Expert Systems With Applications 158 (2020) 113461

Contents lists available at ScienceDirect

Expert Systems With Applications

journal homepage: www.elsevier.com/locate/eswa

A decision support system for urban infrastructure inter-asset management employing domain ontologies and qualitative uncertainty-based reasoning

Lijun Wei[a,*], Heshan Du[a,b], Quratul-ain Mahesar[a,1], Kareem Al Ammari[a], Derek R. Magee[a], Barry Clarke[c], Vania Dimitrova[a], David Gunn[d,2], David Entwisle[d,2], Helen Reeves[d,2], Anthony G. Cohn[a,e,*]

[a] School of Computing, University of Leeds, Leeds, United Kingdom
[b] School of Computer Science, University of Nottingham Ningbo China, Ningbo, China
[c] School of Civil Engineering, University of Leeds, Leeds, United Kingdom
[d] British Geological Survey, Keyworth, Nottingham, United Kingdom
[e] College of Electronic and Information Engineering, Tongji University, Shanghai, China

ARTICLE INFO

Article history:
Received 7 March 2019
Revised 10 March 2020
Accepted 15 April 2020
Available online 17 April 2020

Keywords:
Smart cities
Infrastructure maintenance
Underground utilities
Rule-based system
Reasoning under uncertainty

ABSTRACT

Urban infrastructure assets (e.g. roads, water pipes) perform critical functions to the health and well-being of society. Although it has been widely recognised that different infrastructure assets are highly interconnected, infrastructure management in practice such as planning, installation and maintenance are often undertaken by different stakeholders without considering these dependencies due to the lack of relevant data and cross-domain knowledge, which may cause unexpected cascading social, economic and environmental effects. In this paper, we present a knowledge based decision support system for urban infrastructure inter-asset management. By considering various infrastructure assets (e.g. road, ground, cable), triggers (e.g. pipe leaking) and potential consequences (e.g. traffic disruption) as a holistic system, we model each sub-domain using a modular ontology and encapsulate the interdependence between them using a set of rules. Moreover, qualitative likelihood is assigned to each rule by domain experts (e.g. civil engineers) to encode the uncertainty of knowledge, and an inference engine is applied to predict the potential consequences of a given trigger with location specific data and the encoded rules. A web-based prototype system has been developed based on the above concept and demonstrated to a wide range of stakeholders. The system can assist in the process of decision making by aiding data collation and integration, as well as presenting potential consequences of possible triggers, advising on whether additional information is needed or suggesting ways of obtaining such information. The work shows an intelligent approach to integrate and process multi-source data to pioneer a novel way to aid a complex decision process with a high social impact.

1. Introduction

Urban infrastructure assets, such as roads, ground and utilities (e.g. water, electricity, gas), are critical to the functioning of modern society (Clarke et al., 2017). Without efficient and effective diagnosis and maintenance, asset failures such as ground sinking caused by underground sewer collapse, can lead to significant economic, social, and environmental costs (Hojjati, Jefferson, Metje, & Rogers, 2016). These problems are particularly challenging in urban areas with increasing destructive street works due to extreme weather and ageing infrastructure. Research efforts have been devoted to developing various kinds of decision support systems (DSSs) for proactive urban infrastructure maintenance (Halfawy, 2010; Hojjati et al., 2016; Quintero, Konar, & Pierre, 2005; Rogers et al., 2012). For example, Arsene, Gabrys, and Al-Dabass (2012) proposed a decision support system for water pipe leakage detection, Moazami, Behbahani, and Muniandy (2011) proposed a supporting tool for pavement rehabilitation and maintenance pri-

* Corresponding authors.
E-mail addresses: l.j.wei@leeds.ac.uk (L. Wei), heshan.du@nottingham.edu.cn (H. Du), quratul-ain.mahesar@abdn.ac.uk (Q.-a. Mahesar), k.alammari@leeds.ac.uk (K. Al Ammari), d.r.magee@leeds.ac.uk (D.R. Magee), b.g.clarke@leeds.ac.uk (B. Clarke), v.g.dimitrova@leeds.ac.uk (V. Dimitrova), dgu@bgs.ac.uk (D. Gunn), dce@bgs.ac.uk (D. Entwisle), hjre@bgs.ac.uk (H. Reeves), a.g.cohn@leeds.ac.uk (A.G. Cohn).

[1] Quratul-ain Mahesar is now at the University of Aberdeen; her contribution to this paper was performed whilst employed at the University of Leeds.

[2] The contributions from Helen Reeves, David Entwisle and David Gunn are published with the permission of the Director of the British Geological Survey (UKRI).

https://doi.org/10.1016/j.eswa.2020.113461
0957-4174/

oritisation using fuzzy logic. But in general, these systems are confronted with four practical challenges for achieving proactive maintenance.

The first challenge is that urban infrastructure assets are interdependent at multiple levels (Ouyang, 2014; Rogers et al., 2012), but they are usually constructed and maintained by different stakeholders who plan and conduct street works independently without considering these interdependencies. Construction works or deterioration related to one asset may damage other assets nearby, causing cascading problems (Ouyang, 2014). For example, breaking up or opening a road may damage the underlying ground and buried utilities. Although it has been widely recognised that an integrated (Halfawy, 2010; Quintero et al., 2005) or a "system of systems" approach (Hall, Tran, Hickford, & Nicholls, 2016) is needed for infrastructure management, the lack of explicit knowledge of asset interdependencies makes it difficult for decision makers to have a holistic view of the potential impact of their actions. Moreover, the ground, which supports the road and the buried utilities and transmits actions (e.g. traffic load) between them, is rarely considered by practitioners as an asset (Clarke et al., 2017; Rogers et al., 2012). Successful implementation of an integrated approach largely depends on the ability to share comprehensive multi-sector knowledge, especially the broad knowledge of asset interdependencies.

The second challenge is that decision making in urban infrastructure management requires a variety of data (Quintero et al., 2005), such as underground utility maps, road construction details and road closure regulations. This data is often held by different data owners and stored in disconnected or even incompatible platforms, which makes it difficult for decision makers to gather useful data in a short period of time. The ability to integrate disconnected datasets into one single system would be helpful for decision makers (Michele & Daniela, 2011). Although semantic techniques have been proposed to integrate various buried asset data based on ontologies (Balasubramani et al., 2017; Halfawy, 2010; Quintero & Pierre, 2002), none of these work considered other contextual information in the urban infrastructure system, such as weather, road traffic, and ground conditions, which significantly limits their applicability in complex decision scenarios.

The third challenge is how to devise appropriate methods for proactive infrastructure maintenance, i.e. to predict the potential consequences of actions/observations in infrastructure management and suggest appropriate countermeasures. This requires identifying potential consequences / hazards on infrastructure assets (e.g. road collapse), society (e.g. traffic delays/disruptions, damage to property) and environment (e.g. ground contamination), as well as identifying the causes (e.g. possible behaviours) that may lead to these consequences and the internal mechanism. For example, model-based techniques, such as probabilistic models and neural networks were used for water pipe failure prediction (Arsene et al., 2012; Hadzilacos et al., 2000) and electrical utility maintenance (Bumblauskas, Gemmill, Igou, & Anzengruber, 2017); case-based reasoning techniques were used for selecting infrastructure intervention techniques (Quintero et al., 2005). However, all these techniques require a set of historical data or cases as training samples, which do not always exist in practice. Instead of learning from voluminous historical data, Marlow, Gould, and Lane (2015) used logical rules formulated by domain experts to suggest suitable pipe and road pavements rehabilitation techniques. The advantage of using rule-based approach lie in the fact that rules are based on experts' knowledge underpinned by observation, experiments and theory so they have limited dependence on historical data; the experience from one city's infrastructure can also be easily generalised to another city. But this approach has not been fully examined for diagnosis and predicting consequences in urban infrastructure management yet.

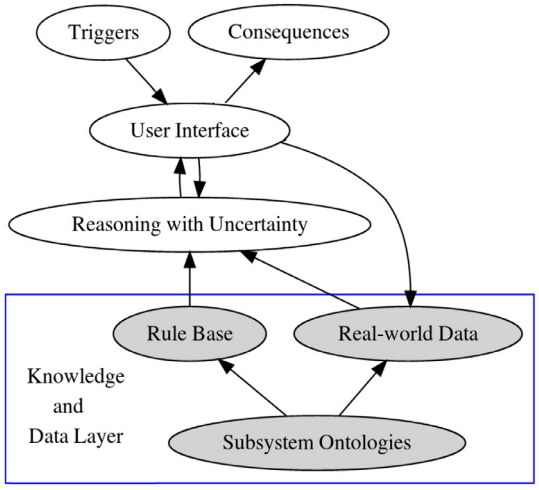

Fig. 1. Building blocks of the ATU-DSS: given a trigger, an inference engine is able to predict the potential consequences by reasoning with uncertain data and rules developed based on a set of modular ontologies.

The last challenge is that in the application of rule-based approach, rules formulated by domain experts are not always certain but require hedging with a confidence. To solve this, approaches such as *Certainty Factors (CFs)* (Shortliffe & Buchanan, 1975) have been proposed by attaching degrees of belief to propositions and rules. However, research warned that certainty factors could yield disastrously incorrect degrees of belief through over-counting of evidence in several circumstances (Heckerman, 1986), especially when the rule sets become larger. So almost all CFs based rule systems were either purely diagnostic (e.g. MYCIN) or predictive (Heckerman, 1986; Heckerman & Shortliffe, 1992). Additionally, the belief in rules is usually specified by domain experts using numerical values, whereas human judgemental reasoning is often more qualitative than numerical (Parsons & Parsons, 2001). Fuzzy rules have also been used to encode the uncertainty of knowledge (Chen, 1994; Malmir, Amini, & Chang, 2017; Moazami et al., 2011), but this method still requires numerical range values for deciding the membership functions and may be challenging for non-academic domain experts without logic background. The domain engineers we consulted in civil engineering also suggested that it is difficult to formulate rules with numeric values, especially in the case of dealing with the ground. In fact, in many cases, precise specification of numerical values may not be necessary for supporting decision making (Goldszmidt & Pearl, 1996; Wellman, 1990).

