

PAPER • OPEN ACCESS

## Indoor haze particulate control using knowledge graphs within self-optimizing HVAC control systems

To cite this article: Kevin Luwemba Mugumya *et al* 2020 *IOP Conf. Ser.: Earth Environ. Sci.* **489** 012006

View the [article online](#) for updates and enhancements.

### You may also like

- [Ventilation requirements and energy aspects: the case of hospitals](#)  
Giorgos Panaras, Risofile Gropca and  
Giannis Papadopoulos
- [Theoretical and experimental investigation of ventilation rates and their relation with IAQ and thermal comfort in university classrooms during SARS-COV-2 pandemic](#)  
Giannis Papadopoulos, Apostolos  
Nikolentzos, Evangelos I. Tolis *et al.*
- [Indoor air quality performance in air-conditioned museum gallery](#)  
R Sulaiman, S N Kamaruzzaman and Y H  
Yau

Join the Society  
Led by Scientists,  
for *Scientists Like You!*



Thomas Edison  
ECS Member



The  
Electrochemical  
Society

Advancing solid state &  
electrochemical science & technology



Allen J. Bard  
ECS Member

# Indoor haze particulate control using knowledge graphs within self-optimizing HVAC control systems

Kevin Luwemba Mugumya<sup>1</sup>, Jing Ying Wong<sup>2</sup>, Andy Chan<sup>3</sup>,  
Chun-Chieh Yip<sup>4</sup>, Shams Ghazy<sup>5</sup>

<sup>1,2,3,5</sup>Department of Civil Engineering, University of Nottingham Malaysia, 43500 Semenyih  
Selangor Malaysia

<sup>4</sup>Department of Civil Engineering, University Tunku Abdul Rahman Sungai Long Campus,  
43500 Kajang, Selangor, Malaysia

E-mail: <sup>1</sup>keey4kly@nottingham.edu.my (corresponding author)

**Abstract.** Transboundary haze pollution in South East Asia is posing a threat to conventional design of buildings yet indoor air pollution from haze particulate infiltration has still received less attention in Malaysia compared to haze pollution outdoors. Because of this minimal research effort, indoor building environments have increasingly become very complex environments for facility managers to monitor and control due to the corresponding growth in heterogeneity within building behavioural information and monitoring (sensory) systems. As a solution to this and part of an ongoing study, this paper presents the preliminary process of modelling heterogeneous building information related to indoor air quality (IAQ) (building envelope, sensors, contaminant properties, geometry, occupancy schedules and weather data) within modular and extensible semantic web knowledge graphs (KG). This work argues that this data model can preserve the existential and latent parametric relationships within such information thereafter availing an accurate representation of the heterogeneous state-space in machine learning workflows of self-learning building monitors and controllers. Compared to the conventional homogenous feature vectors, KGs hold sufficient context-aware semantics for an algorithmic building control system to smartly monitor the IAQ and autonomously learn to adapt air handling units towards occupant comfort in an energy efficient manner. Specifically, this paper highlights the high-level implementation process of KGs within the deep Q-learning process of the aforementioned control systems. Finally, a brief discussion is provided on how this process reduces the complexity that facility managers face while operating their IAQ control systems followed by the conclusions and future work to be carried out in this study.

**Keywords:** Indoor air quality, sensors, building information modelling, deep Q learning, haze, energy efficiency

## 1. Introduction

Over the years, building designers have resorted to utilizing airtight building envelopes that favour optimal energy use and less contamination by outdoor air. To ensure optimum health of building occupants, natural ventilation (using operable windows) is necessary for exchanging indoor air with the outdoor [1]. However, due to the tropical weather and scorching temperatures within many South East Asian (SEA) countries like Malaysia, such an air exchange is done via mechanical air handling units which are an integral part of the heating, ventilation and air conditioning (HVAC) systems installed within a building.



