint` would be identified by the `air`, `flow` and `sp` tags, and `Max Air`  
177 `Flow Setpoint` would be identified by the `air`, `flow`, `sp` and `max` tags. Although suggested by the  
178 set-based relationship, `Air Flow Setpoint` is *not* a superclass of `Max Air Flow Setpoint`: the  
179 former is a setpoint, but the latter is actually a parameter governing the selection of setpoints and therefore  
180 belongs to a distinct subhierarchy.

181 As a result, the rules for defining valid tagsets for subclass relationships must be defined in terms of  
182 which tags can be added to a given tag set to produce a valid subclass relationship. Defining a concept  
183 requires knowing which tags *cannot* be added; without a clear set of rules for validation, a user may use  
184 the `max` and `min` tags to indicate the upper/lower bounds of deadband-based control, which is inconsistent  
185 with the intended usage of these tags.

	Tag	Desc. equip	Desc. point	Desc. mechanism	For AHU	For VAV	For Coil	For Valve	For Chiller	For Boiler
heating	heat	×	×				×	×		
	heating		×							
	hotWaterHeat	×		×	×					
	gasHeat	×		×	×					
	elecHeat	×		×	×					
	steamHeat	×		×	×					
	perimeterHeat	×				×				
reheating	reheat		×			×				
	reheating		×			×				
	hotWaterReheat	×		×		×				
	elecReheat	×		×		×				
cooling	cool	×			×	×	×			
	cooling		×							
	coolOnly	×			×					
	dxCool	×		×	×					
	chilledWaterCool	×		×	×					
	waterCooled	×							×	
	airCooled	×							×	

**Table 1.** An enumeration of the intended use and context of tags relating to heating and cooling, as given by the Haystack documentation. Note the differences in diction across compound tags, and how some compound tags could be assembled from more atomic tags. Some tags are used both for equipment and for points when equipment is modeled as a single point (such as VFDs, Fans, Coils)

186 **3.2 Balancing Composability and Consistency**

187 One of the primary benefits of an tag-based metadata scheme is composability. A dictionary of tags  
 188 provides the vocabulary from which users can draw the terms they need to communicate some concept.  
 189 Composing sets of tags together allows for communication of increasingly complex concepts, and adding  
 190 new tags to the dictionary exponentially increases the number of describable concepts.

191 However, increased composability comes at the cost of lower *consistency*: a clear, unambiguous, one-  
 192 to-one mapping between a set of tags and a concept. Without rules defining composability, consistent  
 193 interpretation of a tag set is dependent upon idiom, convention and other “common knowledge” of the  
 194 community using the tags. As a result, the set of tags used by one individual to describe an entity may have  
 195 multiple meanings or no meaning at all to other individuals. In other words, the intended meaning of a tag  
 196 becomes more ambiguous the more contexts in which it is used. For example, using the tags *heat*, *oil*  
 197 and *equip* on an entity does not state if the equipment heats using oil or if oil is what is being heated.



198 To mitigate this effect, Haystack defines “compound” tags. These are concatenations of existing tags  
199 into new atomic tags with specific semantics distinct from that of its constituents. For example, the  
200 hotWaterHeat compound tag is defined specifically as indicating that an air handler unit has heating  
201 capability using hot water. This trades composability — which tags can be used together — with consistency  
202 — an unambiguous definition for a set of tags. Haystack calls these “semantic conflicts”:

203 “Another consideration is semantic conflicts. Many of the primary entity tags carry very specific  
204 semantics. For example the site tag by its presence means the data models a geographic site. So we  
205 cannot reuse the site tag to mean something associated with a site; which is why use the camel case  
206 tag siteMeter to mean the main meter associated with a site.[33]”

207 The most common types of semantic conflicts concern *process tags* and *substance tags*. We explore how  
208 ambiguities arise for these two family of tags; §4 demonstrates how Brick handles these issues.

209 **Process Tags.** For a metadata scheme to consistently describe a process, it must decompose a process  
210 into entities and capture how each entity relates to the process: does the entity monitor or control the  
211 process? Does it transport a substance or provide a means for two substances to interact?

212 It is difficult for a limited dictionary of tags to capture unambiguously the family of concepts involved.  
213 Consider the case of an air handling unit that heats air by passing it around a coil of hot water. With a  
214 limited dictionary of tags, most of the entities (equipment and points) involved will be tagged with hot,  
215 heat, air and/or water. However, a flat set of tags does not permit any differentiation between concepts  
216 that share the same tags.

217 Table 1 categorizes the intended usage of each Haystack tag containing the word “heat” (including  
218 “reheat”) or “cool.” Without this table or the Haystack documentation in hand, it is difficult to discern when  
219 to use a compound tag or several tags together: chilled+waterCooled, chilledWaterCool or  
220 chilled+water+cool? Furthermore, there is no formal, programmatically accessible form of the  
221 documentation that would allow this to be done in an automated fashion.

222 **Substance Tags.** What constitutes sufficient and consistent descriptions of substances (such as water  
223 and air) depends upon the breadth of intended use. Existing building metadata systems do not model sub-  
224 stances directly; instead, they describe equipment and points in terms of what substances they manipulate,  
225 measure or utilize. Thus, an effective metadata scheme for substances must capture *at least* the nature of  
226 the relationship between substances, equipment and points.

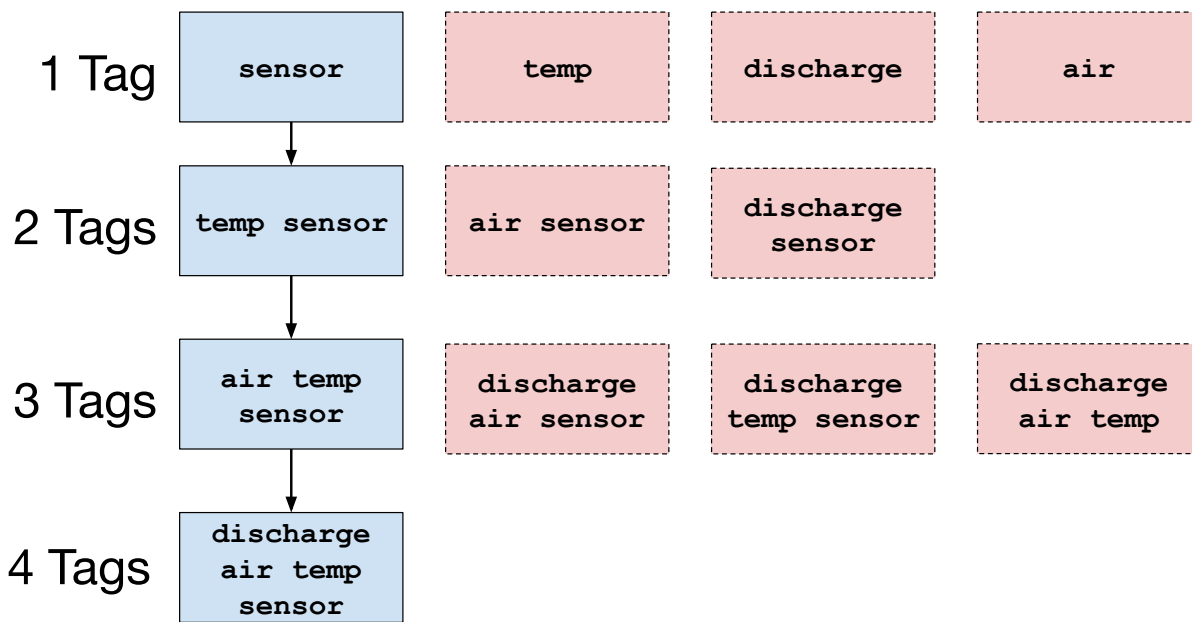
227 Flat tag structures lack the expressive power to make these distinctions unambiguous. To re-  
228 duce ambiguity, Haystack uses substance tags only on points and uses compound tags for equip-  
229 ment. For example hot water valve cmd and chilled water entering temp sensor  
230 use the water substance tag, but an air handler unit with a water-based chiller would use  
231 chilledWaterCool. This means that substance tags cannot be used to identify which points *and*  
232 *equipment* relate to a given substance. Furthermore, to perform such a query, a user would need  
233 to know the entire family of tags that relate to that substance. In the case of “water”, this list  
234 is water, waterCooled, waterMeterLoad, chilledWaterCool, chilledWaterPlant,  
235 hotWaterHeat, hotWaterPlant and hotWaterReheat, not to mention any user-defined tags.

### 236 3.3 Lack of Composition Rules

237 Haystack sacrifices composability of tags for more consistent interpretability, such as through the use  
238 of compound tags. Without a set of rules for how tags *can* be composed, there is no programmatic or

239 automated mechanism to enforce or inform consistent usage of the tag dictionary. Haystack contains a  
 240 small set of explicit rules, but largely relies upon idiom and human interpretation for consistency.

241 **Extending Tag Sets.** In order to encourage consistent usage, metadata schemes need rules for generating  
 242 new concepts and generalizing existing concepts. Rules for *generating* new concepts allow these concepts  
 243 to be qualified by their relation to existing classes. Rules for *generalizing* existing concepts allow users (and  
 244 programs) to reason about the behavior of a group of concepts. Formal mechanisms for generalization and  
 245 specialization aid the discoverability, interpretability and extensibility of a metadata scheme. Unfamiliar  
 246 concepts can be understood or referenced by their behavior or superclasses, and new concepts can be added  
 247 seamlessly.



**Figure 1.** The set of valid (blue + solid outline) and invalid (red + dashed outline) tagsets for a set of four tags. The class hierarchy is established from top to bottom; subclass relationships are indicated by arrows.

248 Concepts in Haystack can be extended through annotation with additional tags; e.g. `temp sensor`  
 249 refines the concept of `sensor`. However, tags cannot be freely combined (Figure 1). One mechanism for  
 250 defining valid tag sets parameterizes existing tag sets with a choice from a set of mutually exclusive tags.  
 251 Haystack explicitly defines several of these. Two examples from many:

- 252 1. The heating method for an AHU, given by one of `gasHeat`, `hotWaterHeat`, `steamHeat` or  
 253 `gasHeat`.
- 254 2. The family of water meters recognized by Haystack can be differentiated by the tags `domestic`,  
 255 `chilled`, `condenser`, `hot`, `makeup`, `blowdown` and `condensate`.

256 Haystack also has many *implicit* rules for defining valid extensions to tag sets. Application of these  
 257 rules largely depends upon domain knowledge – for example an entity will likely not have two distinct  
 258 substance tags such as `air` and `water` – as well as informal idioms conveyed through documentation. An  
 259 example of the latter is the convention that points (sensors, setpoints and commands) will have a “what”  
 260 tag (e.g. `air`), a “measurement” tag (e.g. `flow`) and a “where” tag (e.g. `discharge`). However, this is  
 261 not a hard and fast rule, and many of the tag sets in Haystack’s documentation break with this convention.