1.1. Our contribution

In order to meet the challenges described above, we present an intelligent web-based decision support system for urban infrastructure inter-asset management based on a system-of-systems approach (Fig. 1), especially focusing on the assets in direct contact with the underground world, including road, ground and underground utilities. The system is called Assessing the Underworld DSS, referred to as **ATU-DSS** hereafter.[3] The above four challenges were addressed as below: (1) firstly, to help address the challenge of multi-sector knowledge sharing, urban infrastructure is considered as a system of multiple subsystems and a family of interlinked modular ontologies were developed to capture the domain knowledge on each sector, including assets (e.g. pipe), related triggers (e.g. road cracks), potential consequences (e.g. loss of utility ser-

[3] UK EPSRC funded Assessing the Underworld (ATU) Project.

vice) and investigation techniques (e.g. ground penetrating radar surveys), etc. Then, based on the concepts defined in these domain ontologies, a set of logical rules were developed to encapsulate the interdependencies between different assets (e.g. *"IF RoadBaseWaterContent increases and PipeLeakingRate is Severe, THEN RoadBaseStiffness will definitely decrease"*) and the relations leading to serious asset failures or other hazard consequences. (2) secondly, to help address the challenge of disconnected data, various spatial datasets about the infrastructure assets and their contextual information were sourced from different owners and integrated in a single system to provide instant location specific data retrieval. (3) thirdly, to face the challenge of proactive maintenance and limited historic data, our system adopted a rule based reasoning approach, which could be enriched when more human decisions or real data are fed in. An inference engine is applied to infer the potential consequences of a reported trigger based on the knowledge in the rule base and the retrieved data from the integrated database. (4) lastly, a qualitative uncertainty based reasoning approach is proposed to handle the challenge of uncertainty in human knowledge; the system can also make assumption of the states of missing data to derive potential issues when not all data is available and suggest investigation methods to obtain the missing data. This function allows practitioners to plan for further surveys to reduce the potential risk.

Our contribution in this work is twofold. Theoretically, we proposed a framework (shown in Fig. 1) for developing a knowledge-driven decision support system using a system-of-systems approach, which can be easily generalised to various engineering applications. The framework starts by identification and modelling of subsystems using modular ontologies, and capturing their interdependencies using logical rules, followed by definition of triggers which may affect the subsystems and definition of consequences which may have serious impact on the subsystems or external environment (e.g. social consequences). Then, an inference engine is applied to predict the potential consequences of given triggers and advise on whether additional information is needed and suggest ways of obtaining such information. Practically, a prototype system has been developed based on the above concept by combining real-time site specific data retrieval with automated reasoning. It also allows users to modify data values for alternative analysis.

The prototype system can help decision makers (e.g. incident managers, contractors, local authorities) to gather relevant data in one-stop, store and re-use previously collected data, codify requirements from local authorities (e.g. restrictions), predict possible issues of observations or actions in advance, and learned wisdom (e.g. issues encountered) from previous experience to help data interpretation and decision making. These functions can improve safety in infrastructure management and reduce costs and prevent delays. For example, this system can help institutions pass on knowledge to junior engineers and help answering questions like (Clarke et al., 2017): (*a*) *How will the condition of road surface, adjacent pipes and ground at a specific site change* because of extreme weather, deterioration of assets, or human actions (e.g. planned excavation)? (b) *Will this change cause any undesirable consequences* (e.g. traffic disruption, loss of services or even fatalities)? and (c) *which asset should we maintain/replace* in the first place or *when and where* should we install a new underground asset? This evaluation of the undesirable consequences is different from traditional traffic or environmental impact assessment as it takes into account the knowledge of different factors and their interdependencies to achieve a more comprehensive assessment. The system can also suggest the likelihood and severity of potential consequences which can help decision makers to prioritise their maintenance tasks, preparation of health/safety files and reduce the potential risk of high likelihood and severity.

The rest of the paper is structured as follows: Section 2 briefly introduces the ATU domain ontologies; Section 3 explains the qualitative uncertainty based reasoning approach and the strategy for handling missing data in this system. We then provide a detailed description of the prototype system in Section 4, including how the rule base was developed, what exemplar data sources were integrated and different functions of the user interface; followed by discussions in Section 5 and conclusions in Section 6.

2. ATU Ontologies: A common vocabulary for data integration and reasoning

An ontology is a formal representation of the knowledge within a domain using a set of concepts and relationships between them (Staab & Studer, 2009). It can be used as a common vocabulary and thus plays an important role in information sharing (Gruber, 1993; Noy & McGuinness, 2001). More formally, an ontology consists of a TBox which defines terminologies or knowledge at the conceptual level, and an ABox which describes facts about individuals using terminologies defined in the TBox (Baader, Calvanese, McGuinness, Nardi, & Patel-Schneider, 2007; Du, Alechina, Jackson, & Hart, 2013). A TBox contains definitions of classes or concepts and definitions of roles or object properties, which cover the conceptual hierarchies and relations among classes. An ABox contains assertions about individuals or instances.

There exist several ontologies where concepts of infrastructure assets (such as the ground/soil, roads and buried pipes/cables) are defined. The concept *Soil* or *Ground* is defined in general environmental ontologies or thesauri, including the Semantic Web for Earth and Environment Technology (SWEET) ontology (Raskin & Pan, 2005), the environment ontology (ENVO) (Buttigieg, Morrison, Smith, Mungall, & Lewis, 2013), the General Multilingual Environmental Thesaurus (GEMET),[4] the AGROVOC Multilingual Thesaurus[5], etc., as well as ontologies specialised for describing soil (Das, 2010; Heeptaisong & Srivihok, 2010; dos Santos Aparício, de Farias, & dos Santos, 2006; Shivananda & Kumar, 2013; Zhao, Zhao, Tian, Qian, & Zhang, 2009). Buried assets, such as pipes and cables, are defined in several infrastructure management ontologies or utility ontologies (Fu & Cohn, 2008; Osman & El-Diraby, 2006; Zeb & Froese, 2014). The concept *Road* is defined in several transportation ontologies as part of the transportation network (Corsar, Markovic, Edwards, & Nelson, 2015; Katsumi & Fox, 2019; Lorenz, Ohlbach, & Yang, 2005), and defined in ontologies for urban development (Berdier, 2011; Berdier & Roussey, 2007). A survey of existing transportation ontologies has been provided recently by Katsumi and Fox (2018). The concept *Road* is also defined in general ontologies or vocabularies covering various domains, such as the environment ontology (ENVO) (Buttigieg et al., 2013), the DBpedia Ontology[6], the Linked Open Vocabularies (LOV)[7], etc. However, none of the ontologies above provides a systematic and comprehensive description of soil/pipe/road properties and processes, nor how soil/pipe/road properties and processes affect each other.

At the heart of the ATU-DSS is a suite of interlinked modular ontologies (Fig. 2), developed by following the NeOn methodology (Suárez-Figueroa, Gómez-Pérez, & Fernández-López, 2012), consulting various domain experts (e.g. geotechnical engineers, geophysicists) and extensively reviewing literature, existing ontologies (e.g. NASA's SWEET ontology (Raskin & Pan, 2005)), standards and datasets. The ATU ontologies are written in OWL 2 Web Ontology Language Manchester Syntax (Horridge & Patel-Schneider, 2012), which is based on description logic. Whilst we are aware of

[4] http://www.eionet.europa.eu/gemet. Accessed: 2020-02-26.
[5] http://aims.fao.org/standards/agrovoc. Accessed: 2020-02-26.
[6] https://wiki.dbpedia.org/services-resources/ontology. Accessed: 2020-02-26.
[7] https://lov.linkeddata.es/dataset/lov/. Accessed: 2020-02-26.

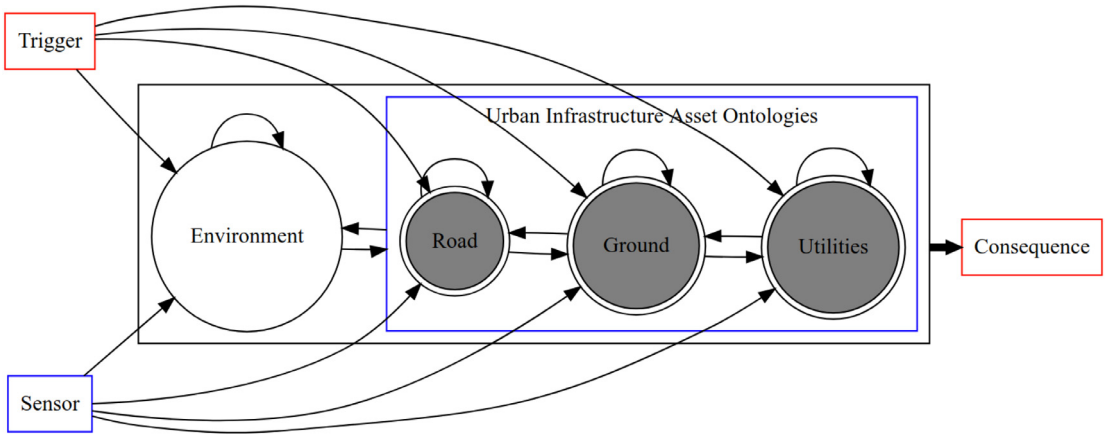

Fig. 2. Inter-dependencies between different domain ontologies.

Table 1
Number of processes and properties in ATU urban infrastructure asset ontologies.

Ontology	Number of Processes	Number of Properties
Ground Ontology	36	61
Road Ontology	17	85
Pipe Ontology	4	50

the recent work on fuzzy ontologies (Morente-Molinera, Kou, Pang, Cabrerizo, & Herrera-Viedma, 2019), which allows fuzzy relations additionally compared to regular ontologies, the ontologies we defined here are regular or crisp ontologies. To ensure the high quality of knowledge defined in them, the ATU ontologies are created semi-automatically, going through an iterative process involving several discussions and meetings with domain experts. As shown in Fig. 2, the ATU ontologies model the knowledge on infrastructure assets, triggers, environment and investigation techniques (e.g. sensors) for urban infrastructure management. Each ontology is introduced in the following sections.