Transboundary haze pollution in SEA has been a major cause of public concern in Malaysia due to the continuous burning of biomass in neighbouring Indonesia from which particulate matter (PM) is transported by oceanic currents to many inland areas of Malaysia [2]. Besides affecting atmospheric visibility, haze particles pose adverse impacts on public health such as coughing, nasal congestion, breathing difficulties, sore eyes and increased risk of lung cancer [3]. During haze episodes in Malaysia, most studies focus on estimating population exposure to PM<sub>2.5</sub> by examining only outdoor PM<sub>2.5</sub> concentrations yet during this time, Malaysians spend most of their time indoors [4]. Because outdoor haze particles of less than 2.5 micrometers can bypass many building filtration shields and get trapped within the building envelope [5, 6], Malaysians face a risk of long term exposure while indoors if no proper measures are taken to minimize the trapped particulates [7, 8]. The high Air Pollution Index (API) values that were recorded in several states around Malaysia during the 2015 haze episode [4] and the following years has raised concern among engineers to design more resilient indoor building environments [9, 6].

Intelligent building systems have progressively been adopted to smartly monitor the indoor built environment while autonomously controlling HVAC systems to achieve acceptable IAQ in an energy efficient manner [10, 11, 12, 13, 14]. Many of these systems rely on deep learning [13] and reinforcement learning methods to ingest building data while learning how to optimize the IAQ in a stochastic indoor built environment.

Contaminant transport and concentration within the building envelope depends on several factors like zone occupancy, infiltration rate, air flow patterns, wind speed, HVAC equipment operation schedules, indoor-outdoor temperature not forgetting the ever-changing weather patterns. Most of these factors are inherently heterogenous and unless captured in a way that self-learning building control systems can make good use of, the rich and complex existential relationships within are underutilized and intended optimum control behavior is rarely achieved [15].

As part of the ongoing efforts by the Linked Building Data Community Group (LBDCG) [16] to semantically enrich the heterogenous Architecture Engineering and Construction (AEC) processes, this paper presents the use of the *Resource Description Framework* (RDF) as the default data model by which self-learning HVAC control systems can holistically consume heterogenous data to achieve smarter and energy efficient minimization of haze contaminant concentrations in indoor built environments.


## 2. Methodology

### 2.1. Data modelling

The RDF data model is utilized to encode the domain information associated with contaminant control within indoor spaces. RDF provides a formal way of semantically describing and linking this heterogeneous information by using statements called triples [17]. Syntactically, these take the ‘*subject-predicate-object*’ format as shown in the figure 1 below.

Every resource (subject, predicate and object) in the RDF graph (*knowledge graph/ontology*) is uniquely identified using a Uniform Resource Identifier (URI) to eliminate any ambiguity when deployed online. The RDF triples are stored/ serialized in a compact form using Turtle [18] which is light-weight on the web and relatively human-readable.

Towards achieving the goal of ontology reuse, this research is going to refer to already existing ontologies that are readily available, extensible and managed by the World Wide Web Consortium (W3C) while modelling information related to contaminants within a building’s interior spaces. The exchange schema for commonly adopted Building Information Models (BIMs) already exists in RDF format [19] as proposed by the Linked Building Data Community Group (LBDCG). However considering that this is a broad ontology encapsulating everything



**Figure 1:** RDF triples in the form **subject-predicate-object**. The arrow implies directionality of the relationship

about a building, it is not used for this work in its entirety. Rather, smaller extensible modules are adopted, extended and integrated accordingly. For this task, specific modelling emphasis is placed on the concepts below;

- Building topology, geometry and zoning.
- Sensors in the building that monitor air quality (PM<sub>2.5</sub> concentrations).
- Indoor-outdoor temperature and occupancy timings.
- HVAC equipment high-level control information and operation schedules.

The choice of modular ontologies adopted is based on efficacy for modelling the information above into one comprehensive ontology referred to arbitrarily as ERLO (Energy Reinforcement Learning Ontology) which will form the basis for developing the learning state-space of the autonomous HVAC controller.

## 2.2. Reinforcement Learning Process for the HVAC monitor and controller

Reinforcement learning (RL) in general consists of an agent interacting in an environment, learning what action  $\mathbf{a}$  to take depending on the state  $\mathbf{s}$  of the environment [20]. The learning process is through trial and error with a reward  $\mathbf{r}$  for taking desirable actions. The goal of the agent is to maximize long term reward through *exploration and exploitation* of the action space.