262 Consequently, there is no clear notion of how concepts can be meaningfully extended or generalized, which  
263 limits the extensibility of Haystack.

264 **Modeling Choices.** The lack of formal structures for constructing tag sets means that enforcement of  
265 consistency – choosing the same set of tags to represent the same concept – relies upon the conventions of  
266 industrial practice and the idioms of the Haystack community. As a result, there is substantial variation in  
267 how the same concept is modeled.

268 One prominent example in Haystack is the choice of whether to model pumps and fans as equipment or as  
269 points. Although pumps and fans are equipment, in many BMS they are represented by only a single point  
270 (usually the speed or power level). Haystack’s documentation encourages simplifying the representation of  
271 such equipment under such circumstances:

272 “Pumps may optionally be defined as either an `equip` or a `point`. If the pump is a VFD then it  
273 is recommended to make it an `equip` level entity. However if the pump is modeled [in the BMS] a  
274 simple on/off point as a component within a large piece of equipment such as a boiler then it is modeled  
275 as just a `point`. [34]”

276 Complex predicates such as these complicate the querying of a Haystack model. In particular, exploratory  
277 queries have to take the family of modeling choices into account: to list all of the pumps in a Haystack  
278 model, it is not sufficient to only look for entities with the `pump` and `equip` tag.

### 279 3.4 Impact on Consistency

280 These issues with tag-based metadata inhibit extensibility and consistency at scale. Most Haystack models  
281 are designed to be used by small teams familiar with the site or sites at hand, so it is enough for these  
282 models to be *self-consistent*. As long as there is agreement on how to tag a given concept, the informality of  
283 the model is not as detrimental; most tag sets in Haystack make intuitive sense to domain experts. However,  
284 the lack of formalization — specifically, a lack of a formal class hierarchy and rules for composability and  
285 extensibility — presents issues for adoption as an industrial standard and basis for automated analysis and  
286 reasoning.

287 In the next sections, we show that the tradeoff between composability and consistency is tied to the  
288 choice to use tags for annotation as well as definition. With an explicit and formal class hierarchy it is  
289 possible to design a system that exhibits the composability of simple tags, while retaining the consistency  
290 and extensibility of an ontology.

## 4 DESIGN OF BRICK+

291 Although Brick [4] establishes a formal class hierarchy and a set of descriptive relationships, it lacks  
292 the structure for inference of classes from tags and exhibits a number of design issues that impede this  
293 development. This motivates the design of *Brick+*, a drop-in replacement ontology for Brick that extends  
294 the hierarchy of described concepts to include fine-grained semantic properties and defines an explicit  
295 mapping from Brick concepts to sets of tags. Together, these enable the *programmatically interpretation* of tag  
296 sets, therefore eliminating the consistency and interpretability issues inherent to a tags-only design (§3).  
297 In conjunction with the structured implementation (§5), the formal design of Brick+ also enables a suite  
298 of supporting tooling for the validation (§6), migration (§7) and construction (§5.3) of Brick+ metadata  
299 models.

### 300 4.1 Limitations of Brick

301 The design and implementation of Brick has several issues which inhibit formalizing the relationship  
302 between classes and tags.

303 **No formal equivalence between tag sets and classes.** Brick models a class hierarchy using a special  
 304 construction called a TagSet. A TagSet has a definition, a set of related tags, and a name composed of each  
 305 of the tags concatenated together. The Brick ontology defines which tags are used with which TagSets, but  
 306 fails to capture bidirectional equivalency between the two definitions. Brick can retrieve the tags associated  
 307 with a TagSet, but given a set of tags, Brick cannot infer the set of possible TagSets.

308 **No modeling of function or behavior.** The Brick class hierarchy relates different TagSets only by a  
 309 “subclass” relationship; there is no semantic information to distinguish classes in terms of their behavior.  
 310 The simple association of tags to TagSets also does not offer any semantic information. Enhancing the class  
 311 definitions with more semantic information would increase the usability of Brick and the discoverability of  
 312 concepts.

313 **Inconsistent modeling and implementation.** The implementation of the Brick ontology consists of a  
 314 set of Turtle<sup>2</sup> files containing the ontology statements. These files are generated by a Python script that  
 315 transforms an CSV-based specification into the Turtle syntax for RDF. This process is brittle, error-prone  
 316 and difficult to test and extend.

## 317 4.2 Overview of Brick+

318 Brick+ has three components: a class lattice defining the family of equipment, points, locations, substances  
 319 and quantities in buildings; a set of expressive relationships defining how entities behave and how they are  
 320 connected, contained, used and located; and a family of tags defining the atomic attributes and aspects of  
 321 entities

322 The implementation of Brick+ relies upon the use of a *semantic reasoner*, piece of software that  
 323 materializes the set of facts deduced through the application of the logical rules contained within an  
 324 ontology. An important implementation factor is the language used to define the ontology: more expressive  
 325 languages can significantly increase the runtime complexity of the reasoning process (decreasing the utility  
 326 of the system in an applied context), whereas less expressive languages may not be able define the necessary  
 327 rules. The formal specification of Brick+ uses the OWL DL language to define rules for the operation and  
 328 usage of Brick+ and to achieve the desired runtime properties.

## 329 4.3 Brick+ Class Lattice

330 Brick+ organizes all concepts into a class structure rooted in a small number of high-level concepts.  
 331 Brick defines this structure as a tree-based hierarchy; Brick+ refines this structure into a *lattice*. Both the  
 332 lattice and the hierarchy are defined in terms of a “subclass” relationship (§2), but differ in how they define  
 333 relationships between concepts. A class hierarchy captures how concepts can be specialized, but does not  
 334 encode how these concepts behave and relate to one another. In contrast, a lattice captures how concepts  
 335 can be composed from sets of properties. This offers greater flexibility in the definition of concepts in  
 336 Brick+ and facilitates the tag decomposition and mapping to Haystack detailed in §5.

337 Brick+ has six primary concepts. **Point** is the root class for all points of telemetry and actuation. There  
 338 are six immediate subclasses of `Point` categorized by the high-level semantics of how each point behaves:

- 339 • `Sensor` points are outputs of transducers recording the state of the physical world, e.g. `Air`  
 340 `Temperature Sensor`
- 341 • `Setpoints` points are control points used to guide the operation of a feedback-driven control system,  
 342 e.g. `Air Flow Setpoint`

<sup>2</sup> <https://www.w3.org/TR/turtle/>

- 343 • Command points are control points that directly affect the state of equipment, e.g. Fan Speed  
344 Command
- 345 • Status points report the current logical status of equipment, e.g. Damper Position Status
- 346 • Alarm points are high-priority indicators conveying non-nominal behavior, e.g. Water Loss  
347 Alarm
- 348 • Parameter points are configuration settings used to guide the operation of equipment and control  
349 systems, e.g. Max Air Flow Setpoint.

350 Brick+ refines the design of the Brick ontology to differentiate between parameters and setpoints. This  
351 avoids conflating the concepts of the minimum and maximum setpoints used in deadband control (such  
352 as to configure a thermostat to maintain a temperature within that band) and the minimum and maximum  
353 allowed values for a setpoint (for example to place a lower bound on permitted air flow setpoints).

354 **Equipment** is the root class for the lattice of mechanical equipment used in a building. The Brick+  
355 equipment lattice covers equipment for HVAC, lighting, electrical and water subsystems. Brick+ extends  
356 the modeling of equipment in Brick to include how classes of equipment relate to substances and processes  
357 in the building.

358 **Location** is the root class for the lattice of spatial elements of a building. The lattice includes physical  
359 elements such as floors, rooms, hallways and buildings as well as logically-defined physical extents such as  
360 HVAC, lighting and fire zones.

361 **Substance** is the root class for the lattice of physical concepts that are measured, monitored, controlled  
362 and manipulated by building subsystems. Examples of physical substances are air, water and natural gas.  
363 These can be further subclassed by their usage within the building, for example “mixed air” is a subclass of  
364 “air” that refers to the combination of outside and return air in an air handler unit.

365 **Quantity** is the root class for the lattice of quantifiable properties of substances and equipment.  
366 Examples of physical properties include temperature, conductivity, voltage, luminance and pressure.  
367 Subclassing quantities enables differentiation between types of quantities, such as between Dry Bulb  
368 Temperature and Wet Bulb Temperature.

369 **Tag** is a root class for the flat namespace of atomic tags supported by Brick+. The majority of these tags  
370 are drawn from the Haystack tag dictionary, and are *instances* of the Tag class.

#### 371 4.4 Brick+ Relationships

372 Relationships express how entities and concepts can be composed with one another; this is key to  
373 the consistent and extensible usage of Brick+. For entities – the “things” in a building – composition  
374 encapsulates functional relationships such as monitoring, controlling, manipulation, sequencing within a  
375 process, and physical and logical encapsulation. Concepts are identified by classes and are organized into a  
376 lattice by relationships.

377 As in Brick, relationships in Brick+ exist between a *subject* (the entity possessing the relationship’s  
378 indicated property) and an *object* (the entity that is the target of the property). Brick+ defines a set of  
379 constraints for each relationship to ensure correct and consistent usage between subject and object entities,  
380 without constraining the application of the relationship to yet unknown scenarios.

381 All Brick relationships have at least one domain or range constraint determining the allowed classes for  
382 the subject or object. Domain constraints limit the class of entities that can be the subject of a relationship;  
383 range constraints limit the class of entities that can be the object of a relationship. Brick defines domains and

Relationship	Definition	Domain	Range	Inverse	Transitive?
hasLocation	Subject is physically located in the object entity	*	Location	isLocationOf	yes
feeds	Subject conveys some media to the object entity in some sequential process	Equipment Equipment	Equipment Location	isFedBy	no
hasPoint	Subject has a monitoring, sensing or control point given by the object entity	Equipment Location	Point Point	isPointOf	no
hasPart	Subject is composed – logically or physically – in part by the object entity	Equipment Location	Equipment Location	isPartOf	yes
measures	Subject measures a quantity or substance given by the object entity	Sensor Sensor	Substance Quantity		no no
regulates	Subject informs or performs the regulation of the substance given by the object entity	Setpoint Equipment	Substance Substance		no
hasOutputSubstance	Subject produces or exports the object entity as a product of its internal process	Equipment	Substance		no
hasInputSubstance	Subject receives the object entity to conduct its internal process	Equipment	Substance		no

**Table 2.** List of high-level relationships supported by Brick+.

384 ranges of relationships in terms of classes from the lattice. Brick+ supports these definitions (enumerated  
 385 in Table 2) and extends them such that domains and ranges can be defined in terms of the properties of  
 386 the subject and object, rather than which sublattice they belong to. This allows the definition of more  
 387 fine-grained *sub-relationships* with additional semantics.