The ***ATU Urban Infrastructure Asset Ontologies*** describe the main concepts and relations of underground related urban infrastructure assets using three ontologies, including the soil/ground, road (surface infrastructure) and pipe (buried utilities) (Du et al., 2017; 2016). Each asset ontology models the investigated asset using a set of properties (e.g. ground clay content), processes (e.g. ground biological process) and simple relations about how properties and processes affect each other (Du et al., 2016). The processes and properties are grouped into categories as the characteristics of the asset, such as *GroundBiologicalProcess* and *GroundChemicalProperty*. The number of processes and properties in the three asset ontologies are shown in Table 1. An example of the hierarchy of the ATU *Pipe Ontology* is shown in Fig. 3.

A key feature of these ontologies is that within an asset ontology, a change in a property would activate a process which leads to a change in other property(s). This cascading structure was achieved by encoding the dependency between classes in each sub-ontology with six relationships, including *hasImpactOn* and its inverse *influencedBy*, as well as *increases, decreases* and their inverse *increasedBy* and *decreasedBy* (Du et al., 2017; 2016). An example of the cascading relationships between different ontology concepts is "*GroundSwelling* decreases the *GroundStiffness* which *hasImpactOn* some *GroundS-WaveVelocity*". The complex relationships between multiple concepts are defined with rules and will be explained in Section 3. The urban infrastructure asset ontologies are publicly available at https://doi.org/10.5518/190.[8]

The ***ATU Trigger Ontology*** defines the categories and properties of events that may cause cascading effect on infrastructure assets (Clarke et al., 2017). They are often human actions such as planned construction works, or observable phenomena like natural phenomena (e.g. excessive rainfall, extreme temperatures), abnormal observations on road (e.g road cracks), on ground (e.g. ground movement) or on buried utilities (e.g. drop in water pressure). For example, rainfall is an external trigger which infiltrates the ground and leading to an increase in *GroundWaterContent*. 32 types of triggers were included in the trigger ontology as shown in Fig. 4, each of which relates to a scenario that we want to tackle in urban infrastructure management. The properties of triggers include general properties like severity, location/spatial geometry (e.g. point, linestring) and time (a time point or a period), and specific properties like the type of a construction work. In the prototype DSS system, a decision process starts with a report of a trigger by users through the user interface.

The ***ATU Consequence Ontology*** identifies various consequences in infrastructure management, such as direct consequences on stakeholders (e.g. cost overrun) or on infrastructure assets (e.g. pipe burst, trench collapse), and indirect legal, social, economical (e.g. traffic disruption) or environmental consequences (e.g. water pollution) on the general public. Each consequence is attached with a context-dependent severity level (e.g. Negligible, Marginal, Critical, Catastrophic).

ATU Investigation Ontology (Sensors). Being different from the well known *Semantic Sensor Network (SSN) ontology*,[9] the *ATU Investigation Ontology (Sensors)* encodes the knowledge of currently available techniques for obtaining different infrastructure asset properties. These techniques can include a pointer to an external institution (e.g. website), a laboratory test, or different types of sensor surveys. By working with domain experts in geotechnical engineering and geophysics and reviewing literature (sen, 1988), the current ontology includes seven classes (e.g. seismic methods, electrical methods)and 26 types of geophysical techniques (e.g. Ground Penetrating Radar), together with their relationships with the ATU Urban Infrastructure Asset Ontologies. Two relationships were defined to describe the suitability of an investigation method for measuring different asset properties in shallow (0 – 5 m depth) surveys. The relationships are *measures* and its inverse *measuredBy*, which means "SensorA *Measures* PropertyB with usefulness_score N", and "PropertyB is *MeasuredBy* SensorA with usefulness_score N". The investigation suggestions and the corresponding usefulness score *N* (i.e. an integer between 0 and 4, where 0 means "not considered applicable" and 4 means "gen-

[8] Accessed: 2020-02-23.

[9] SSN ontology: https://www.w3.org/TR/vocab-ssn/. Accessed: 2020-02-23.

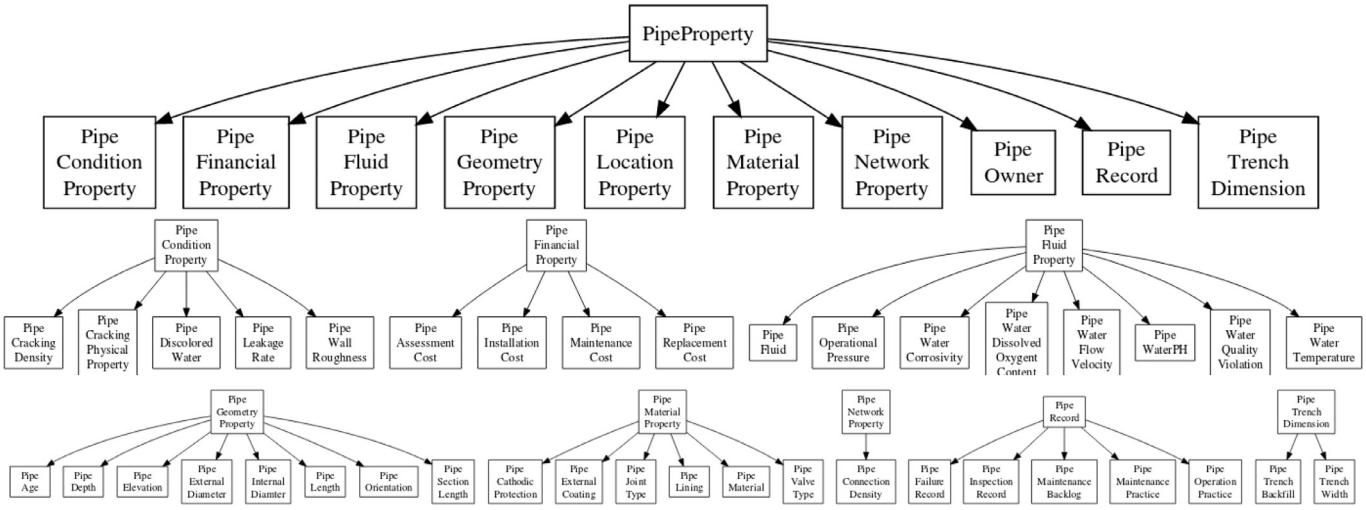

Fig. 3. An example of the ontology concept hierarchy of pipe properties.

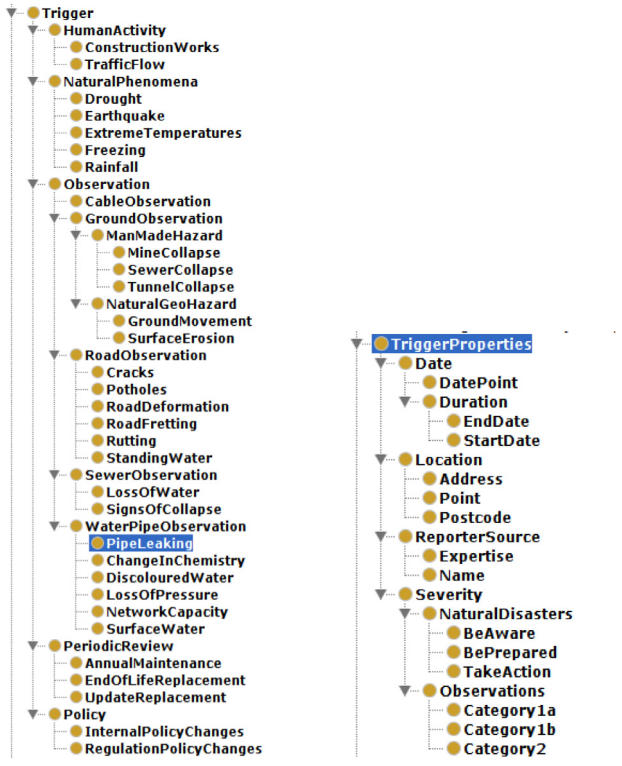

(a) Trigger types. (b) Trigger properties.

Fig. 4. Triggers and their properties included in the Trigger Ontology.

erally considered an excellent and well developed approach") of each relation were assigned by domain experts and implemented in OWL Protégé as annotations (which can be queried). With these relationships established, appropriate investigation techniques can be recommended to users in a prototype decision support system when the data of an asset property is missing. Other investigation methods can also be added into the knowledge base in the future.

The ***ATU Environment Ontology*** models the environment factors (e.g. rainfall, drought) affecting or being affected by the infrastructure assets. Instead of building this ontology from scratch like the modular ontologies presented above, the *Environment Ontology* was created based on several existing external ontologies (e.g. NASA's SWEET Ontology,[10] the Environment Ontology,[11] Ordnance Survey's Buildings and Places Ontology[12]). This is because our work does not need a thorough modelling of the environment, but only need several essential concepts like rainfall, drought, contamination.

The concepts in these ATU ontologies are used to guide relevant data sourcing, and as a common vocabulary for defining inference rules (complex relationships between multiple ontology concepts) and integrating various datasets from different domains such that heterogeneous data can be used seamlessly in automated reasoning. Though ontologies have been used widely for knowledge modelling and information retrieval (Munir & Anjum, 2018), most of the existing approaches use a single domain ontology rather than a series of ontologies covering various domains as we do here. For more details of the ATU ontologies, interested readers are referred to one of our earlier works (Du et al., 2016) on soil ontology. The authors are also preparing a separate paper to introduce the modelling process of ATU ontologies.

3. Qualitative uncertainty based reasoning in rule-based systems

Based on the concepts defined in the ATU ontologies, we continue to develop logical rules in collaboration with domain experts to encapsulate the broad knowledge of internal dependencies in one subsystem, as well as the external dependencies between different infrastructure assets, environment factors and human activities. For example, a rule *"Heavy and Long rainfall will infiltrate the road if the road crack penetrates the road surface."* is defined referring to the concepts in the *Environment Ontology* and the *Road Ontology*, written as: "*Environment Rainfall Intensity (Heavy)* $\wedge$ *EnvironmentRainfallDuration (Long)* $\wedge$ *RoadCrackingDepth (High)* $\xrightarrow{Definite}$ *RoadWaterInfiltration (Active)*".

As discussed previously, rules are often conditional and require augmenting with confidence, but the domain engineers we consulted in urban infrastructure management found it difficult to formulate rules with numeric probabilities. Instead they preferred expressing uncertain information using qualitative linguistic expressions, which is in accordance with the theory proposed by Wallsten and Budescu (1995) on human reasoning. In this work,

[10] https://bioportal.bioontology.org/ontologies/SWEET. Accessed: 2020-02-23.
[11] http://environmentontology.org/. Accessed: 2020-02-23.
[12] http://bit.ly/2Fq5u0F. Accessed: 2020-02-23.