For a stochastic indoor building environment with many possible contamination states, this RL process is combined with a *deep neural network/ function approximator* that can deal more efficiently with such large state-spaces. Specifically the *Q-learning model-free RL algorithm* is adopted with a neural network through a process called *deep Q-learning* [21] that accepts the state-space as input and outputs an approximation of the Q-values that determine how good it is for the learning agent (HVAC controller) to perform a certain action in a given indoor environment state (contamination level). These Q-values are learned progressively and updated periodically according to a specified learning rate  $\alpha \in (0,1)$  which determines the extent to which newly learned Q-values override old ones. This process happens iteratively according to equation (1) until convergence at optimal Q-values is achieved, a state at which the agent now knows the best control actions to take for each possible environment state.


$$Q^{new}(s_t, a_t) = (1 - \alpha) \underbrace{Q(s_t, a_t)}_{\text{old value}} + \alpha \overbrace{\left( r_{t+1} + \gamma \max_{a'} Q(s', a') \right)}^{\text{learned value}} \quad (1)$$

This contaminant control RL problem is modelled as a *finite Markov Decision Process (MDP)* [22] briefly defined as a 5-tuple  $(\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma)$ .  $\mathcal{S}$  being the set of states/ state-space the indoor building environment can be in,  $\mathcal{A}$  being the set of possible actions (action-space) the HVAC

controller can take to reduce the PM concentrations indoors,  $\mathcal{P}$  being the probability distribution that governs the environment dynamics/ transition from one state to another,  $\mathcal{R}$  being the reward function that returns the reward received when the agent transitions from one state to the next and  $\gamma \in (0,1)$  being the discount rate deciding the value of future rewards earned by the agent.

### 2.3. Adoption strategy of the RDF data model within the deep Q learning framework of the HVAC controller

Rather than using conventional feature vectors to populate the input state-space tensor of the HVAC learning agent, the RDF knowledge graph (ERLO) from section 2.1 will be translated into a *third order tensor* that this research argues to preserve the latent relationships within the concepts presented earlier affecting contaminant control within indoor building spaces (high-level schematic process shown in figure 2).



**Figure 2:** Schematic implementation of ERLO in the learning process of the HVAC controller

## 3. Results and discussion

### 3.1. Ontology development

This section presents the workflow and results from modelling the heterogeneous building information (envelope, sensors, contaminant properties, geometry, schedules and weather data) using RDF within semantic web knowledge graphs. A discussion is also provided on how such a data model accurately represents the heterogeneous state-space ingestible by an algorithmic HVAC control system to smartly monitor the indoor environment and maintain acceptable IAQ in the presence of trapped haze particulates. The modular ontologies used to model the above information are presented below.

#### 3.1.1. Building Ontology Topology

BOT [23] is the central AEC ontology for defining relationships between the major sub-components of a building (site, building, storey, space and element) while offering extensibility mechanisms to other domains. It is chosen because of its simplicity and adaptability to existing and non-existing buildings. In BOT, a building consists of zones in a hierarchy. The subclass of a zone being a site which contains a building(s) which contains storey(s) which contain space(s) in that hierarchical order. A zone can be adjacent to another zone or even contain other zones. It can also be bounded by physical building elements or even contain them. Explicit details about this ontology can be found in [24]. The goal is to extend BOT to other domain ontologies like sensors and automation either by specifying subclasses or sub properties of BOT concepts as shown in figure 3 and 4. The classes adopted for this data modelling are shown in table 1.

#### 3.1.2. Semantic Sensor Ontology (SSN)

SSN [25] is chosen for describing sensors that monitor indoor air quality their observations,

```

1 @prefix bot: <https://w3id.org/bot#> .
2 @prefix sosa: <http://www.w3.org/ns/sosa/> .
3 @prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .
4 @prefix owl: <http://www.w3.org/2002/07/owl#> .
5
6 sosa:Sensor a owl:Class ;
7 rdfs:subClassOf bot:Element .
8

```

**Figure 3:** Turtle snippet showing a `bot:Element` being extended by a SOSA element (`sosa:Sensor`) using the `rdfs:subClassOf` property.

```

1 @prefix bot: <https://w3id.org/bot#> .
2 @prefix sosa: <http://www.w3.org/ns/sosa/> .
3 @prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .
4 @prefix owl: <http://www.w3.org/2002/07/owl#> .
5
6 sosa:Sensor a owl:Class .
7 sosa:isObservedBy a owl:ObjectProperty ;
8 rdfs:subPropertyOf bot:containsElement ;
9 rdfs:range sosa:Sensor .
10

```

**Figure 4:** Turtle snippet showing a BOT property `bot:containsElement` being extended by a SOSA property (`sosa:isObservedBy`) using the `rdfs:subPropertyOf` property.