388 For example, as in Table 2, the `feeds` relationship indicates the passage of some substance between  
 389 two pieces of equipment or between an equipment and a location. If the subject of the `feeds` relationship  
 390 has the property that it outputs air, then the `feeds` relationship can be specialized to the `feedsAir`  
 391 sub-relationship.

#### 392 4.5 Brick+ Tags

393 Brick+ addresses the consistency and interpretability issues of tag-based metadata by explicitly binding  
 394 Brick classes to sets of tags. In Brick, classes are human-interpretable because they have clear textual  
 395 definitions; in Brick+, classes are additionally programmatically-interpretable because they are identified  
 396 by their position in the class lattice and by the set of properties that define their behavior. Clear definitions  
 397 promote consistent usage.

398 Binding classes to tag sets effectively bounds the family of possible tag sets to those that have clear  
 399 definitions. *This removes the burden of definition, validation and interpretation from the tag structure by*  
 400 *outsourcing it to the class lattice, which permits the inference of Brick+ classes from unstructured Haystack*  
 401 *tags.*

402 Although Brick also defines tags, Brick+ advances the implementation in several ways. Firstly, Brick+  
 403 removes the need for tags to be lexically contained within the name of the class (the “TagSet” construct in  
 404 Brick). This decoupling allows the definition of classes beyond what can be assembled through concatena-  
 405 tion of tags, or classes that do not have a straightforward tag decomposition; for example, a Rooftop Unit  
 406 equipment in Haystack has the `rtu` tag.

407 Secondly, Brick+ encodes tags so they can be inferred from a Brick+ class and vice versa, even if a given  
 408 entity’s definition is given only by one or the other. Figure 2 illustrates three different methods for instanti-  
 409 ating an `Air Temperature Sensor` demonstrating the flexibility of the Brick+ implementation. The  
 410 classification of an entity can be performed *explicitly* using the `a` or `rdf:type` predicate in conjunction  
 411 with a Brick class, *implicitly* through annotating an entity with the set of tags equivalent to a Brick class,  
 412 *descriptively* by annotating an entity with its behavioral properties, or through a combination of these.

---

```

1 # instantiate class explicitly
2 :sensor1 a brick:Air_Temperature_Sensor .
3
4 # instantiate a class implicitly through application of tags
5 :sensor1 brick:hasTag tag:Air .
6 :sensor1 brick:hasTag tag:Temp .
7 :sensor1 brick:hasTag tag:Sensor .
8
9 # combination of explicit class and tags
10 :sensor1 a brick:Temperature_Sensor .
11 :sensor1 brick:hasTag tag:Air .
12
13 # instantiation from behavior
14 :sensor1 a brick:Sensor .
15 :sensor1 brick:measures brick:Air .
16 :sensor1 brick:measures brick:Temperature .
17
18 # alternative instantiation from behavior
19 :sensor1 a brick:Temperature_Sensor .
20 :sensor1 brick:measures brick:Air .

```

---

**Figure 2.** Five equivalent methods of declaring sensor1 to be an instance of the Brick Air Temperature Sensor class.

---

```

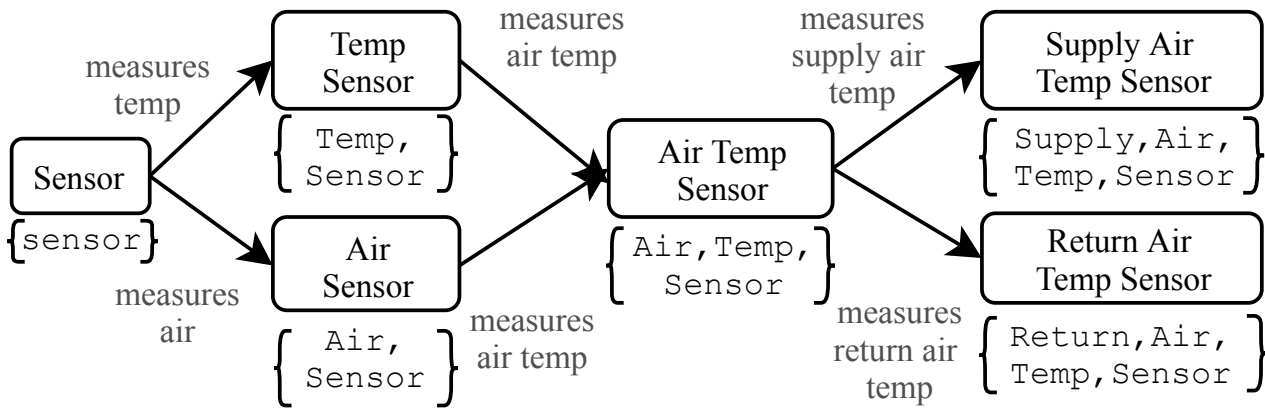
1 brick:Supply_Air_Temperature_Sensor a owl:Class ;
2   rdfs:subClassOf brick:Air_Temperature_Sensor ;
3   rdfs:subClassOf [ owl:intersectionOf (
4     [ a owl:Restriction ; owl:hasValue tag:Sensor ;
5       owl:onProperty brick:hasTag ]
6     [ a owl:Restriction ; owl:hasValue tag:Temperature ;
7       owl:onProperty brick:hasTag ]
8     [ a owl:Restriction ; owl:hasValue tag:Air ;
9       owl:onProperty brick:hasTag ]
10    [ a owl:Restriction ; owl:hasValue tag:Supply ;
11      owl:onProperty brick:hasTag ] ) ],
12   [ owl:intersectionOf (
13     [ a owl:Restriction ; owl:hasValue brick:Temperature ;
14       owl:onProperty brick:measures ]
15     [ a owl:Restriction ; owl:hasValue brick:Supply_Air ;
16       owl:onProperty brick:measures ] ) ] .

```

---

**Figure 3.** OWL DL-compatible definition of the Brick Supply Air Temperature Sensor class showing the explicit class structure, tag equivalence and the use of substance and quantity classes to model behavior

413 Figure 4 illustrates how tags, classes and properties define the lattice for some subclasses of the  
414 Sensor class. Figure 3 shows the implementation of the Supply Air Temperature Sensor  
415 class: line 2 defines how Supply Air Temperature Sensor fits into the Brick class lattice.  
416 Lines 4-17 defines the Supply Air Temperature Sensor class as equivalent to entities that  
417 have the sensor, temperature, air and supply tags. Lines 18-25 define the Supply Air  
418 Temperature Sensor class as equivalent of entities that measure the Temperature property of the  
419 Supply Air substance.



**Figure 4.** Portion of Brick+ class lattice illustrating the equivalence between tags and classes. Edges indicate which properties are added to each concept (class) to produce a new class. Reverse edges (not pictured) are the subclass relationships.

#### 420 4.6 Brick+ Substances and Quantities

421 Brick+ defines a lattice of substances and quantities that can be used to describe the functionality of  
 422 equipment and points. This permits inference of more fine-grained semantic information from existing  
 423 Brick models and allows equipment and points to be classified by their behavior rather than by explicit  
 424 classification.

425 The Brick+ substance class lattice is based upon the hierarchy developed by Project Haystack. It classifies  
 426 substances by phase of matter (*Gas*, *Liquid*, *Solid*) and supports substances qualified by their usage  
 427 within a process: *Air* is a subclass of *Gas*, and *Outside Air* and *Mixed Air* are subclasses of *Air*.  
 428 This construction can be extended to include new substances and subclasses of those substances as used in  
 429 different processes.

430 A key principle of the Brick+ implementation is every property associated with a class must be inferrable  
 431 from instances of that class. Properties associated with classes include the set of tags that are equivalent to  
 432 the class (indicated by the `hasTag` relationship) and the behavioral annotations of the class (indicated by  
 433 relationships like `measures`).

## 5 BRICK+ IMPLEMENTATION

434 Recall that the Brick+ class lattice models concepts by their behavior and related tags as well as by explicit  
 435 subclass relationships. The lattice is defined by a family of relationships which are supported by a set of  
 436 constraints that ensure correct and consistent usage between subject and object entities without constraining  
 437 the application of the relationship to yet unknown scenarios. This enables Brick+ to define a formal  
 438 mapping from Haystack’s informal tags to formal Brick classes.

439 To facilitate the development, testing and debugging of Brick+, we created a Python framework that  
 440 interprets a structured and extensible abstract ontology specification into a Turtle-based implementation.  
 441 The framework is open source and available online<sup>3</sup>. The implementation of the inference engine described  
 442 in §5 and the validation tool described in §6 are published as part of the open source `brickschema`

<sup>3</sup> <https://github.com/BrickSchema/Brick>



443 Python package<sup>4</sup>. The migration tool described in §7 is distributed as part of the Brick+ source code, and is  
444 also open source.

445 This section presents an overview of the implementation of the Brick+ ontology with a focus on the  
446 implementation of substances and the inference procedure for converting Haystack tags to Brick classes.

## 447 **5.1 Substance Implementation**

448 The inferred properties concerning substances are more complex to account for the differences in usage.  
449 Because substances are classes, it is possible to associate instances of substances with Brick entities.  
450 This association helps applications model how entities behave in relation to the same substance instance.  
451 For example, a `Mixed Air Temperature Sensor` and a `Mixed Air Damper` could be related  
452 through their respective measurement and regulation of the same instance of `Mixed Air`. However, if a  
453 shared instance is not given in the definition of the Brick model, an OWL DL reasoner cannot infer the  
454 instantiation of an appropriate substance. Brick+ solves this with *punning* [41].

455 Punning is a mechanism by which a class name can represent a canonical instance of that class. This  
456 allows an OWL DL reasoner can relate a punned substance to a property of an equipment or point.  
457 Importantly, this does not prohibit the instantiation of substance instances if and when a Brick model  
458 supplies those. Line 15 of Figure 2 contains an example of an inferred substance for instances of the  
459 `brick:Air_Temperature_Sensor` class.

## 460 **5.2 Formalizing the Tag-Class Equivalence**

461 There are several possible approaches for formalizing the mapping between tag sets and classes, each  
462 with non-obvious tradeoffs. Fundamentally, the formalization has to grapple with the same issue captured  
463 in Definition 3.1. The use of the subset relationship between the set of tags describing an entity and the  
464 set of tags defining a class can produce ambiguous or erroneous results. In fact, the process described in  
465 Definition 3.1 is only sufficient if there is no pair of disjoint classes whose tags are a subset of each other.  
466 This implies that under such semantics the only permitted subset relationships between tag sets are those  
467 that also obey the subclass relationship between the corresponding classes.