Table 2
Definitions of confidence levels for an event E happening.

Confidence level	Definition
Definite (D)	100% sure that E will happen
Very likely (V)	not 100% but only rarely will E not happen
Likely (L)	E happens more than 50% of the time but it is not surprising if E does not happen
Unlikely (U)	E happens less than 50% of the time but it is not surprising if E happens
Very unlikely (VU)	not 0% chance of E happening but it will happen rarely
Impossible (I)	E never happens

we present a qualitative confidence levels based uncertainty management scheme, which is considered to be more accessible by our domain experts compared with the linguistic quantifiers used with fuzzy sets theory (Bonissone, Gans, & Decker, 1987; Cid-Lpez, Hornos, Carrasco, Herrera-Viedma, & Chiclana, 2017; Zadeh, 1984) or kappa calculus (Clinton & David, 2004; Goldman & Maraist, 2015; Goldszmidt & Pearl, 1996; Poole & Smyth, 2005). The proposed confidence levels can also be interpreted in terms of the degrees of surprise using kappa calculus or using probability. Different calculation formulae are proposed to sequentially propagate the qualitative confidences of data and rules, as well as combine parallel chains leading to the same conclusion. It is an extension of our previous work (Mahesar et al., 2017) by adding probabilistic interpretations of the approach, adding a mechanism to avoid multi-counting of the same fact in different rules, checking potential contradictions in the rule base and considering diagnostic rules for abductive inference.

3.1. Definition of qualitative confidence levels and confidence vectors

3.1.1. Qualitative confidence levels

In the scope of this paper, six qualitative confidence levels {*Impossible, Very Unlikely, Unlikely, Likely, Very Likely, Definite*} were empirically selected to describe experts' confidence in a rule, i.e. the degree of people's (e.g. domain experts) belief of the conclusion is true given the premise. The confidence levels can be any other ordered linguistic lists (i.e. words/phrases) if the list follows an order from impossible to definite. For example, a list {*impossible, very improbable, improbable, probable, very probable, definite*} could be used to replace the one used in this work. Furthermore, the confidence levels granularity depends on the requirements of different applications. For example, the designer of a decision support system could simplify the confidence levels of rule sets from six to four as {*Impossible, Unlikely, Likely, Definite*} if they want to ease the workload of domain experts.

Definitions of the six confidence levels used in our work are listed in Table 2. These confidence levels constitute three pairs of symmetrical confidence levels, including *{Impossible vs Definite}, {Very Unlikely vs Very Likely}*, and *{Unlikely vs Likely}*. With this feature, domain experts are free to use any of the two states of a binary variable (e.g. On/Off, Active/Inactive) for authoring rules, since the complementary confidence level can be automatically inferred to guarantee other related rules could be fired. For example, let's assume a binary variable B with two states *On* and *Off*, and two rules:

$$\begin{aligned} &R1: \text{“}\textit{If A, Then it is } \textbf{\textit{Definite}} \textit{ that B is On}\text{”} \\ &R2: \text{“}\textit{If B is Off, Then it is } \textbf{\textit{Likely}} \textit{ that C will increase}\text{”} \end{aligned} \tag{1}$$

Originally, if we observe A, the first rule will be fired but the second not. However, based on our definition above, we can generate a new rule based on the first rule as: "*If A, Then it is* ***Impossible*** *that B is Off*". With this rule added and fired, the second rule will also be fired to further infer the state of C. Such complementary rules are added to the system automatically in the rule definition phase. We can also attach qualitative confidences to input facts to reflect their imprecision due to the inaccuracies in observation or measurement limitations (e.g. instrument precision).

Confidence levels can also be interpreted in terms of the degrees of surprise using kappa calculus (Goldszmidt & Pearl, 1996), an order of magnitude calculus in which each component is an order of magnitude more surprising than the next; a probability distribution P can be mapped to a kappa ranking κ such that P/ϵ^{κ} is finite but not infinitesimal for an infinitesimal ϵ. Kappa rankings can be interpreted as an approximation to probabilities through the following relations (Darwiche & Goldszmidt, 1994):

$$\kappa(\alpha \wedge \beta) = \kappa(\alpha|\beta) + \kappa(\beta) \tag{2}$$

$$\kappa(\alpha \vee \beta) = min(\kappa(\alpha), \kappa(\beta)) \tag{3}$$

where multiplication in probabilities ($P(\alpha \wedge \beta) = P(\alpha|\beta)P(\beta)$) is replaced by addition of kappa values (Eq. (2)), and addition of probabilities ($P(\alpha \vee \beta) = P(\alpha) + P(\beta) - P(\alpha \wedge \beta)$) is replaced by minimisation in the kappa calculus (Eq. (3)) (Clinton & David, 2004; Goldszmidt & Pearl, 1996; Poole & Smyth, 2005).

3.1.2. Qualitative confidence vectors of facts and rules

For computation purpose, instead of directly using the linguistic confidence levels or a single confidence value (i.e., a scalar), we use an ordered vector of four numerical elements (abbreviation of {*Very Unlikely, Unlikely, Likely, Very Likely*}) to encode the confidence level of a rule or a fact, called confidence vectors as:

$$C = \langle VU,\ U,\ L,\ V \rangle \tag{4}$$

The four elements of the confidence vector C_F of a fact can have any non-negative integer elements, while the confidence vector C_R of a predictive rule can have at most one positive element corresponding to the confidence levels defined in Table 2. For example, a rule *"If A happens, then it is* ***likely*** *that B will happen"* is attached with a confidence level *Likely* and the corresponding confidence vector is $C_R = \langle 0, 0, 1, 0 \rangle$. It should be noted that confidence level *D (Definite)* is implicit in the definition when *VU, U, L, V* are all 0, and confidence level *I (Impossible)* is implicit in the definition when any of the vector elements is with an extreme large number (for example, $VU = 1000$ is assigned to rules marked as *Impossible* in our application).[13] This vector representation is designed to store the accumulated uncertainties from data and rules using simple bit-addition. We can also easily restore the approximate probability from a confidence vector at the end of an inference process. However, if the uncertainties are represented using one single scalar, addition cannot be used any more as it will be difficult to separate and reconstruct the uncertainty from one single accumulated scalar. More explanations will be given in the next section.

Interpretation of confidence vectors. Given the confidence vector $C = \langle VU, U, L, V \rangle$ of a fact E (or a rule R), assuming that each element of this vector can be represented by a numerical value a_i between 0 and 1, the probability p (a measure of an expert's belief) of this fact (rule) to happen can be calculated as:

$$p_E = a_1^{VU} \cdot a_2^{U} \cdot a_3^{L} \cdot a_4^{V} \tag{5}$$

where $\{a_1, a_2, a_3, a_4\} \in [0, 1]$, $a_1 < a_2 < a_3 < a_4$ and $\{VU, U, L, V\} \in Z^+$, $p_E \in [0, 1]$. It can be noted that p_E approximates 0 when any of the vector elements gets extremely large (e.g. a conclusion is inferred from a long sequence of uncertain rules), especially for *VU*

[13] More details about *Impossible* will be given in the following sections. For example, assume $C = \langle 1000, 0, 0, 0 \rangle$ and the probability related to *VU* is $a_1 = 0.01$, the corresponding probability of this confidence vector is $P = a_1^{1000} = 0.01^{1000} \approx 0$; therefore, we consider this fact/rule as impossible.

since a_1 is close to zero; while p_E will be close to 1 when the sum of $\{VU, U, L, V\}$ is getting close to 0. The extreme case is when *VU, U, L, V* are all 0, p_E is equal to 1, which is consistent with the definition of "Definite" in Eq. (4). The numerical values of $\{a_1, a_2, a_3, a_4\}$ can vary for different applications or for different experts (though should be consistent in one knowledge base). The confidence vector provides an easy and intuitive way to encode and propagate the uncertainties in a rule-based system.

In terms of kappa-calculus, by taking logs of the probability defined in Eq. (5) (with a positive infinitesimal base ϵ) and replacing multiplication by addition, we get the surprise level of E (or R) happening as:

$$\begin{aligned}\kappa(E) &= (VU * \log_\epsilon a_1 + U * \log_\epsilon a_2 + L * \log_\epsilon a_3 + V * \log_\epsilon a_4)\\ &= \langle VU, U, L, V\rangle * \langle \log_\epsilon a_1, \log_\epsilon a_2, \log_\epsilon a_3, \log_\epsilon a_4\rangle^T\\ &= \langle VU, U, L, V\rangle * K \end{aligned} \tag{6}$$

The surprise that E happens comes from the surprise of each element of $\langle VU, U, L, V\rangle$, and can be re-written as the multiplication of the confidence vector $\langle VU, U, L, V\rangle$ and a constant vector $K = \langle \log_\epsilon a_1, \log_\epsilon a_2, \log_\epsilon a_3, \log_\epsilon a_4\rangle^T$ which represents the surprise of different confidence levels. As ϵ is infinitesimal, $\{a_1, a_2, a_3, a_4\} \in [0, 1]$ and $\langle VU, U, L, V\rangle$ are non-negative numbers, the kappa values are always non-negative. When all elements of $\langle VU, U, L, V\rangle$ are zero, the surprise of E happening is zero which is consistent with the definition of *Definite* in Table 2. In the following sections, we explain how to propagate the confidence vectors of data and rules in different situations.

3.2. Formula 1: The confidence level of a rule's conclusion given one fact

For a rule $R1$ with one premise,

$R1:$ *If* a, *then* b;

if our confidence in a is $C_a = \langle VU_a, U_a, L_a, V_a\rangle$, and the experts' confidence in this rule is $C_{R1} = \langle VU_{R1}, U_{R1}, L_{R1}, V_{R1}\rangle$, then our confidence in the conclusion b, denoted as $C_{a,b}$ (based on a), can be computed by adding the confidence vectors of the premise and the rule, written as:

$$\begin{aligned}C_{a,b} &= C_a + C_{R1}\\ &= \langle VU_a + VU_{R1}, U_a + U_{R1}, L_a + L_{R1}, V_a + V_{R1}\rangle\end{aligned} \tag{7}$$

For example, if we have $C_a = \langle 0, 0, 0, 0\rangle$ and $C_{R1} = \langle 0, 0, 0, 1\rangle$ (*Very Likely*); then, we can estimate the confidence in b as $C_{a,b} = \langle VU_a + VU_{R1}, U_a + U_{R1}, L_a + L_{R1}, V_a + V_{R1}\rangle = \langle 0, 0, 0, 0\rangle + \langle 0, 0, 0, 1\rangle = \langle 0, 0, 0, 1\rangle$. How to combine the uncertainty of a fact inferred from parallel rules will be discussed in Section 3.4.