**Table 1:** BOT classes and properties to be adopted

Classes (domain)	Properties	Classes (range)
<code>bot:Zone</code>	<code>bot:containsZone</code>	<code>bot:Zone</code>
	<code>bot:adjascentZone</code>	<code>bot:Zone</code>
<code>bot:Site</code>	<code>bot:hasBuilding</code>	<code>bot:Building</code>
<code>bot:Building</code>	<code>bot:hasStorey</code>	<code>bot:Storey</code>
<code>bot:Storey</code>	<code>bot:hasSpace</code>	<code>bot:Space</code>
<code>bot:Space</code>	<code>bot:containsElement</code>	<code>bot:Element</code>
<code>bot:Element</code>	<code>bot:hostsElement</code>	<code>bot:Element</code>

properties and actuators while utilizing a lightweight but self-contained ontology at its core (Sensor, Observation, Sample and Actuator-SOSA) [25] for its elementary classes. This can be used to extend the building contamination zones from BOT with sensory information as already demonstrated in figure 3 and 4. Some of the classes from SSN and SOSA utilized within this research are `sosa:Observation`, `sosa:Sensor`, `sosa:ActuatableProperty`, `ssn:Stimulus`, `sosa:FeatureOfInterest`, `sosa:ObservableProperty`, `sosa:sample`, `ssn:Property` while the properties adopted include `sosa:observes`, `sosa:observedProperty`, `sosa:isObservedBy`, `sosa:madeObservation`, `sosa:madebySensor`, `ssn:detects`, `sosa:actsOnProperty`. Further details on the explicit property restrictions, domains and ranges can be found in [25].

### 3.1.3. Ontology for Property Management (OPM)

In addition to SSN and SOSA, the OPM ontology [26] is adopted to model properties that evolve over time like indoor-outdoor temperature variations, PM concentrations, zone infiltration rate and occupancy schedule information. An example is shown in figure 5 declaring `<someROOM>` as a `bot:Space` and assigning an area property via `opm:hasPropertyState`, `opm:Assumed` and `opm:CurrentPropertyState`. This `<someROOM>` can now contain a sensor



via the `bot:containsElement` property since `sosa:Sensor` has already been declared in figure 3 as a subclass of `bot:Element`. This means that the sensor in `<someROOM>` can now benefit from the existing relationship of area assigned by OPM. The specific classes and properties to be utilized from OPM depend on the required merge with SSN and SOSA however, full details of this ontology are presented in [26].

```

1  @prefix cdt:    <http://w3id.org/lindt/custom_datatypes#> .
2  @prefix ex:     <https://example.org/opmTest/> .
3  @prefix bot:    <https://w3id.org/bot#> .
4  @prefix opm:    <https://w3id.org/opm#> .
5  @prefix prov:   <http://www.w3.org/ns/prov#> .
6  @prefix props:  <https://w3id.org/product/props/> .
7  @prefix xsd:    <http://www.w3.org/2001/XMLSchema#> .
8
9  <someROOM> a bot:Space;
10 props:area ex:property_efafede9-53d3-425f-89f3-ed8a40b5995a .
11
12 ex:property_efafede9-53d3-425f-89f3-ed8a40b5995a
13 a opm:Property ;
14 opm:hasPropertyState ex:state_8e889bba-bc25-4b1a-ac34-330f776892fa .
15
16 ex:state_8e889bba-bc25-4b1a-ac34-330f776892fa
17 a opm:Assumed , opm:CurrentPropertyState ;
18 schema:value "20 m2"^^cdt:area ;
19 prov:generatedAtTime "2018-05-28T16:41:17.711+02:00"^^xsd:dateTime .' }
20

```

**Figure 5:** Turtle snippet showing a BOT (`bot:Space`) being extended by OPM

### 3.1.4. Building Product Ontology (BPO)

This ontology is utilized for the description of building elements like HVAC without the inclusion of unnecessary geometry. This ontology works in unison with the buildingSMART Data Dictionary (bsDD) for classification purposes and should geometric descriptions of HVAC systems be required, the File Ontology for Geometry formats (FOG) can be employed as an extension [27].