468 To see why this is the case, consider the following (informal) definitions:

- 469 • All `Setpoint` classes are disjoint from all `Parameter` classes; that is, the set of instances of the  
470 `Setpoint` type is disjoint with the set of instances of the `Parameter` type.
- 471 • The `Air Flow Setpoint` class (a subclass of `Setpoint`) is defined by the tags `air`, `flow` and  
472 `setpoint`.
- 473 • The `Supply Air Flow Setpoint` class (a subclass of `Setpoint`) is defined by the tags  
474 `supply`, `air`, `flow` and `setpoint`.
- 475 • The `Max Air Flow Setpoint Limit` class (a subclass of `Parameter`) is defined by the tags  
476 `max`, `air`, `flow`, `setpoint` and `limit`.

477 The tags for `Air Flow Setpoint` are a subset of the tags of *both* `Supply Air Flow`  
478 `Setpoint` and `Max Air Flow Setpoint Limit`. The former is permissible because `Supply`  
479 `Air Flow Setpoint` is a subclass of `Air Flow Setpoint`. However, the subset relationship in-  
480 correctly implies that `Max Air Flow Setpoint Limit` is a subclass of `Air Flow Setpoint`,  
481 which is a violation of the disjoint relationship between their respective parent classes (`Parameter` and

---

<sup>4</sup> <https://brickschema.readthedocs.io/>

482 Setpoint). The primary challenge in formalizing the mapping between tag sets and classes is how to  
483 overcome this limitation with respect to the following constraints:

- 484 1. **Build on the RDF data model and OWL ontology language.** The mapping between tag sets and  
485 classes should be expressed in the RDF data model for compatibility with the existing entities,  
486 annotations and properties already defined in the Brick+ ontology. Additionally, inferring a class  
487 from a set of tags (and vice versa) should be possible through the mechanisms in the OWL ontology  
488 language, i.e. via a reasoning engine. This keeps the formalism consistent with the design of the  
489 ontology and the mapping to the space of tags that Brick+ incorporates and also minimizes the amount  
490 of external machinery required to make use of the developed mappings.
- 491 2. **The set of tags mapped to a class should be sensible.** There is a strong usability argument for  
492 maintaining a “sensible” set of tags that maps to each class; a “sensible” set of tags is defined as a  
493 set of tags that “makes sense.” This may include the descriptive terms in the class name or words  
494 commonly associated with the class concept. Importantly the interpretability of a “sensible” set of tags  
495 does not depend on the large class structure. As long as the set of tags for every class is unique, the  
496 mapping developed here is sound.

497 However, developing a “sensible” mapping that can be expressed in RDF and that is supported by  
498 OWL-based inference is a tricky task. Because the Brick+ formal model is built on the set logic of OWL, a  
499 naive approach to designing a mapping between tags and classes quickly encounters the tension between  
500 tag set overlap and the disjointness of the class hierarchy described above. It is tempting to design the  
501 tag sets associated with classes in a manner that sidesteps this issue; for example, one could define  
502 the tags associated with `Max Air Flow Setpoint Limit` to be `max`, `air`, `flow` and `limit`  
503 (excluding `setpoint`) to avoid the logical violation given by the subsumption of the tags for `Air Flow`  
504 `Setpoint`. This is possible due to the perscriptive nature of the mapping, but it may feel unnatural to  
505 users of Brick+ to avoid the `setpoint` tag when describing a concept related to a setpoint, even if the  
506 concept itself is not a setpoint.

507 How, then, can both properties be achieved in the same design? The current understanding of the authors  
508 is that implementing a sensible mapping is fundamentally incompatible with a solution expressed *entirely*  
509 in the set logic of OWL. A full proof of this assertion is beyond the scope of this paper, but we present  
510 the essential intuition in the course of explaining the design and implementation of the Brick+ tag-class  
511 mapping below.

512 The tag-class mapping in Brick+ decouples the informal specification of which tags are associated with  
513 which classes from how this mapping is formalized in the ontology. This allows the two representations to  
514 evolve independently. The Brick+ generation framework incorporates a list of tags for each defined class.  
515 A set of unit tests verifies that there is a unique set of tags for each class in the Brick+ class hierarchy and  
516 raises a warning to the developer if this is not the case. Invoking the framework generates Brick+ ontology,  
517 including the tag-class mapping. We discuss each direction of the mapping separately below.

518 **Class to Tag Set Mapping:** The Brick+ ontology defines a `hasTag` property which associates an entity  
519 (an instance of a class) with a tag. The mapping must enable the population of the appropriate `hasTag`  
520 properties for each instance of a Brick+ class. The key construct in the modeling approach involves the  
521 use of OWL Restriction classes (marked by `owl:Restriction`); these are an OWL construct which  
522 defines the members of a class as those that possess certain properties. Brick+ defines a Restriction class  
523 for each tag in the ontology that, by definition, counts its members as those entities that possess the given  
524 tag. The Restriction classes enable a semantic reasoner to generate a `hasTag` property with a given tag

---

```

1 # OWL Restriction definition
2 brick:has_Temperature a owl:Restriction ;
3   owl:hasValue tag:Temperature ;
4   owl:onProperty brick:hasTag .
5
6 # model statement
7 example:thing a brick:has_Temperature .
8
9 # generated by the reasoner
10 example:thing brick:hasTag tag:Temperature .

```

---

**Figure 5.** An example of an OWL Restriction class encoding the association with the Temperature tag in Brick.

525 for each member of the class. Figure 5 demonstrates this mechanism: lines 2-4 contain the definition of a  
 526 Restriction class for the Temperature tag. Line 7 declares an entity as an instance of the Restriction  
 527 class; this allows a semantic reasoner to produce the association with the tag as captured on line 10.

528 Brick+ represents the set of tags associated with a class as the intersection of each of the constituent  
 529 Restriction classes. The intersection is realized as a class formed by the `owl:intersectionOf` property.  
 530 Lines 3-11 of Figure 3 contain an example of this construction. By definition, the intersection class has as  
 531 its members the set of entities that fulfill the requirements of each of the constituent tag Restriction classes  
 532 — that is, the set of entities that have the required tags. This construction means that any entities declared to  
 533 be instances of this intersection class can, through the application of a semantic reasoner, have the required  
 534 tags automatically associated.

535 The last piece of the implementation involves how to associate a Brick+ class with the anonymous  
 536 intersection class that represents the set of tags (referred to as a *tag set class*). There are two available  
 537 approaches: marking the Brick+ class as a *subclass* of the tag set class, or marking the Brick+ class as *equiv-*  
 538 *alent* to the tag set class. Marking the tag set class as *equivalent* (using the `owl:equivalentClass`  
 539 property) is not sufficient because it captures entities with a superset of the required tags; this is exactly  
 540 the logical violation problem described above. Therefore, we use the other approach. Brick+ encodes a  
 541 Brick+ class as a subclass (using the RDFS `subClassOf` property) of the tag set class. This means that  
 542 an entity declared to be a member of a Brick+ class will inherit membership of each of the Restriction  
 543 classes in the tag set class construct; the application of a semantic reasoner can then populate the required  
 544 tag associations.

545 **Tag Set to Class Mapping:** Encoding the mapping from a set of tags to a Brick+ class requires reasoning  
 546 about what tags an entity does *not* have in addition to which tags it does have. This is incompatible with  
 547 the *open world assumption* [36], which is employed by the underlying set logic of OWL. Informally, the  
 548 open world assumption tells us that a statement may be true regardless of whether or not that statement  
 549 is included in a given knowledge base. In the context of tag-class mapping, the open world assumption  
 550 means that the *absence* of a tag on an instance cannot be interpreted as the instance not having that tag;  
 551 that statement may exist elsewhere. A deeper discussion about the utility of the open world assumption  
 552 in the context of buildings — where a knowledge base may indeed be the authoritative data source — is  
 553 beyond the scope of this paper and is the subject of ongoing work. To circumvent this issue, the tag set to

```

1 id: 'd83664ec RTU-1 OutsideDamper'
2 air: ✓
3 cmd: ✓
4 cur: ✓
5 damper: ✓
6 outside: ✓
7 point: ✓
8 regionRef: '67faf4db'
9 siteRef: 'a89a6c66'
10 equipRef: 'd265b064'

```

**Figure 6.** Original Haystack entity from the Carytown reference model

```

1 :d83664ec brick:hasTag tag:Command . # cmd
2 :d83664ec brick:hasTag tag:Damper .
3 :d83664ec brick:hasTag tag:Outside .
4 :d83664ec brick:hasTag tag:Point .

```

**Figure 7.** Intermediate RDF representation of the Haystack entity; Haystack software-specific tags (e.g. `cur`, `tz`) are dropped.

```

1 :d83664ec_point brick:hasTag tag:Damper .
2 :d83664ec_point brick:hasTag tag:Command .
3 :d83664ec_point a brick:Damper_Position_Command . # inferred
4 :d83664ec_equip brick:hasTag tag:Air .
5 :d83664ec_equip brick:hasTag tag:Outside .
6 :d83664ec_equip brick:hasTag tag:Damper .
7 :d83664ec_equip a brick:Outside_Damper . # inferred
8 :d83664ec_point brick:isPointOf :d83664ec_equip . # inferred
9 :d83664ec_point brick:isPartOf :d265b064 # inferred

```

**Figure 8.** Brick inference engine splits the entity into two components: the explicit point and the implicit outside damper equipment.

554 class mapping in Brick+ is accomplished using a simplified version of the inference procedure described in  
 555 §5.3. The implementation is captured in the external `brickschema` Python package<sup>5</sup>.

### 556 5.3 Brick-Haystack Inference Procedure

557 In order to apply the Brick+ inference to Haystack entities, some preprocessing is required. Firstly,  
 558 the engine filters out Haystack tags that do not contribute to the definition of the entity, including data  
 559 historian configuration (`hisEnd`, `hisSize`, `hisStart`), current readings (`curVal`) and display names  
 560 (`disMacro`, `navName`). Figure 6 shows an example of a “cleaned” Haystack entity containing only the  
 561 marker and `Ref` tags from the Carytown reference model.

562 Next, the engine transforms the Haystack entity into an RDF representation that can be understood by the  
 563 inference engine. The engine translates each of the marker tags into their canonical Brick form: for example,  
 564 Haystack’s `sp` becomes `Setpoint`, `cmd` becomes `Command` and `temp` becomes `Temperature`. The  
 565 engine creates a Brick entity identified by the label given by the Haystack `id` field, and associates each of  
 566 the Brick tags with that entity using the `brick:hasTag` relationship. Figure 7 contains the output of this  
 567 stage executed against the entity in Figure 6.