Interpretation of Formula 1. As rule $R1$ defines the belief in b when a happens ($P(b|a)$), we can obtain the joint probability of b and a through rule $R1$ as:

$$P(a, b) = P(a)P(b|a) \tag{8}$$

Note we can not derive $P(b)$ since $P(b|\neg a)$ is not defined in this rule. In terms of kappa-calculus, the surprise of $(a \wedge b)$ happening can be obtained based on the addition formula in Eq. (2), and linked to confidence vectors based on the definition of kappa values in Eq. (6):

$$\begin{aligned}\kappa(a \wedge b) &= \kappa(a) + \kappa(b|a) = C_a * K + C_{b|a} * K\\ &= \langle VU_a + VU_{R1}, U_a + U_{R1}, L_a + L_{R1}, V_a + V_{R1}\rangle * K\\ \Rightarrow\ C_{a\wedge b} &= \langle VU_a + VU_{R1}, U_a + U_{R1}, L_a + L_{R1}, V_a + V_{R1}\rangle\end{aligned} \tag{9}$$

The derived confidence vector is the same as the formula in Eq. (7), thus it can be considered as a simple way to accumulate the kappa values (the first part of $\kappa(E)$) of different facts and rules. We can decide the surprise level of a derived fact by multiplying K with the derived confidence vector and tuning the values of a_1, a_2, a_3, and a_4.

3.3. Formula 2: The confidence level of a rule's conclusion given a conjunction of facts

For a rule $R2$ with n uncertain premises A_i and confidence vector $C_{R2} = \langle VU_{R2}, U_{R2}, L_{R2}, V_{R2}\rangle$:

$R2:$ *If* $A_1 \wedge A_2 \cdots \wedge A_n$, *then* B;

If our confidence in each fact A_i is $C_{A_i} = \langle VU_i, U_i, L_i, V_i\rangle$ and assuming all the facts are independent, our confidence in the conjunction of $(A_1 \wedge \cdots \wedge A_n)$ can be calculated by taking the sum of the confidence vectors of all these facts, written as:

$$C_{(A_1\wedge\cdots\wedge A_n)} = \sum_i C_{A_i} = \langle \sum_i VU_i, \sum_i U_i, \sum_i L_i, \sum_i V_i\rangle \tag{10}$$

After obtaining the confidence $C_{(A_1\wedge\cdots\wedge A_n)}$ in the conjunction of premises, our confidence in the conclusion B can be derived based on Formula 1 (Eq. (7)) by adding the confidence vectors of the rule and the premises:

$$C_{(A_1\wedge\cdots\wedge A_n),B} = C_{(A_1\wedge\cdots\wedge A_n)} + C_{R2} \tag{11}$$

For example, if our confidence in a rule *"If c and d, then h"* is $C_R = \langle 0, 1, 0, 0\rangle$ (*Unlikely*) and the two facts $C_c = \langle 0, 0, 1, 0\rangle$, $C_d = \langle 0, 0, 0, 1\rangle$, we have:

$$\begin{aligned}C_{(c\wedge d),h} &= (\langle 0, 0, 1, 0\rangle + \langle 0, 0, 0, 1\rangle) + \langle 0, 1, 0, 0\rangle\\ &= \langle 0, 0, 1, 1\rangle + \langle 0, 1, 0, 0\rangle = \langle 0, 1, 1, 1\rangle\end{aligned}$$

In the above formula, we assume all premises are independent. In terms of risk we may want to know if there is any chance the premises are not independent as this may suggest a higher confidence of the inferred fact[14] The following solution of multi-count problem partially handles the dependencies of premises (when one premise is used to infer other premises).

Avoidance of "Multi-count" Problem of Uncertainties. In rule based systems, whenever we have a rule of the form $A \Rightarrow B$, we can conclude B given A without worrying about other rules; and once B is proved, it can be used regardless of how it was derived. However, in dealing with probabilities, the source of the premises of a conclusion is important for subsequent reasoning (Russell & Norvig, 2010). Ignoring this may result in "multi-count" of the uncertainties of several facts during inference (Heckerman, 1986; Heckerman & Shortliffe, 1992). For example, assume there are two input facts *{a, b}* and two rules in the knowledge base:

$$\begin{aligned}&\textit{Rule } 1: \textit{ if } \mathbf{a} \textit{ and } \mathbf{b}, \textit{ then } \mathbf{c};\\ &\textit{Rule } 2: \textit{ if } \mathbf{c} \textit{ and } \mathbf{a}, \textit{ then } \mathbf{d}.\end{aligned} \tag{12}$$

Based on **Formula 2** (Eqs. (10) and (11)), we can derive the uncertainty of c and d and record all the facts used to deduce each fact (Mahesar, Dimitrova, Magee, & Cohn, 2017):

$$\begin{aligned}&\textit{Rule } 1 :\Rightarrow C_c = \langle C_a + C_b + C_{R1}\rangle (GF:\ a, b)\\ &\textit{Rule } 2 :\Rightarrow C_d = \langle C_c + C_a + C_{R2}\rangle (GF:\ a, b; a)\end{aligned} \tag{13}$$

where ***(GF)*** records the antecedents of an inferred fact. It can be seen that for inferring d, the uncertainty of a has been accumulated twice. To avoid multi-counting the uncertainty of a fact, we can first propagate the confidence vectors of rules to conclusions (inferred facts) and store the corresponding antecedents of all premises (i.e. unique values) in an inference process; then, the confidence of each inferred fact can be calculated by adding the confidence vectors of rules and unique antecedents when inference finishes.

[14] Proof of the lower and upper bound confidences under independent and non-independent assumptions are detailed in Section Appendix C.

Table 3
Propagation of confidence vectors to avoid multi-count problem in an example (Eq. (12)).

Fact	C_F	C_R	GF	Rule
a	C_a	[]	[a]	/
b	C_b	[]	[b]	/
c	[]	C_{R1}	[a,b]	R1: $a \wedge b \xrightarrow{C_{R1}} c$
d	[]	$C_{R1}+C_{R2}$	$[a,\ b]\vee a$	R2: $c \wedge a \xrightarrow{C_{R2}} d$

To implement this, a given fact d can be initialised as "$gf(C_F = C_d)(C_R = [])(GF = [d])$", where C_F stores the confidence vector of this fact, C_R stores the accumulated confidence vector of rules used to infer this fact, GF stores the antecedents of this fact. For an inferred fact, its C_R is the sum of the C_R of all its premises (based on Formula 1) and its GF is the union of the given facts GF of all its premises. For the example in Eq. (12), the reasoning process is shown in Table 3 from top to bottom.

Then, we can infer the confidence vector of c as $C_c = C_R + C_{GF} = C_{R1} + (C_a + C_b)$, and $C_d = C_R + C_{GF} = (C_{R1} + C_{R2}) + C_{[a,b]\vee a} = (C_{R1} + C_{R2}) + (C_a + C_b)$.

In summary, the "over-counting" problem can be avoided by checking the antecedents of all the premises of a rule and this solution deals in part with the dependencies between premises.

Interpretation of Formula 2. For the example in Eq. (12), since the confidence in Rule 2 suggests the conditional probability $P(d|c, a)$, we can obtain the joint probability $P(c, a; d)$ as

$$P(c, a; d) = P(c, a)P(d|c, a) = P(c)P(a|c)P(d|c, a) \tag{14}$$

In our system, we consider a premise a to be dependent on a premise c if a was used for deriving c (i.e a is the antecedent of c) and $P(a|c) = 1$; otherwise, we assume that different premises are independent to each other (e.g. $P(a|c) = P(a)$). Then, the joint probability can be written as:

$$P(c, a; d) = \begin{cases} P(c) * 1 * P(d|c, a), & if\ a\ is\ an\ antecedent\ of\ c \\ P(c)P(a)P(d|c, a), & if\ a\ is\ not\ an\ antecedent\ of\ c \end{cases} \tag{15}$$

In terms of kappa-calculus, the surprise of $(c \wedge a \wedge d)$ happening can be obtained using the kappa values (Eq. (6)) and the addition Eq. (2) as:

$$\begin{aligned} &\kappa(c, a; d) = \kappa(c) + \kappa(a|c) + \kappa(d|c, a) \\ &= \begin{cases} \kappa(c) + 0 + \kappa(d|c, a), & if\ a\ is\ an\ antecedent\ of\ c \\ \kappa(c) + \kappa(a) + \kappa(d|c, a), & if\ a\ is\ not\ an\ antecedent\ of\ c \end{cases} \end{aligned} \tag{16}$$

$k(a|c) = 0$ means there is no surprise that a happens if we observe c happens since a is one of the antecedents of c. Let Q denote the set of independent premises of rule $R2$, then the kappa value of the conclusion (and the premises) happening can be written as:

$$\begin{aligned} \kappa(c, a; d) &= \left[\sum_{i\in Q} \kappa(i)\right] + \kappa_{d|c,a} = [(\sum_{i\in Q} C_i) + C_{d|c,a}] * K \\ &= \left[\langle \sum_{i\in Q} VU_i, \sum_{i\in Q} U_i, \sum_{i\in Q} L_i, \sum_{i\in Q} V_i\rangle + C_{d|c,a}\right] * K \end{aligned} \tag{17}$$

So the corresponding confidence vector of the conclusion can be derived as:

$$C_{c,a;d} = \left\langle \sum_{i\in Q} VU_i, \sum_{i\in Q} U_i, \sum_{i\in Q} L_i, \sum_{i\in Q} V_i \right\rangle + C_{R2} \tag{18}$$

This formula is in accordance with Formula 2; moreover the multi-count problem has also been considered by excluding the dependent premises.