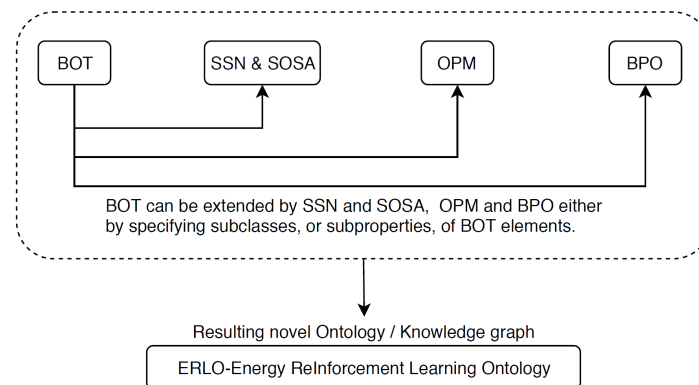
### 3.1.5. The encapsulating ontology (ERLO-Energy Reinforcement Learning Ontology)

After extending BOT in 3.1.1 to 3.1.4 (see figure 6), the resulting ontology/ knowledge graph ERLO, aims to encapsulate all the domain concepts that can accurately capture the heterogeneous learning state-space while preserving all the complex latent relationships for an autonomous HVAC control system to learn how to monitor and minimize haze  $PM_{2.5}$  concentration indoors in a smart and energy-efficient manner. Because this is part of an ongoing study, in the next section 3.2 only the high-level process is provided on how ERLO is incorporated into the deep Q-learning process of an HVAC controller.

## 3.2. High-level implementation of the knowledge graph (ERLO) within the reinforcement learning process of an HVAC controller

### 3.2.1. The learning state-space

This write-up is choosing to ignore the HVAC control action-space not because it is not important, but rather because the choice of optimal control actions taken by the learning agent highly depend on the richness of the learning state-space and how well the heterogeneous relationships within are preserved. This research argues that ERLO is a data model that can




**Figure 6:** Development of ERLO by extending the core ontology BOT with SSN, SOSA, OPM and BPO

achieve this as compared to conventional feature vectors that only fair well with homogenous information like sound, text and images. The states that ERLO can represent include building zone occupancy, indoor-outdoor temperature, current physical time, PM<sub>2.5</sub> levels, HVAC operation scheduling, variations in weather data and geolocation.

### 3.2.2. Modelling the learning state-space tensor for the HVAC controller using ERLO

The deep Q-learning process introduced in 2.2 accepts as input, the state-space from which an HVAC agent can learn to predict the best control actions in each observed contamination state of the indoor building environment by continuously updating the learned Q-value in equation 2.0 until convergence is achieved at optimality. The input state-space is represented as a tensor which is simply a vector or matrix of n-dimensions. Rather than using conventional feature vectors to populate this tensor, the RDF knowledge graph (ERLO) from section 2.1 will be translated into a *third order tensor* for this purpose. A typical RDF triple has 3 elements (see figure 1) and a third-order tensor can be used to map them, 2 orders for the entities (**subject** and **object**) and another order for the **predicate** (property). The intersection of all three orders within the tensor matrix represents a single RDF statement from ERLO as shown in figure 7.



**Figure 7:** Third-order tensor representation of a knowledge graph with one RDF statement “Sensor ReadsValue TempValue” shown for clarity.



It is important to note that the above tensor holds all possible combinations of RDF statements even those which are false. To deal with this issue, principles from adjacency matrix are used whereby the value at any intersection holds the truth value of that statement (1.0 if it is true and 0 otherwise).