<sup>5</sup> <https://brickschema.readthedocs.io/>

568 At this stage, the engine naïvely assumes a one-to-one mapping between a Haystack entity and a Brick  
569 entity. This is usually valid for equipment entities which possess the `equip` tag, but Haystack point entities  
570 (with the `point` tag) may implicitly refer to equipment that is not modeled elsewhere. Figure 6 is an  
571 example of a Haystack point entity that refers to an outside air damper that is not explicitly modeled in the  
572 Haystack model. The last stage of the inference engine performs the “splitting” of a Haystack entity into an  
573 equipment and point.

574 First, the inference engine attempts to classify an entity as an equipment. The engine temporarily replaces  
575 all point-related tags from an entity – `Point`, `Command`, `Setpoint`, `Sensor` – with the `Equipment`  
576 tag, and finds Brick classes with the smallest tag sets that maximize the intersection with the entity’s  
577 tags. This corresponds to the *most generic Brick class*. In our running example, the inference engine  
578 would transform the entity in Figure 7 to the tags `Damper`, `Outside` and `Equipment`. There are 12  
579 Brick classes with the `Damper` tag, but only one class with both the `Damper` and `Outside` tags; thus,  
580 the minimal Brick class with the maximal tag intersection is `Outside Air Damper`. If the inference  
581 engine cannot find a class with a non-negligible overlap (such as the `Equipment` tag), then the entity is  
582 not equipment.

583 Secondly, the inference engine attempts to classify the entity as a point. In this case, the engine does  
584 not remove any tags from the entity, and finds the Brick classes with the smallest tag sets that maximize  
585 the intersection with the entity’s tags. In our running example, the minimal class with the maximal tag  
586 intersection is `Damper Position Command`.

587 Figure 8 contains the two inferred entities output by this methodology. In the case where a Haystack  
588 entity is split into an equipment and a point, the Brick inference engine associates the two entities with  
589 the `brick:isPointOf` relationship (line 10 of Figure 8). Additionally, the inference engine translates  
590 Haystack’s `Ref` tags into Brick relationships using the simple lookup-table based methodology established  
591 in [5]. The inference engine applies these stages to each entity in a Haystack model; the union of the  
592 produced entities and relationships constitutes the inferred Brick model.

## 6 BRICK+ VALIDATION

593 While the formal definition of the Brick+ ontology enables the exact specification of concepts and the  
594 relationships between them, it does not directly provide a means for validating correct or idiomatic usage  
595 within a model. Recall that a *Brick+ model* (sometimes referred to as an *instance* of Brick+) is an RDF  
596 graph that uses the Brick+ ontology to represent and describe the entities and relationships within a  
597 building.

598 We begin by defining *correct* and *idiomatic* usage of the Brick+ ontology and provide examples of  
599 where the need for such validation arises in practice. We then show how we apply the SHApes Constraint  
600 Language to this task by defining “shapes” that encode correct and idiomatic practices. Finally, we describe  
601 the implementation of the Brick+ validation library and tool.

602 The implementation and execution of Brick+ validation is made possible because of the formalization of  
603 the underlying model. In Brick and Haystack, the lack of formal rules means that there is no specification  
604 of what “correct” or “idiomatic” usage looks like — this must instead be determined through forum posts  
605 and human-readable documentation. Brick+ advances the state-of-the-art of building metadata by enabling  
606 such a validation process to be performed and in an automatic manner.

---

```

1 brick:isPointOf a owl:ObjectProperty ;
2   rdfs:domain brick:Point ;
3   owl:inverseOf brick:hasPoint .
4
5 brick:Point a owl:Class ;
6   owl:disjointWith brick:Equipment .
7
8 building:ahu1 a brick:Equipment .
9 building:vav1 a brick:Equipment .
10 building:ahu1 brick:isPointOf building:vav1 .

```

---

**Figure 9.** An example of a logical violation in an instance of a Brick+ model.

## 607 6.1 The Role of Validation

608 We define *correct* usage of an ontology to mean that the terms and properties defined in the ontology are  
609 used appropriately within an instance and do not result in any logical violations of the formal model. For  
610 example, consider the set of RDF statements in Figure 9. Lines 1-6 are a partial implementation of the  
611 Brick+ ontology. Lines 1-3 specify that any subject of the `brick:isPointOf` property is implied to  
612 be an instance of the `brick:Point` class. lines 5-6 state that the set of instances of `brick:Point` is  
613 disjoint with the set of instances of `brick:Equipment`; that is, no entity can be both a Point and an  
614 Equipment.

615 We now turn to the definition of a violating Brick+ model. Lines 8 and 9 define two pieces of equipment  
616 (we use the high-level `brick:Equipment` for illustrative purposes; in reality these instances would be  
617 members of more specific, descriptive classes within the Brick+ Equipment class lattice). Line 10 introduces  
618 the logical violation: the use of the `brick:isPointOf` property implies that `building:ahu1` is  
619 a member of `brick:Point` which conflicts with the statement on line 8 that `building:ahu1` is a  
620 member of the disjoint class `brick:Equipment`. Note that without external information — such as the  
621 domain knowledge that an entity named “ahu1” is likely an air handling unit and thus an equipment — it is  
622 impossible to tell which statement is erroneous.

623 *Idiomatic* usage of an ontology such as Brick+ differs from *correct* usage in that idiomatic violations are  
624 still logically valid. Instead, such violations are failures to meet structural and organizational expectations.  
625 The specification of modeling idioms is essential for normalizing the use of an ontology to a higher degree  
626 than can reasonably be provided by the ontology definition itself. Because the Brick+ ontology is meant  
627 to generalize to many different kinds of buildings, subsystems, equipment and organizations thereof, the  
628 ontology definition makes very few statements about what information is *required* to be present in a given  
629 building instance for it to be considered valid. Idioms fill this gap by encoding “best practices” of what  
630 *should* be contained in a given model.

631 Modeling idioms are diverse in form because they can fulfill many roles. For example, modeling idioms  
632 may include the expectation that

- 633 • all VAVs in an instance should refer to an upstream AHU and a downstream HVAC zone
- 634 • all VAVs of a particular make and model should have five associated monitoring and control points
- 635 • all temperature sensors should be reporting in Celsius

636 Because modeling idioms are not tied directly to the formal definition Brick+, their enforcement is not a  
637 requirement for the usage of Brick+. We expect modeling idioms to be defined, distributed and applied on



---

```

1 bsh:isPointOfDomainShape a sh:NodeShape ;
2   sh:targetSubjectsOf brick:isPointOf ;
3   sh:message "Subject of isPointOf should be an instance of Point" ;
4   sh:class brick:Point .
5
6 bsh:hasPointRangeShape a sh:NodeShape ;
7   sh:targetSubjectsOf brick:hasPoint ;
8   sh:property [
9     sh:class brick:Point ;
10    sh:message "Object of hasPoint should be an instance of Point" ;
11    sh:path brick:hasPoint ] .

```

---

**Figure 10.** SHACL shapes defining correct usage of the `brick:isPointOf` and `brick:hasPoint` relationships

638 a per-project or per-building basis; in the future, equipment manufacturers may distribute Brick+-encoded  
639 modeling idioms for how their equipment should be represented in a Brick+ model.

## 640 **6.2 Brick+ Validation with SHACL**

641 The abstract Brick specification described in §5 lends a means to generate constraints enforcing correct  
642 and/or idiomatic usage. These constraints are defined using the SHACLs Constraint Language (SHACL [21]),  
643 a W3C standard for validating RDF graphs against a set of conditions or constraints. The SHACL standard  
644 comprises a specification for a *shapes graph*, an RDF graph containing the constraint definitions —  
645 including but not limited to expected properties, values and types associated with properties and arity of  
646 properties — and a method for verifying if a target RDF graph meets those constraints.

647 A shapes graph contains a collection of *shapes*. A shape consists of a list of constraints and a *target*  
648 *declaration* which specifies which node or group of nodes the constraints apply to. SHACL constraints  
649 have many forms; rather than review the full range of possibilities (we refer the reader to [21] for detailed  
650 documentation on SHACL), we concentrate on the two main categories of SHACL shapes used in Brick+  
651 validation: relationship shapes and class shapes.

652 **Relationship shapes** are constraints that validate use of Brick+ relationships enumerated in Table 2.  
653 There is one shape for each domain and range property defined for each Brick+ relationship. The domain  
654 and range properties (denoted by `rdfs:domain` and `rdfs:range`) imply the class of the subject and  
655 object of the relationship, respectively. Validating against relationship shapes can alert authors of Brick+  
656 models of potential logical violations (see Figure 9).

657 Figure 10 contains two relationship shape definitions for the inverse relationships `brick:isPointOf`  
658 and `brick:hasPoint`. The top shape — `bsh:isPointOfDomainShape`, lines 1-4 — demonstrates  
659 the typical structure of a shape constraining the class of a relationship's subject. The implementation  
660 in SHACL is straightforward: line 2 indicates that the shape targets all nodes which are subjects of the  
661 `brick:isPointOf` relationship. The targeted nodes are called the *focus nodes* in SHACL parlance.  
662 Line 4 indicates that the focus node's class should be `brick:Point`.

663 The bottom shape — `bsh:hasPointRangeShape`, lines 6-11 — demonstrates the typical shape  
664 structure for constraining the type of the object of a relationship. Line 7 indicates that the shape targets all  
665 nodes which are subjects of the `brick:hasPoint` relationship. The composite structure on lines 8-11  
666 states that the objects of the `brick:hasPoint` relationship should have the class `brick:Point`.

---

```

1 bsh:vavRelationshipShape a sh:NodeShape ;
2 sh:targetClass brick:Variable_Air_Volume_Box ;
3 sh:property [
4   sh:path brick:feeds ;
5   sh:message "VAV boxes should feed an HVAC Zone" ;
6   sh:class brick:HVAC_Zone ] ;
7 sh:property [
8   sh:path brick:isFedBy ;
9   sh:message "VAV boxes should be fed by an AHU" ;
10  sh:class brick:Air_Handler_Unit ] .