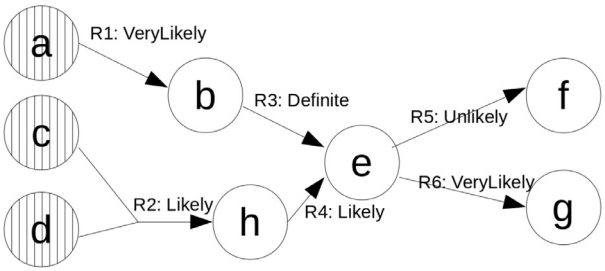

Fig. 5. A direct graph representing six predictive rules in a knowledge base (shaded circles for given facts and white circles for derived facts.).

3.4. Formula 3: Combining confidence levels of the same conclusion derived from parallel rules

Let $C_F = \langle VU_F, U_F, L_F, V_F\rangle$ denote our confidence in an inferred fact F. If the same fact F is inferred from two separate rules and the confidences are $C_F^1 = \langle VU_1, U_1, L_1, V_1\rangle$ and $C_F^2 = \langle VU_2, U_2, L_2, V_2\rangle$, then C_F can be computed by:

$$C_F = min(\langle VU_1, U_1, L_1, V_1\rangle, \langle VU_2, U_2, L_2, V_2\rangle) \tag{19}$$

where $min(\langle VU_1,\ U_1,\ L_1,\ V_1\rangle,\ \langle VU_2,\ U_2,\ L_2,\ V_2\rangle)$ returns the minimum of the two arguments as:

$$\begin{aligned} &\langle VU_1, U_1, L_1, V_1\rangle < \langle VU_2, U_2, L_2, V_2\rangle \ \Leftrightarrow \\ &(VU_1 < VU_2) \vee [(VU_1 = VU_2) \wedge (U_1 < U_2)] \\ &\vee [(VU_1 = VU_2) \wedge (U_1 = U_2) \wedge (L_1 < L_2)] \\ &\vee [(VU_1 = VU_2) \wedge (U_1 = U_2) \wedge (L_1 = L_2) \wedge (V_1 < V_2)] \end{aligned} \tag{20}$$

For example, in Fig. 5 (a knowledge base with six predictive rules), nodes in the figure represent uncertain variables, and arrows represent the rules from premises to conclusions.

Fact **a, c** and **d** are observed with confidence vectors C_a, C_c and C_d. It can be seen from the figure that fact e can be inferred from two different rules *Rule3 (Definite)* and *Rule4 (Likely)*. Based on the propagation **Formula 1** defined in Eq. (7), the confidence vectors of e inferred separately from the two rules are:

$$Rule\ 3 : if\ b,\ then\ e;\ C_{e,b}^1 = C_b + C_{R3}$$

$$Rule\ 4 : if\ h,\ then\ e;\ C_{e,h}^2 = C_h + C_{R4}$$

If our confidence in b is $C_b = \langle 0, 0, 0, 1\rangle$ and in h is $C_h = \langle 0, 0, 2, 1\rangle$, based on rules $C_{R3} = \langle 0, 0, 0, 0\rangle$ (*Definite*) and $C_{R4} = \langle 0, 0, 1, 0\rangle$ (*Likely*), we have:

$$C_e^1 = C_b + C_{R3} = \langle 0, 0, 0, 1\rangle + \langle 0, 0, 0, 0\rangle = \langle 0, 0, 0, 1\rangle$$

$$C_e^2 = C_h + C_{R4} = \langle 0, 0, 2, 1\rangle + \langle 0, 0, 1, 0\rangle = \langle 0, 0, 3, 1\rangle$$

Based on Formula 3 in Eq. (20), the confidence vector of e is: $min(C_e^1, C_2^2) = \min(\langle 0, 0, 0, 1\rangle, \langle 0, 0, 3, 1\rangle) = \langle 0, 0, 0, 1\rangle$

Interpretation of Formula 3. For the above example, let the conjunction of b and e as one event A with $P(A) = P(b \wedge e) = P(e|b)P(b)$, the conjunction of h and e as one event B with $P(B) = P(h \wedge e) = P(e|h)P(h)$, and b and h are mutually exclusive (as shown in Fig. 5). Given two mutually exclusive events A and B and assume $P(A) \geq P(B)$, we have:

$$\begin{aligned} P(e) &= P(A \vee B) = P(A) + P(B) - P(A)P(B) \\ &= P(A)\left[1 + P(B)\left(\frac{1}{P(A)} - 1\right)\right] \geq max(P(A), P(B)) \end{aligned} \tag{21}$$

As b (Rule 3) and h (Rule 4) provide two independent reasons to believe e, the two observations together should infer e with a belief that is stronger than either component in isolation. So $P(e)$ is always larger than the maximum probability of individual event.

Table 4
Contradictory confidence levels.

	U	VU	IM
D	X	X	X
VL		X	X
L			X

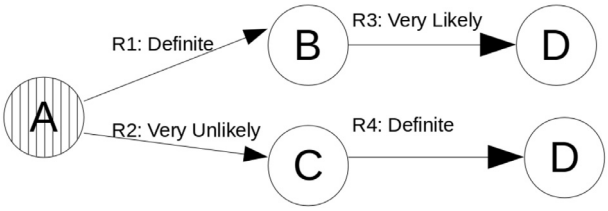

Fig. 6. A direct graph representing the predictive rules in a knowledge base.

In terms of kappa-calculus, the surprise of e happening is based on the smaller surprise (kappa value) of $A(e \wedge b)$ and $B(e \wedge h)$ according to Eq. (3) as:

$$\kappa(e) = min(\kappa(e,b), \kappa(e,h)) \tag{22}$$

Replacing the kappa values with confidence vectors (Eq. (6)), we have

$$\begin{aligned} &\kappa(e) = min(C_{e,b} * K, C_{e,h} * K) = min(C_{e,b}, C_{e,h}) * K \\ \Rightarrow\ &C_e = min(C_{e,b}, C_{e,h}) \\ &\quad = min(\langle VU_A, U_A, L_A, V_A\rangle, \langle VU_B, U_B, L_B, V_B\rangle) \end{aligned} \tag{23}$$

The minimum confidence vector is related to the maximum probability of A and B and the derived confidence vector is in accordance with the formula in Eq. (19). Derivation of the minimum confidence vector is detailed in Appendix B.

3.5. Rule base consistency validation with qualitative confidence levels

In this section, a mechanism is proposed to check the consistency of a rule base based on the definition of qualitative confidence levels in Section 3.1.1. The principle is to ensure that given a knowledge base and assuming all the rules are satisfied, no contradictory conclusions will be inferred from the same group of facts. For example, assume both fact p and its negation $\neg p$ are derived with confidence levels *Definite* from the same group of facts, since "$\neg p$ *is Definite*" implies that "p *is Impossible*", then p *(Definite)* and p *(Impossible)* are contradictory. Based on the six confidence levels defined in Table 2, six pairs of contradictory confidence levels are defined in Table 4, such as *Definite vs Impossible*.

For example, if we have four rules in a rule base (Fig. 6), and A is observed with confidence vector ⟨0, 0, 0, 0⟩ (Definite).

Then the four rules imply that:

$$C_{B;A} = C_A + C_{R1} = \langle 0,0,0,0\rangle;\ \ C_{D;B} = C_B + C_{R3} = \langle 0,0,0,1\rangle$$
$$C_{C;A} = C_A + C_{R2} = \langle 1,0,0,0\rangle;\ \ C_{D;C} = C_C + C_{R4} = \langle 1,0,0,0\rangle$$

It can be seen that D is considered to be *very likely* to happen through *Rule 1* and *Rule 3*; but it is also considered to be *very unlikely* to happen through *Rule 2* and *Rule 4*. If such contradictions are found in the rule definition phase, an alert will be shown to domain experts; the experts can either accept these inconsistencies[15], or add more conditions on the left hand side of the relevant rules or adjust the confidence levels of these rules.

[15] Note that, as mentioned by our domain experts, inconsistencies are not unusual in infrastructure engineering which is one reason for the level of uncertainty observed.

3.6. Extended confidence vector for diagnostic rules

Rules in the system can be classified as predictive rules or diagnostic rules. Predictive rules describe the relationship from cause to effect ("*If Cause, then Effect*"); and diagnostic rules describe the relationship from evidence to hypothesis ("*If Effect, then Possible Cause*"). For example, given an observed infrastructure defect, the confidence in a predictive rule reflects a stable property of this defect (i.e. the likelihood of a consequence to happen, given the defect). In contrast, the confidence in a diagnosis rule (i.e. the likelihood of a defect, given a consequence) depends on the incidence rates of that defect and other reasons that may cause the same consequence (Heckerman & Shortliffe, 1992). Generally, domain experts feel more comfortable when asked to formulate predictive rules than diagnostic rules since the incidence rate of different causes (*prior* probabilities) is often hard to define.

3.6.1. Extended confidence vector with an abduction count

Although the confidence level of diagnostic/abductive rules are hard to define, in order to warn users with all possible causes and the potential consequences of an observation but not mix the confidence in predictive and diagnostic rules, we extend the qualitative confidence vectors described in the previous sections by adding an extra element $\boldsymbol{A}$ to the front of a confidence vector, written as:

$$C = \langle \boldsymbol{A}; VU, U, L, V\rangle \tag{24}$$

For predictive rules, $\boldsymbol{A}$ is fixed as 0 whilst ⟨*VU, U, L, V*⟩ are defined by domain experts; for diagnostic rules, $\boldsymbol{A}$ is fixed as 1 whilst ⟨*VU, U, L, V*⟩ are fixed as ⟨0, 0, 0, 0⟩. For a fact, $\boldsymbol{A}$ records how many diagnostic rules have been used for inferring this fact, whilst ⟨*VU, U, L, V*⟩ represent its uncertainty (either provided at the beginning of an inference or propagated from other facts and predictive rules). A fact is also attached with a list of given facts (*GF*) and abductive facts (*AF*) for storing the corresponding antecedents from predictive and diagnosis rules. The three formulae (Eq. (7), Eq. (10) and Eq. (19)) defined in previous sections also apply to the extended (ordered) confidence vector by adding one element $\boldsymbol{A}$ into the calculation.