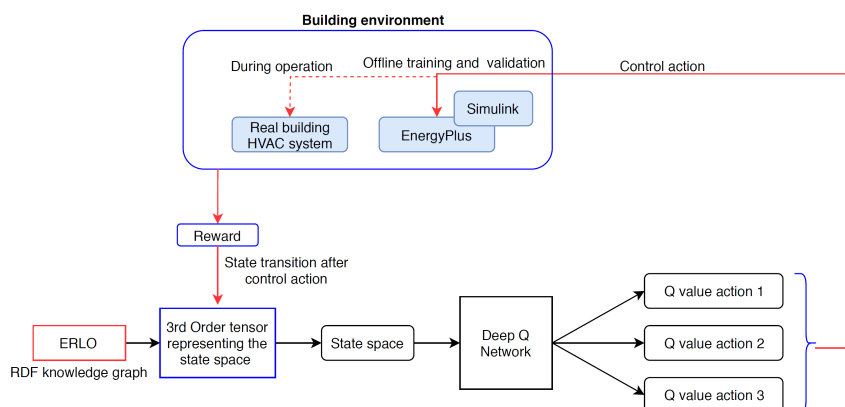
Having such a semantically rich learning state-space that preserves all the latent relationships between the heterogeneous building information, sensory data, HVAC operation and parameters affecting contaminant control within indoor building spaces, smarter and more energy efficient indoor monitors can be developed to work with HVAC systems to maintain acceptable indoor air quality within air tight building envelopes which are typical of Malaysia's indoor environments especially during haze episodes.

#### 4. Conclusions and future work

Many studies in neighbouring Singapore have already shown that the air filters currently installed in many indoor mechanical ventilation systems are not only insufficient in protecting building occupants from high exposures to fine particles of outdoor origin but are also energy intensive due to the pressure drops of air going through them which makes them draw more power during operation. Improving indoor air quality without compromising thermal comfort and optimal energy use requires engineers to identify and understand the existential relationships between the parameters affecting the dynamics of haze PM indoors.

This paper has presented the preliminary efforts from an ongoing research study to model such heterogeneous building information (envelope, sensors, contaminant properties, geometry, schedules and weather data) within semantic web knowledge graphs that more accurately represent the heterogeneous learning state-space ingestible by algorithmic building control systems to smartly monitor the IAQ and autonomously adapt HVAC units towards occupant comfort in an energy efficient manner. By testing such a knowledge graph within the learning architecture of deep Q-learning models for HVAC controllers, a deeper understanding will be added to the Malaysian research community in response to designing more resilient indoor HVAC systems that minimize infiltrated haze particulates within the building envelope.

For future progress of this work, the RDF data model ERLO will be extended to a simulation environment like EnergyPlus and Simulink where algorithm learning and control can be further tested and bench-marked as shown schematically in figure 8. The state of the semantically-enriched simulation (building) environment changes after each trial of a learned control action. These state updates are periodically and iteratively represented in the ERLO state-space model during learning.



**Figure 8:** Future work

## Acknowledgements

This work is supported by the Department of Civil Engineering, University of Nottingham, Malaysia and MES Asia Sdn Bhd.