```

---

**Figure 11.** A SHACL shape encoding the idiom that a VAV must be interposed between an AHU and an HVAC Zone through the use of the `brick:feeds` property

---

```

1 bsh:modelXYZ_VAVShape a sh:NodeShape ;
2 sh:targetClass brick:Model_XYZ_VAV ;
3 sh:property [
4   sh:path brick:hasPart ;
5   sh:message "Model XYZ VAVs must have a damper" ;
6   sh:class brick:Damper ] ;
7 sh:property [
8   sh:path brick:hasPoint ;
9   sh:message "Model XYZ VAVs must have a supply air temp sensor" ;
10  sh:class brick:Supply_Air_Temperature_Sensor ] ;
11 sh:property [
12  sh:path brick:hasPoint ;
13  sh:message "Model XYZ VAVs must have a heating temp setpoint" ;
14  sh:class brick:Heating_Temperature_Setpoint ] ;
15 sh:property [
16  sh:path brick:hasPoint ;
17  sh:message "Model XYZ VAVs must have a cooling temp setpoint" ;
18  sh:class brick:Cooling_Temperature_Setpoint ] ;
19 sh:property [
20  sh:path brick:hasPoint ;
21  sh:message "Model XYZ VAVs must have an air flow sensor" ;
22  sh:class brick:Supply_Air_Flow_Sensor ] ;
23 sh:property [
24  sh:path brick:hasPoint ;
25  sh:message "Model XYZ VAVs must have an air flow setpoint" ;
26  sh:class brick:Supply_Air_Flow_Setpoint ] .

```

---

**Figure 12.** A SHACL shape encoding the required parts and points for a theoretical “Model XYZ” variable air volume box.

667 **Class shapes** specify conditions on the properties and property values for a Brick+ class. Correspondingly,  
668 the focus nodes for a class shape are the set of instances of that class. We expect that most idiomatic shapes  
669 will be class shapes.

670 Figure 11 contains an example of an idiomatic class shape that encodes the requirement that all  
671 VAVs in a model instance must refer to a downstream HVAC zone and an upstream air handling  
672 unit. Validating a Brick+ model instance against this shape involves examining all of the instances  
673 of `brick:Variable_Air_Volume_Box` and its subclasses to see if the required properties exist and if  
674 the objects of those properties fulfill the class requirements.

675 Figure 12 is an example of a shape encoding the expected parts and points for a theoretical variable air  
676 volume box of a particular make and model. Applying this shape to a Brick+ model instance can help  
677 ensure that all instances of the equipment are modeled consistently.

678 Unlike the shapes in Figure 10, the shapes in Figure 11 and Figure 12 are not required for correct usage  
679 of Brick. Instead, adherence to these shapes might be a commissioning requirement for the Brick+ model  
680 produced for a site.

### 681 **6.3 Implementation**

682 We now describe how we have incorporated the SHACL standard into the development and usage of  
683 Brick+.

684 The abstract specification of Brick+ developed in §5 allows us to automatically generate relationship  
685 shapes for verifying correct usage of Brick relationships. There are currently 23 relationship shapes  
686 distributed with Brick+, but we expect this number will increase over time as the number of relationships  
687 and properties supported by Brick+ expands. These shapes are included as part of the Brick+ distribution  
688 and are organized under the abstract RDF namespace <https://brickschema.org/schema/1.1/>  
689 `BrickShape`, commonly abbreviated as `bsh`.

690 To perform validation of a model instance, we incorporate the excellent open-source PySHACL library [2]  
691 into a Brick-specific software module and augment it with some features specific to Brick+. The module  
692 exposes the validation functionality through a command-line tool, `brick_validate`, as well as a Python  
693 library. Shapes for validating correct usage of Brick+ are included in the library so validation against these  
694 shapes is always performed by default, without any additional configuration.

695 The primary feature offered by the Brick+ validation software module is a post-processing step applied  
696 to the output of the underlying PySHACL module. When the PySHACL validation process is complete,  
697 the Brick+ validation software attempts to find the offending triples and relevant context within the model  
698 instance. The software can then provide suggestions for how to repair the model to pass validation.

### 699 **6.4 Evaluation of Validation Tool**

700 To evaluate the efficacy of the validation approach and tools, we applied the Brick+ validation module to  
701 five reference models from the original Brick release [4]. Each of the models was converted to the Brick+  
702 edition of Brick through the migration process described in §7. The validation process found correctness  
703 violations in each graph, including:

- 704 • Incorrect type of subject or object as required by the property: this is one type of error that can be  
705 found through the application of relationship shapes
- 706 • Incorrect use of a relationship; for example, `brick:hasLocation` is used where  
707 `brick:hasPart` is more appropriate.
- 708 • Using a class declared to be in the Brick+ namespace that is not actually defined in the official Brick+  
709 release: this may not be a severe violation because we do expect that ad-hoc extensions to the Brick+  
710 ontology will take place “in the field”, but it is good to raise a warning that potentially unsupported  
711 classes exist in a model instance
- 712 • Failure to declare a type for an entity: this is an example of a correctness constraint that does not fall  
713 under the relationship shapes defined above

714 Validation of Brick model instances has long been a desired feature, but has been difficult to implement  
715 due to the lack of formalization of the Brick model itself. The introduction of Brick+ and its abstract  
716 specification makes description of correct and idiomatic usage natural to express as “shapes” within the

717 SHACL language. The validation of a Brick+ model instance using these shapes is simple to perform  
718 through the Brick+ software library. The Brick+ shapes and the validation process are captured online at  
719 <https://brickschema.readthedocs.io/en/latest/validate.html>.

## 7 BRICK MODEL MIGRATION

720 As the Brick ontology evolves it becomes increasingly important to handle the migration of a particular  
721 building model from one version to another. The migration should fulfill the following properties:

- 722 • **Complete:** The migration should handle the translation of *all* classes and relationships from one  
723 version of the ontology to another.
- 724 • **Semantics-preserving:** The migration should preserve the semantics of the original model when  
725 updating it to the new ontology wherever possible; the extent to which this can be fulfilled is determined  
726 in part by how well the ontology itself preserves the semantics of the older version.
- 727 • **Automatic:** The migration should minimize the amount of input and manual translation effort required  
728 of the model developer.

729 In this section, we present the design, implementation and evaluation of a tool for migrating Brick model  
730 developed against the prior Brick 1.0.x ontology versions [4, 5] to the Brick+ ontology developed in this  
731 paper.

732 While the older versions of Brick have more structure than Haystack, we can still adopt a similar approach  
733 for formalizing the relationship between Brick concepts and Brick+ concepts. Both the Brick-migration  
734 described in this section and the Haystack-inference described in §8 describe how these non-formal  
735 metadata standards can be defined in terms of the formal Brick+ definition.

### 7.1 Migration Strategies for Brick

737 We adopt two strategies for migrating models developed against Brick to the newer Brick+ ontology:  
738 migration of classes and migration of relationships.

739 **Class migration** consists of referring instances of Brick classes to the most appropriate Brick+ class.  
740 Most of the class names stayed the same between Brick and Brick+, meaning the migration can be  
741 performed through a simple 1-to-1 mapping of namespaces. Cases where the name of the class changed  
742 while the role and definition stayed the same are handled through the same mechanism.

743 For cases where there is not a 1-to-1 mapping between classes in the two versions, we adopt  
744 a *parametric* approach to migration. The most common case where 1-to-1 mapping fails is that  
745 of so-called “equipment-flavored” classes. The original Brick class structure included many class  
746 names — the majority of them Point classes — that incorporated the name of equipment. For  
747 example, the Brick class `AHU_Zone_Air_Temperature_Sensor` represents the concept of an  
748 `Zone_Air_Temperature_Sensor` associated with an air handling unit. The existence of this class  
749 raised an issue for practitioners: should they use the `Zone_Air_Temperature_Sensor` class with an  
750 `brick:isPointOf` relationship to an instance of the `Air_Handling_Unit` class, or should they sim-  
751 ply use the `AHU_Zone_Air_Temperature_Sensor` class? Brick+ addresses this issue by eliminating  
752 all “equipment-flavored” classes, preferring the explicit association of points to equipment through the  
753 `brick:isPointOf` relationship.

754 In order to preserve the semantics of the “equipment-flavored” classes, the migration tool must go  
755 beyond simply translating class names and now must add relationships as well. The migration tool adds the  
756 requisite relationships where the instance of the “equipment-flavored” point class already had a relationship

---

```

1 PREFIX brick_v_1_0_3: <https://brickschema.org/schema/1.0.3/Brick#>
2 PREFIX brick: <https://brickschema.org/schema/1.1/Brick#>
3 DELETE {
4   ?subject ?predicate brick_v_1_0_3:AHU_Zone_Air_Temperature_Sensor
5 } INSERT {
6   ?subject ?predicate brick_plus:Zone_Air_Temperature_Sensor
7 } WHERE {
8   ?subject ?predicate brick_v_1_0_3:AHU_Zone_Air_Temperature_Sensor
9 }

```

---

**Figure 13.** An example of a SPARQL 1.1 UPDATE query migrating a Brick 1.0.3 class to a Brick+ class.

757 to an instance of the appropriate equipment class. When the instance does *not* have a relationship to an  
 758 equipment instance, the migration module can either generate a temporary placeholder instance of the  
 759 equipment or raise a flag to the user to indicate the lack of one.

760 The change in class structure from a strict hierarchy (Brick) to a lattice (Brick+) is transparent to the  
 761 building models and thus does not need to be addressed by the migration process.

762 **Relationship migration** consists of replacing Brick relationships with the most appropriate Brick+  
 763 relationship. Brick+ preserves the relationships defined in Brick, so the migration tool only needs to handle  
 764 the translation of the namespace. Although Brick+ incorporates some new relationships, these are either  
 765 used solely within the definition of the class lattice (e.g. `brick:measures`), are not yet used by Brick  
 766 instances, or can be added automatically (e.g. `brick:hasTag`); therefore, the migration tool does not  
 767 need to handle these new relationships.

## 768 7.2 Implementation

769 The migration tool is implemented in Python and is included as part of the open source Brick+ distri-  
 770 bution<sup>6</sup>. The tool includes a set of *conversion queries* that implement the translation from one version of  
 771 Brick to another. A conversion query is phrased in the SPARQL 1.1 UPDATE language and is generated  
 772 from an underlying dictionary of the simple and parametric migrations described above. Figure 13 contains  
 773 an example of a conversion query. Each conversion query is parameterized by the source and destination  
 774 version of Brick; the query converts a given term from the source version of Brick to its migrated form  
 775 in the destination version. The conversion queries are incorporated into the migration algorithm, which  
 776 consists of the following steps:

- 777 1. Accept as input the source model, the source version of Brick, and the intended output version of  
 778 Brick.
- 779 2. If the migration tool contains a set of conversion queries for the provided source and destination  
 780 version, then continue with the *direct* translation. Otherwise, perform an *indirect* translation (described  
 781 below).
- 782 3. Ingest the source model into a database (such as a triplestore) that supports the required SPARQL 1.1  
 783 queries
- 784 4. Execute all of the conversion queries against the triple store
- 785 5. Serialize the edited model to the provided output file

---

<sup>6</sup> <https://github.com/BrickSchema/Brick>

Model Name	Classes	% Translated	Relationships	% Translated
Soda	34/34	100%	7/7	100%
GHC	79/80	98.75%	8/8	100%
Rice	57/57	100%	5/5	100%
EBU3B	217/217	100%	3/3	100%
GTC	570/582	97.94%	8/9	88.89%

**Table 3.** The completeness of the migration tool against five Brick reference models in terms of the *unique classes and relationships* in the source model.