In our applications, predictive rules are defined by domain experts and diagnosis rules are automatically generated by reversing the cause and effect. For ease of clarification, in the following sections, we use $\overset{R}{\rightarrow}$ to represent predictive rules, and $\underset{R}{\rightarrow}$ to represent diagnostic rules. For example, a road crack could be triggered by several factors, such as traffic overloading, extreme temperature (e.g. freezing), water infiltration into the road due to nearby road cracks and rainfall, etc. A predictive rule *"Surface deformation can cause road cracks"* can be reversed as *"Road cracks could be caused by surface deformation"*, the two rules are written: "RoadDeformation (Active) $\xrightarrow{\langle 0;\ 0,0,0,0\rangle}$ RoadCracking (Active)" and "RoadCracking (Active) $\underset{\langle 1;\ 0,0,0,0\rangle}{\rightarrow}$ RoadDeformation (Active)".

3.6.2. Reasoning with both predictive and diagnostic rules

In practice, predictive and diagnostic rules are often used seamlessly. For example, given a defect, we hypothesise what is happening in the world to explain why this defect appears; then, we apply the predictive rules to infer all consequences potentially caused by deterioration of this defect. But if both predictive and diagnostic rules exist in one knowledge base, inter-causal reasoning may happen. For example, suppose we have one predictive rule "*Sprinkler (on)* $\xrightarrow{Definite}$ *Grass (wet)*" and one diagnostic rule "*Grass (wet)* $\underset{Likely}{\rightarrow}$ *Rain*" in a conventional rule-base, if we see the sprinkler is on, chaining forward through the rules, this will increase the belief that the grass will be wet, which in turn increases the belief

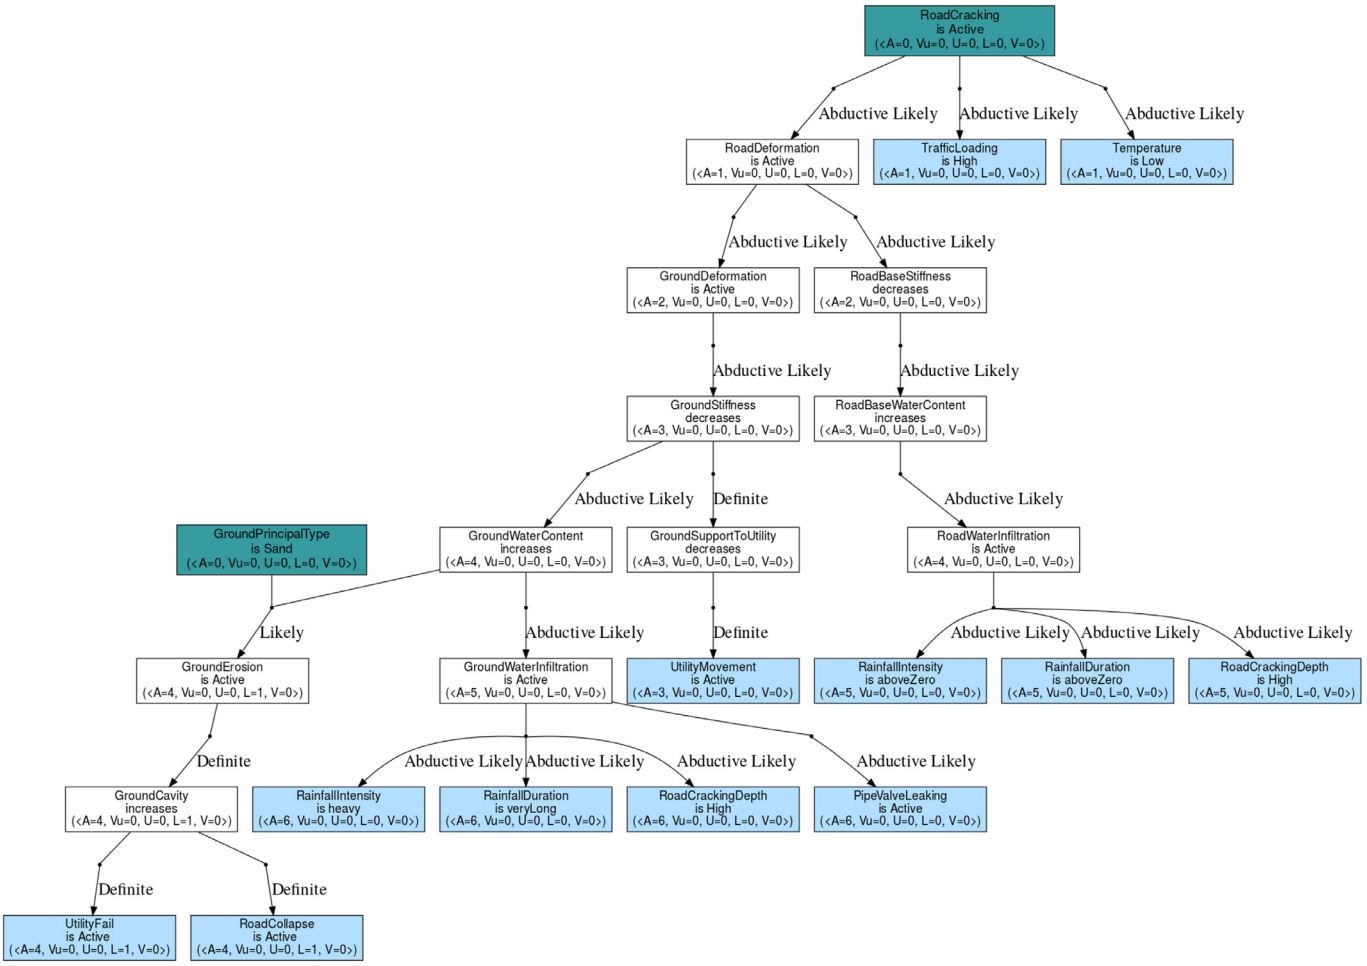

Fig. 7. An example of city infrastructure assessment with diagnosis and predictive inference. Colour scheme: turquoise for input facts, white for inferred facts, and light blue for final conclusions. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

that it is raining. To mitigate this interaction, pre-defined salience scores are added to all the rules such that abductive inference (using diagnostic rules) are performed first to find all possible causes of observed facts, followed by predictive inference (using predictive rules) to find all potential consequences. For example, as shown in Fig. 7, if road cracks are observed and the ground principal type is sand, all possible causes of the cracks are inferred based on diagnostic rules. The additional potential consequences of these inferred facts are then inferred through predictive rules.

3.7. Reasoning with missing facts

In addition to the uncertainty in rules, the domain experts we consulted also wanted to know what facts were assumed to be present in the derivation of a potential consequence so they could conduct further investigations to check whether these missing facts hold or not (Mahesar et al., 2017). To do this, a mechanism is provided in our system to handle incomplete data, i.e. if any premises (facts) of an inference rule are missing[16], the system will make assumptions of all possible states of the missing facts so related rules can still be fired. These missing facts (with assumed values) are also attached to the inferred fact as **({MF})**, the same as the given facts and abduced facts. The facts used for inferring a consequence will be displayed on the user interface for further guidance.

In order to infer all potential consequences, assumed facts are added into the knowledge base with all possible states of the missing facts (e.g. subgrade type = sand/rock/clay/gravel); these facts are combined with the rule base for reasoning. However, feeding different values of a fact (inconsistent information) into the same knowledge base may cause inter-causal problems since one inferred fact can be used regardless of its justification. For example, suppose we have one given fact A (Severe) and three rules in a knowledge base as below:

Rule 1: A is Severe $\wedge$ B is Clay $\xrightarrow{Likely}$ C increases;
Rule 2: A is Severe $\wedge$ B is Sand $\xrightarrow{VeryLikely}$ C increases;
Rule 3: C increases $\wedge$ B is Sand $\xrightarrow{Unlikely}$ *D* decreases;

As *A* is given, the system will make two assumptions about the missing fact B as “B is Clay” and “B is Sand”. Ideally, we only want rule 1 to fire to infer *C*(*increases*) or rule 2 and rule 3 to fire together to infer *C*(*increases*) and *D*(*decreases*). However, the *C*(*increases*) inferred from rule 1 will also cause rule 3 to fire in which case “B is Clay” and “B is Sand” cannot hold together. To avoid this, in Truth Maintenance System (Reason Maintenance Systems), a dependency network is often constructed to record the dependencies of derived facts so as to retract the inconsistent facts at the end of inference (Doyle, 1979), but this post-processing approach may allow exponentially many subsequent in-

[16] Currently, in order to limit the number of assumptions made, at least one premise of a rule must be not assumed.

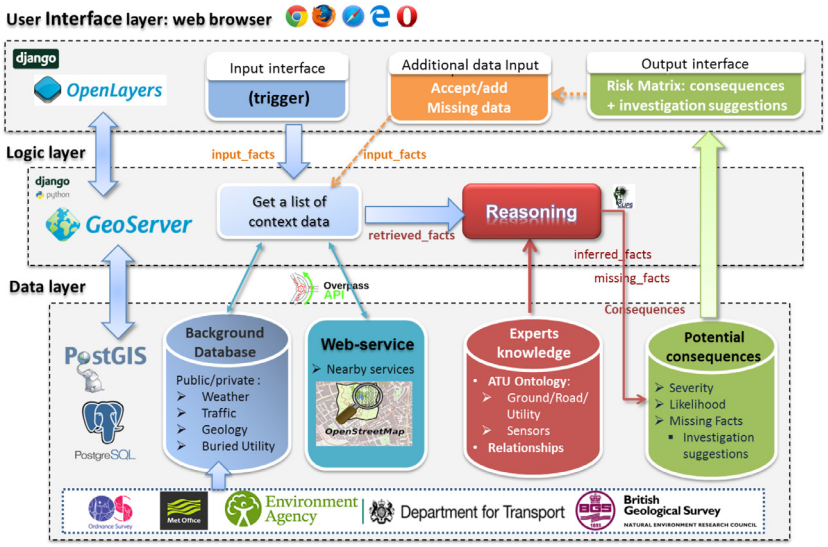

Fig. 8. A three-layer system architecture of the **ATU-DSS** prototype.

consistent facts to be inferred during the inference process. Therefore, in our system, we avoid logically inconsistent derivation on the fly by checking whether there are any contradictory antecedents in a rule. As shown previously in the example in Table 3, all given facts are initialised by including themselves in the given fact list, whilst all assumed facts (missing facts) are initialised by adding themselves in the missing fact list. For the above example, since $C(Increases)\{GF[A(Severe)], MF[B(Clay)]\}$ is not consistent with $B(Sand)$ $\{MF[B(Sand)]\}$, rule 3 will not fire with the C inferred from rule 1. This step can guarantee that each derived fact has a consistent antecedents.