## References

- [1] Tham K W, Parshetti G K, Balasubramanian R, Sekhar C and Cheong D K 2018 *Building and Environment* **128** 96–106 ISSN 03601323 URL <https://doi.org/10.1016/j.buildenv.2017.11.036>
- [2] Field R D, Van Der Werf G R and Shen S S 2009 *Nature Geoscience* **2** 185–188 ISSN 17520894
- [3] Ramakreshnan L, Aghamohammadi N, Fong C S, Bulgiba A, Zaki R A, Wong L P and Sulaiman N M 2018 *Environmental Science and Pollution Research* **25** 2096–2111 ISSN 16147499 URL <https://doi.org/10.1007/s11356-017-0860-y>
- [4] Sharma R and Balasubramanian R 2017 *Environmental Research Letters* **12** 114026 ISSN 17489326 URL <https://doi.org/10.1088/1748-9326/aa86dd>
- [5] Zhou J, Chen A, Cao Q, Yang B, Chang V W and Nazaroff W W 2015 *Building and Environment* **93** 14–23 ISSN 03601323 URL <http://dx.doi.org/10.1016/j.buildenv.2015.04.029>
- [6] Chen C and Zhao B 2011 *Atmospheric Environment* **45** 275–288 ISSN 13522310 URL <http://dx.doi.org/10.1016/j.atmosenv.2010.09.048>
- [7] Johnston F H, Purdie S, Jalaludin B, Martin K L, Henderson S B and Morgan G G 2014 Air pollution events from forest fires and emergency department attendances in Sydney, Australia 1996–2007: A case-crossover analysis
- [8] Reid C E, Brauer M, Johnston F H, Jerrett M, Balmes J R and Elliott C T 2016 Critical review of health impacts of wildfire smoke exposure
- [9] Liu C, Yang J, Ji S, Lu Y, Wu P and Chen C 2018 *Building and Environment* **144** 357–364 ISSN 03601323 URL <https://doi.org/10.1016/j.buildenv.2018.08.039>
- [10] Du D and Fei M 2008 *Applied Mathematics and Computation* **205** 26–36 ISSN 00963003
- [11] Lu S, Wang W, Lin C and Hameen E C 2019 *Building and Environment* **156** 137–146 ISSN 03601323 URL <http://doi.org/10.1016/j.buildenv.2019.03.010>
- [12] Barrett E and Linder S 2015 Autonomous HVAC control, a reinforcement learning approach *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* vol 9286 (Springer Verlag) pp 3–19 ISBN 9783319234601 ISSN 16113349
- [13] LeCun Y, Bengio Y and Hinton G 2015 *Nature* **521** 436–444 ISSN 0028-0836 URL <http://www.nature.com/articles/nature14539>
- [14] Yang L, Nagy Z, Goffin P and Schlueter A 2015 *Applied Energy* **156** 577–586 ISSN 03062619 URL <http://dx.doi.org/10.1016/j.apenergy.2015.07.050>
- [15] Mason K and Grijalva S 2019 *arXiv e-prints* 1–20 (Preprint arXiv:1806.00540v1)
- [16] W3C 2014 Linked Building Data Community Group URL <https://www.w3.org/community/lbd/>
- [17] Studer R, Grimm S and Abecker A 2007 *Semantic web services: Concepts, technologies, and applications* ISBN 9783540708933 URL <http://dx.doi.org/10.1007/3-540-70894-4>
- [18] W3C 2006 W3C RDF Validation Service URL <https://www.w3.org/RDF/Validator/>
- [19] Pauwels P and Terkaj W 2016 *Automation in Construction* **63** 100–133 URL <http://dx.doi.org/10.1016/j.autcon.2015.12.003>
- [20] Sutton R S and Barto A G 2018 *Reinforcement Learning: An Introduction* (MIT Press) ISBN 9780262039246
- [21] Mnih V, Kavukcuoglu K, Silver D, Rusu A A, Veness J, Bellemare M G, Graves A, Riedmiller M, Fidjeland A K, Ostrovski G, Petersen S, Beattie C, Sadik A, Antonoglou I, King H, Kumaran D, Wierstra D, Legg S and Hassabis D 2015 *Nature* **518** 529–533 ISSN 14764687 URL <http://dx.doi.org/10.1038/nature14236>
- [22] Bellman R 1957 *Journal of Mathematics and Mechanics* **6** 679–684 URL <https://www.jstor.org/stable/24900506>
- [23] Rasmussen M H, Pauwels P, Hviid C A and Karlshøj J 2017 Proposing a Central AEC Ontology That Allows for Domain Specific Extensions *Joint Conference on Computing in Construction* vol 1 pp 237–244 URL <http://doi.org/10.24928/jc3-2017/0153>
- [24] Rasmussen M H, Pauwels P, Lefrançois M and Schneider G F 2019 Building Topology Ontology (BOT)-Draft Community Group Report 07 January 2019 URL <https://w3c-lbd-cg.github.io/bot/>
- [25] Haller A, Janowicz K, Cox S, Le Phuoc D, Taylor K and Lefrançois M 2017 Semantic Sensor Network (SSN) Ontology-W3C Recommendation 19 October 2017 URL <https://www.w3.org/TR/vocab-ssn/>
- [26] Rasmussen M H, Lefrançois M, Bonduel M, Hviid C A and Karlshøj J 2018 OPM: An ontology for describing properties that evolve over time *CEUR Workshop Proceedings* vol 2159 pp 23–33 ISSN 16130073
- [27] Wagner A, Moeller L K, Leifgen C and Eller C 2019 Building Product Ontology URL <https://www.projekt-scope.de/ontologies/bpo/>