Model Name	Classes	% Translated	Relationships	% Translated
Soda	1693/1693	100%	2078/2078	100%
GHC	9103/9112	99.90%	36458/36458	100%
Rice	632/632	100%	718/718	100%
EBU3B	6174/6174	100%	8392/8392	100%
GTC	1524/1526	99.21%	5309/5311	99.99%

**Table 4.** The completeness of the migration tool against five Brick reference models in terms of the *instances of classes and relationships* in the source model.

786 The migration tool maintains an RDF graph containing the details of all available conversions. In the case  
 787 where a direct migration is not available, the migration tool runs a shortest-path algorithm to determine a  
 788 sequence of intermediate versions through which the source model can be migrated so as to arrive at the  
 789 desired output version. Currently, the shortest-path algorithm uses the number of intermediate versions as  
 790 the distance metric.

### 791 **7.3 Evaluation of Migration Tool**

792 To evaluate the efficacy and completeness of the migration tool, we apply it to five of the original Brick  
 793 reference models published as part of [4] to translate them from version 1.0.2 to Brick+.

794 The results of applying the migration tool to the source models are enumerated in Tables 3 and 4.  
 795 The migration tool successfully converts 98% of the unique classes used in the models and 91% of the  
 796 relationships. The unmapped classes — those for which no conversion query existed — were left as-is: no  
 797 information was lost from the original models. These unmapped classes exist where the model authors  
 798 defined their own extensions to Brick.

799 This evaluation demonstrates that the migration tool is effective in handling the translation of classes and  
 800 relationships between past and current version of Brick. We believe that the methodology developed here  
 801 will allow the migration tool to perform effective migrations of models through future versions of Brick.

## 8 HAYSTACK–BRICK+ INFERENCE

802 To further evaluate how well a formal approach to metadata enables consistency, we examine how well the  
 803 Brick+ inference engine is able to extract and classify entities from a set of five Haystack models.

### 804 **8.1 Source Haystack Models**

805 We assemble a set of five Haystack models, each consisting of a set of tagged entities. Haystack model 1  
 806 is the “Carytown” reference model published by Project Haystack for a 3000 sq ft building in Richmond,  
 807 VA. Haystack models 2 and 3 are sample Haystack data models with for complex buildings, and thus  
 808 contain large numbers of specialized and non-standard tags [11]. Haystack models 4 and 5 represent two  
 809 office buildings on the UC Davis campus. Together, these five Haystack models represent a diverse family  
 810 spanning small to large buildings, differing numbers of custom tags, and different model modelers.



Site Name	Haystack Entities	Inferred Brick Entities	% Classified Entities	Unclassified Entities	Avg % Custom Tags per Entity	Unique Custom Tags
1	22	23	86.4%	3	7.4%	4
2	147	168	89.8%	15	5.0%	6
3	149	145	73.8%	39	6.6%	7
4	2183	1755	86.7%	290	17.6%	46
5	6474	6236	93.0%	451	19.5%	41

**Table 5.** Results of inferring Brick entities from tagged Haystack entities.

## 811 8.2 Haystack Inference Results

812 Table 5 contains the results of applying the Brick inference engine to the five Haystack models. When the  
813 inference engine splits Haystack entities into equipment and a point, the number of inferred Brick entities  
814 can exceed the number of original Haystack entities

815 The *% Classified Entities* column indicates the percentage of Haystack entities that were successfully  
816 classified by the Brick inference engine; the *Unclassified Entities* column contains the number of entities  
817 that were not classified. The majority of unclassified entities were such due to the use of non-standard  
818 tags that have no provided definition, and thus were not included in the Brick tag structure. The lowest-  
819 performing Haystack model, Site 3, represents a data center and contained a number of specialized lighting,  
820 HVAC and data center equipment and points that are not covered by the existing Haystack tag dictionary.

821 To understand the impact of informal modeling practices on interpretability and consistency, we examine  
822 the occurrence of non-standard tags in the five Haystack models; the results are contained in the *Avg %*  
823 *Custom Tags per Entity* column and *Unique Custom Tags* column, which shows the number of user-defined  
824 tags in each building, showing the same trend. Models 4 and 5 contain a higher incidence of custom  
825 tags because they contain detailed representations of HVAC systems, thus requiring additional vocabulary  
826 beyond what is defined in Haystack. The required vocabulary includes HVAC concepts not yet defined in  
827 Haystack (e.g., *differential* for *differential pressure*) and functional relationships outside  
828 the Haystack's scope, such as capturing spatial relationships.

829 Examination of the Haystack models reveals three patterns of inconsistent tagging. Firstly, the lexical  
830 overlap of tags (detailed in Table 1) leads to one tag being used incorrectly in place of another; for example,  
831 using *heat* instead of *heating*. Secondly, because there is no notion of a "sufficient" tag set for a  
832 concept, several entities have ambiguous interpretations due to partial tagging. For example, several  
833 entities have the *differential* tag, but do not have a tag to clarify the quantity (e.g. *pressure*,  
834 *temperature*). Thirdly, the lack of compositional rules resulted in the ad-hoc creation of site-specific  
835 "compound" tags: models 4 and 5 use a custom *spMax* tag instead of the Haystack-defined *sp* and *max*  
836 tags to differentiate between setpoints and parameters.

## 837 8.3 Brick Inference Results

838 To complete our evaluation of Brick+, we measure the number of properties that can be inferred from the  
839 entities in existing Brick models. Because Brick models already have a formal representation, the inference  
840 engine does not need to apply the cleaning or splitting phases of the inference procedure (§5) and can rely  
841 entirely upon the existing features of the OWL DL reasoner.

Ontology	Inferred Properties	
	(Total)	(Avg per entity)
Brick	122,552	2.94/35.44
Brick+	201,266	4.79/35.55

**Table 6.** Number of inferred properties for all entities across 104 Brick models in Brick and Brick+.

842 We executed the Hermit [16] OWL reasoner on 104 existing Brick models from the Mortar testbed [14]  
 843 using the existing Brick ontology and our proposed Brick+ ontology, and computed the number of inferred  
 844 properties. The results are summarized in Table 6: Brick+ was able to infer almost 80,000 more properties  
 845 than Brick over the 42,681 entities contained in the Brick models. Brick+ was able to infer all the same  
 846 properties as Brick, but was able to infer tags and behavioral properties as well.

#### 847 **8.4 Discussion of Inference Results**

848 Our results demonstrate that Brick+ is able to infer 73-93% of entities in Haystack models that follow a  
 849 canonical tagging scheme, and can infer more semantic properties about entities in Brick models than the  
 850 previous release of Brick. Recall that Brick+'s inference engine does not currently infer all possible classes  
 851 from a Haystack model; rather, it formalizes a particular *interpretation* and *organization* of Haystack tags  
 852 applied to entities. Haystack tags in real-world Haystack models are highly idiosyncratic, due in part to  
 853 site-specific invention of tags to cover concepts and relationships not defined in the Haystack tag dictionary.  
 854 This suggests that Brick+'s inference engine will not be able to fully classify each Haystack entity without  
 855 additional automated metadata construction techniques [22, 9]. Our results support this hypothesis: an  
 856 ontology-based inference engine demonstrates decent performance against the informal Haystack data  
 857 model, but, as expected, custom tags inhibit inference.

## 9 CONCLUSION

858 Interoperability for building applications requires metadata standards that are semantically sound, rich and  
 859 extensible. Tags provide an intuitive and informal model, but lack rules for composition and validation that  
 860 enable consistent, interpretable usage. Brick+ constructs a compositional model of metadata where tags are  
 861 part of a type system with an underlying formalism based on lattice theory. This enables new algorithmic  
 862 methods for checking validity, consistency and compositional correctness that is necessary for building a  
 863 new class of scalable and portable building applications.

864 This paper presents a qualitative analysis of the popular Haystack tagging system and demonstrates  
 865 how its ad-hoc nature inhibits the consistent description of building systems. To address these issues, we  
 866 have introduced Brick+, a refinement of the Brick ontology with clear formal semantics that permits the  
 867 inference of well-defined classes from unstructured tags. Brick+ helps to bridge the gap between existing  
 868 ad-hoc, informal metadata practices and interoperable formal systems; this establishes a foothold for the  
 869 continued co-development of the Brick and Haystack metadata standards.

870 Brick+ is open-source and is in the process of being adopted as the authoritative implementation of Brick.  
 871 The Brick+ ontology, generation framework, source code of the inference engine, and the Haystack dataset  
 872 are all available online at <https://github.com/BrickSchema/Brick>.

## 10 ACKNOWLEDGEMENTS

873 This research is supported in part by California Energy Commission EPC-15-057, Department of Energy  
 874 grant EE-0007685, NSF grants CNS-1526841 and CSR-1526237, and the CONIX Research Center, one of  
 875 six centers in JUMP, a Semiconductor Research Corporation (SRC) program sponsored by DARPA. The

876 opinions expressed belong solely to the authors, and not necessarily to the authors' employers or funding  
877 agencies.

## CONFLICT OF INTEREST STATEMENT

878 The authors declare that the research was conducted in the absence of any commercial or financial  
879 relationships that could be construed as a potential conflict of interest.

## AUTHOR CONTRIBUTIONS

880 The Author Contributions section is mandatory for all articles, including articles by sole authors. If an  
881 appropriate statement is not provided on submission, a standard one will be inserted during the production  
882 process. The Author Contributions statement must describe the contributions of individual authors referred  
883 to by their initials and, in doing so, all authors agree to be accountable for the content of the work.

## FUNDING

884 This research is supported in part by California Energy Commission EPC-15-057, Department of Energy  
885 grant EE-0007685, NSF grants CNS-1526841 and CSR-1526237, and the CONIX Research Center, one of  
886 six centers in JUMP, a Semiconductor Research Corporation (SRC) program sponsored by DARPA. The  
887 opinions expressed belong solely to the authors, and not necessarily to the authors' employers or funding  
888 agencies.

## DATA AVAILABILITY STATEMENT

889 The datasets generated for this study can be found on GitHub at [https://github.com/  
890 BrickSchema/brick-examples](https://github.com/BrickSchema/brick-examples)

## REFERENCES

- 891 [1] [Dataset] (2018). Project Haystack  
892 [2] [Dataset] (2020). PySHACL  
893 [3] [Dataset] American Society of Heating, Refrigerating and Air-Conditioning Engineers  
894 (2018). ASHRAE's BACnet Committee, Project Haystack and Brick Schema Collaborating to Provide Unified Data Semantic Modeling Solution. [http://web.archive.org/web/20181223045430/https://www.ashrae.org/about/news/2018/  
895 ashrae-s-bacnet-committee-project-haystack-and-brick-schema-collaborating](http://web.archive.org/web/20181223045430/https://www.ashrae.org/about/news/2018/ashrae-s-bacnet-committee-project-haystack-and-brick-schema-collaborating)  
896  
897  
898  
899 [4] Balaji, B., Bhattacharya, A., Fierro, G., Gao, J., Gluck, J., Hong, D., et al. (2016). Brick: Towards  
900 a unified metadata schema for buildings. In *Proceedings of the ACM International Conference on  
901 Embedded Systems for Energy-Efficient Built Environments (BuildSys)*. ACM  
902 [5] Balaji, B., Bhattacharya, A., Fierro, G., Gao, J., Gluck, J., Hong, D., et al. (2018). Brick: Metadata  
903 schema for portable smart building applications. *Applied energy* 226, 1273–1292  
904 [6] Balaji, B., Verma, C., Narayanaswamy, B., and Agarwal, Y. (2015). Zodiac: Organizing large  
905 deployment of sensors to create reusable applications for buildings (ACM), 13–22  
906 [7] Bechhofer, S., van Harmelen, F., Hendler, J., Horrocks, I., McGuinness, D. L., Patel-Schneider, P. F.,  
907 et al. (2004). *OWL Web Ontology Language Reference*. Tech. rep., W3C, [http://www.w3.org/TR/owl-  
908 ref/](http://www.w3.org/TR/owl-ref/)

- 909 [8] Bhattacharya, A., Ploennigs, J., and Culler, D. (2015). Short paper: Analyzing metadata schemas for  
910 buildings: The good, the bad, and the ugly. In *Proceedings of the 2nd ACM International Conference*  
911 *on Embedded Systems for Energy-Efficient Built Environments* (ACM), 33–34
- 912 [9] Bhattacharya, A. A., Hong, D., Culler, D., Ortiz, J., Whitehouse, K., and Wu, E. (2015). Automated  
913 metadata construction to support portable building applications. In *Proceedings of the 2nd ACM*  
914 *International Conference on Embedded Systems for Energy-Efficient Built Environments* (ACM), 3–12
- 915 [10] Capozzoli, A., Piscitelli, M. S., Gorrino, A., Ballarini, I., and Corrado, V. (2017). Data analytics  
916 for occupancy pattern learning to reduce the energy consumption of hvac systems in office buildings.  
917 *Sustainable Cities and Society* 35, 191–208
- 918 [11] [Dataset] Coffey, P. (2019). Project Haystack Example Data Models. [http://web.archive.](http://web.archive.org/web/20190626161742/https://patrickcoffey.bitbucket.io/)  
919 [org/web/20190626161742/https://patrickcoffey.bitbucket.io/](http://web.archive.org/web/20190626161742/https://patrickcoffey.bitbucket.io/)
- 920 [12] Dong, B., Lam, K., Huang, Y., and Dobbs, G. (2007). A comparative study of the ifc and gbxml  
921 informational infrastructures for data exchange in computational design support environments. In  
922 *Building Simulation 2007, BS 2007*
- 923 [13] Fierro, G., Koh, J., Agarwal, Y., Gupta, R. K., and Culler, D. E. (2019). Beyond a house of sticks:  
924 Formalizing metadata tags with brick. In *Proceedings of the 6th ACM International Conference on*  
925 *Systems for Energy-Efficient Buildings, Cities, and Transportation*. 125–134
- 926 [14] Fierro, G., Pritoni, M., AbdelBaky, M., Raftery, P., Peffer, T., Thomson, G., et al. (2018). Mortar:  
927 an open testbed for portable building analytics. In *Proceedings of the 5th Conference on Systems for*  
928 *Built Environments* (ACM), 172–181
- 929 [15] Gao, J., Ploennigs, J., and Berges, M. (2015). A data-driven meta-data inference framework for  
930 building automation systems (ACM), 23–32
- 931 [16] Glimm, B., Horrocks, I., Motik, B., Stoilos, G., and Wang, Z. (2014). HermiT: an OWL 2 reasoner.  
932 *Journal of Automated Reasoning* 53, 245–269
- 933 [17] [Dataset] Guha, R. and Brickley, D. (2014). RDF schema 1.1
- 934 [18] Hardin, D., Stephan, E. G., Wang, W., Corbin, C. D., and Widergren, S. E. (2015). *Buildings*  
935 *interoperability landscape*. Tech. rep., Pacific Northwest National Lab.(PNNL), Richland, WA  
936 (United States)
- 937 [19] Hong, D., Wang, H., Ortiz, J., and Whitehouse, K. (2015). The building adapter: Towards quickly  
938 applying building analytics at scale (ACM), 123–132
- 939 [20] Jahn, M., Schwartz, T., Simon, J., and Jentsch, M. (2011). Energypulse: tracking sustainable behavior  
940 in office environments. In *Int. Conf. on Energy-Efficient Computing and Networking* (ACM), 87–96
- 941 [21] Knublauch, H. and Kontokostas, D. (2017). Shapes constraint language (shacl). *W3C Candidate*  
942 *Recommendation* 11
- 943 [22] Koh, J., Balaji, B., Sengupta, D., McAuley, J., Gupta, R., and Agarwal, Y. (2018). Scrabble:  
944 transferrable semi-automated semantic metadata normalization using intermediate representation. In  
945 *Proceedings of the 5th Conference on Systems for Built Environments* (ACM), 11–20
- 946 [23] Lange, H., Johansen, A., and Kjærgaard, M. B. (2018). Evaluation of the opportunities and limitations  
947 of using ifc models as source of building metadata. In *Proceedings of the 5th Conference on Systems*  
948 *for Built Environments*. 21–24
- 949 [24] Lassila, O. and Swick, R. R. (1999). Resource description framework (RDF) model and syntax  
950 specification
- 951 [25] Mathes, A. (2004). Folksonomies - Cooperative Classification and Communication Through Shared  
952 Metadata , 14

- 953 [26] Mims, N., Schiller, S. R., Stuart, E., Schwartz, L., Kramer, C., and Faesy, R. (2017). Evaluation  
954 of U.S. Building Energy Benchmarking and Transparency Programs: Attributes, Impacts, and Best  
955 Practices doi:10.2172/1393621
- 956 [27] OSTI (2016). *The National Opportunity for Interoperability and its Benefits for a Reliable, Robust,  
957 and Future Grid Realized Through Buildings*. Tech. rep. doi:10.2172/1420233
- 958 [28] Passant, A. (2007). Using ontologies to strengthen folksonomies and enrich information retrieval in  
959 weblogs. In *International Conference on Weblogs and Social Media*
- 960 [29] Passant, A. and Laublet, P. (2008). Meaning of a tag: A collaborative approach to bridge the gap  
961 between tagging and linked data. *LDOW* 369
- 962 [30] Pauwels, P. and Terkaj, W. (2016). Express to owl for construction industry: Towards a recommendable  
963 and usable ifcowl ontology. *Automation in Construction* 63, 100–133
- 964 [31] Piette, M. A., Ghatikar, G., Kiliccote, S., Koch, E., Hennage, D., Palensky, P., et al. (2009). *Open  
965 automated demand response communications specification (Version 1.0)*. Tech. rep., Ernest Orlando  
966 Lawrence Berkeley National Laboratory, Berkeley, CA (US)
- 967 [32] Privara, S., Cigler, J., Váňa, Z., Oldewurtel, F., Sagerschnig, C., and Žáčková, E. (2013). Building  
968 modeling as a crucial part for building predictive control. *Energy and Buildings* 56, 8–22
- 969 [33] [Dataset] Project Haystack (2019). Project Haystack Documentation: Defs. [http:  
970 //web.archive.org/web/20190629183024/https://project-haystack.dev/  
971 doc/docHaystack/Defs](http://web.archive.org/web/20190629183024/https://project-haystack.dev/doc/docHaystack/Defs)
- 972 [34] [Dataset] Project Haystack (2019). Project Haystack Documentation: VFDs. [http:  
973 //web.archive.org/web/20190629182856/https://project-haystack.org/  
974 doc/VFDs](http://web.archive.org/web/20190629182856/https://project-haystack.org/doc/VFDs)
- 975 [35] Rasmussen, M. H., Pauwels, P., Hviid, C. A., and Karlshøj, J. (2017). Proposing a central aec ontology  
976 that allows for domain specific extensions. In *2017 Lean and Computing in Construction Congress*
- 977 [36] Reiter, R. (1981). On closed world data bases. In *Readings in artificial intelligence* (Elsevier).  
978 119–140
- 979 [37] Roth, S. (2014). Open green building XML schema: A building information modeling solution for our  
980 green world, gbXML schema (5.12)
- 981 [38] Schein, J., Bushby, S. T., Castro, N. S., and House, J. M. (2006). A rule-based fault detection method  
982 for air handling units. *Energy and Buildings* 38, 1485–1492
- 983 [39] Sturzenegger, D., Gyalistras, D., Morari, M., and Smith, R. S. (2012). Semi-automated modular  
984 modeling of buildings for model predictive control (ACM), 99–106
- 985 [40] Tang, S., Shelden, D. R., Eastman, C. M., Pishdad-Bozorgi, P., and Gao, X. (2020). Bim assisted  
986 building automation system information exchange using bacnet and ifc. *Automation in Construction*  
987 110, 103049. doi:https://doi.org/10.1016/j.autcon.2019.103049
- 988 [41] [Dataset] W3C (2007). Punning
- 989 [42] Wille, R. (1992). Concept lattices and conceptual knowledge systems. *Computers & mathematics  
990 with applications* 23, 493–515
- 991 [43] Wille, R. (2009). Restructuring lattice theory: an approach based on hierarchies of concepts. In  
992 *International Conference on Formal Concept Analysis* (Springer), 314–339
- 993 [44] Yang, Q. and Zhang, Y. (2006). Semantic interoperability in building design: Methods and tools.  
994 *Computer-Aided Design* 38, 1099 – 1112. doi:https://doi.org/10.1016/j.cad.2006.06.003