Maximally consistent sets of assumptions. With the conclusions (e.g. consequences in our ATU-DSS) inferred from different combinations of given and assumed facts, we need to group consistent consequences based on different assumptions: 1) firstly, an adjacency matrix (undirected graph) of all consequences is generated: two consequences are connected if they are with consistent antecedents (i.e. no fact hold different values); 2) then, all maximal cliques in this undirected graph are identified, in each of which consequences are consistent and the clique cannot be extended by including one more adjacent vertex. Whilst typically systems would choose a preferred set of assumptions automatically based on certain criteria, we are looking at all sets simultaneously and our user interface in the prototype will let users compare the reasoning chains of different consequences and decide which group of assumptions is more reasonable (Fig. 12).

4. A prototype decision support system for urban infrastructure inter-asset management

A prototype has been developed based on the uncertain reasoning approach and the ATU ontologies presented in previous sections. The system architecture is shown in Fig. 8. It includes a data layer, a logic lager and an interface layer. In the following sections, we will first introduce the ATU-DSS rule base, then present the data sources, and finally present the user interface and demonstrate how to use the system with an example. A video demonstrating the prototype is available at: http://bit.ly/2MRHMCc.[17]

4.1. Developing a rule-base for ATU-DSS based on scenarios

Since the information needed to be encoded as rules in ATU-DSS is very extensive and the time of domain experts is precious, we adopted a scenario-based strategy for rule base development. First, several representative scenarios were selected by our domain experts, such as rainfall with road cracking and underground pipe leakage with active traffic loading. New scenarios can be added gradually subsequently. Then, for each scenario, rules were defined by domain experts by following the deterioration process of assets and considering all contextual possibilities. The flowcharts of these processes were sent to external domain experts and practitioners for validation. For example, in a scenario about *rainfall with road cracking*, the relevant variables include road construction properties, road cracking depth, rainfall duration/intensity and subgrade (ground) type. If the road surface is cracked and the cracks extend to the underlying road base, then any rainfall event will lead to infiltration into the road construction and underlying subgrade; if the subgrade is clay and the ground water level is low, it is likely that the ground water level will rise softening the clay; if the subgrade is a soluble rock it is possible that solution cavities could form, etc.

Rules with confidence levels were first created by domain experts in an agreed format and stored in text files; then they were automatically converted to a format recognisable by inference engines using a piece of code written in Python. We used the rule inference engine *Jess*[18] in our prototype (this can be interchanged with CLIPS[19] or other rule engines) for rule implementation and reasoning. A full-length exemplar code and explanation of the implementation are given in Appendix D. Diagnosis rules were also automatically generated by inverting the predictive rules.

4.2. Data sources and spatial criteria for data retrieval

Informed by the ATU ontologies and rules, various infrastructure and contextual datasets were sourced from different owners (e.g. UK Met-Office, UK Department for Transport, British Geological Survey, utility companies) and integrated in the prototype system to provide instant location/time specific data retrieval. The

[17] Accessed: 2020-02-23.

[18] http://www.jessrules.com/jess/docs/71/. Accessed: 2020-02-23.

[19] http://clipsrules.sourceforge.net/. Accessed: 2020-02-23.

sourced datasets are mostly in (or converted to) the form of GIS tables and stored in a PostgreSQL database.

The *Meteorological data* is sourced from the UK Met-Office.[20] When a trigger is reported through the system user interface, the 30-day weather data (e.g. daily/hourly rainfall, maximum and minimum air/concrete/soil temperatures) up to the occurrence day of the trigger is calculated based on the data from its nearby (less than 10km away) weather stations. Historic flood outline map is also sourced from the UK Environment Agency to provide the information of flood risk around the site of the trigger. Weather forecast data may also be added in the future.

The *Road and Traffic Information* is sourced from the UK Department for Transport[21] and Ordnance Survey. For a reported trigger, its nearest road segment is first retrieved from the road network database to identify the corresponding traffic counting point for obtaining the historical traffic data, based on which the weighted annual traffic on this road is calculated according to the wear factors of seven types of heavy vehicles (e.g. buses).[22] The road designed traffic loading can be provided by external road datasets (e.g. the National Street Gazetteer) or added by users. By comparing the traffic volume on a road with the average volume in the locality, the importance of this road can be assessed, which is also an indicator of the effect that the trigger and subsequent consequences or mitigation measures could have on the traffic flow. The past and future planned roadworks on this road are also retrieved (data from Highways England website[23]) to help evaluate the vulnerability of the road system.

The *Ground Conditions* data (e.g. ground water level, geological faults) is sourced from the British Geological Survey (BGS) and local councils. For example, the *BGS 50K dataset*[24] provides geological information like superficial and bedrock geology; the *corrosivity dataset*[25] gives an indication as to whether the ground conditions below the top soil are likely to cause corrosion of underground iron asset, and the *SuDS*[26] dataset suggests the potential presence of geological and hydrogeological hazards that could be initiated or worsened by water infiltration to the ground. Brownfield information sourced from different local councils is also added in the system.

Buried utilities data is sourced from different asset owners (e.g. United Utilities, National Grid, North West Electricity), vectorised and integrated in a back-end spatial database,[27] in which multiple attributes of the buried assets are recorded (e.g. utility type, location, depth, material, diameter, pressure/voltage, year of installation, owners, operation status). Based on the location of a reported trigger, data of buried utilities in the general vicinity of this trigger (within a 200m radius) is retrieved from the database (Fig. 11(a)).

Nearby Services. Information of the nearby services such as hospitals, schools, banks, is also important for estimating the potential social and economic impact of a trigger. For example, the people at schools and hospitals may be more vulnerable to harm and harder to evacuate in an emergency, such as a gas explosion due to damage to a gas pipe. The data of sensitive services around a trigger ($\leq 2km$) is fetched from the OpenStreetMap when a trigger is reported (Fig. 11(b)). To assess the potential impact on each service, their shortest driving distances to the trigger is calculated using the widely used path planning A* algorithm.

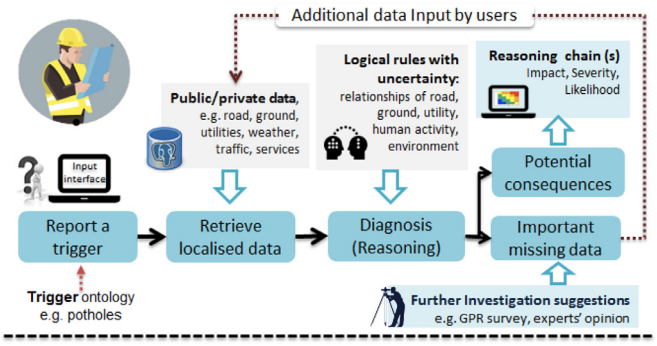

Fig. 9. Workflow in the ATU-DSS (from the left to right) (Wei et al., 2018).

Mapping Data to ATU Ontology Concepts. The sourced data is mapped to corresponding ontology concepts based on a predefined correspondence table so that it can be used for inference together with the logical rules.

Although automatic methods exist for matching data and ontologies (Munir & Anjum, 2018), in our case, all the correspondences were manually defined by experts to guarantee their correctness. We carefully read the related documents of each dataset, especially their definitions, the meaning of each table column (or attribute, field), the data unit, and how the value in each cell was derived. Some examples of mapping are given below:

- For some cases, the name of a data table (or a column) can be similar to an ontology concept, which gives a hint to find the correct match. For example, the table column *length* in the OS Open Roads Dataset[28] (table *uk_road_network*) also suggests the length of a road segment, thus can be mapped to the ontology concept *RoadLength* in the Road Ontology;
- But for some cases, similar names may not suggest a correct correspondence. For example, there exists a geological dataset *Depth_to_water_table*[29] and an ontology concept *GroundWaterTableDepth* in the Ground Ontology. The two names look similar but the depth data cannot be mapped to the ontology concept as the data values were not the real depth but categorised into [1, 2, 3] (1: $> 5m$ below ground surface; 2: 3-5m below ground surface; 3: $< 3m$ below ground surface);
- For some cases, there are no similarities between the names, correspondences can only be established after checking the definition of data and ontology concepts. For example, the table column *Function* in the *OS Open Roads Dataset* can be mapped to the ontology concept *RoadType* in the Road Ontology because both of them give the information of road types such as A road, B road, etc. Another example is about a geological dataset called *GroundWaterLevels*[30]. Its document suggests that this dataset provides the information of depth to groundwater level and the unit of its table column *Value* is in metres which is the same as the definition of the ontology concept *GroundWaterTableDepth*, therefore, a correspondence is established between them.

All these correspondence relations are stored in a predefined correspondence table (Table 5).

[20] Met Office Integrated Data Archive System (MIDAS) Land and Marine Surface Stations Data (1853-current).

[21] UK Department for Transport (DfT) Traffic Statistics.

[22] UK Design Manual for Roads and Bridges (2006).

[23] https://data.gov.uk/dataset/5b3267d8-4307-4eef-a9af-3a4c28224694/planned-road-works-on-the-he-road-network. Accessed: 2020-02-23.

[24] http://www.bgs.ac.uk/products/digitalmaps/digmapgb_50.html. Accessed: 2020-02-23.

[25] http://www.bgs.ac.uk/products/groundconditions/corrosivity.html. Accessed: 2020-02-23.

[26] http://www.bgs.ac.uk/suds/. Accessed: 2020-02-23.

[27] For proof-of-concept purpose only. In the UK, automated retrieval of utility statutory records is only available in Scotland with the VAULT system.

[28] https://www.ordnancesurvey.co.uk/business-government/products/open-map-roads. Accessed: 2020-03-01.

[29] http://www.bgs.ac.uk/suds/. Accessed: 2020-02-23.

[30] https://www.bgs.ac.uk/products/hydrogeology/depthToGroundwater.html. Accessed: 2020-02-23.

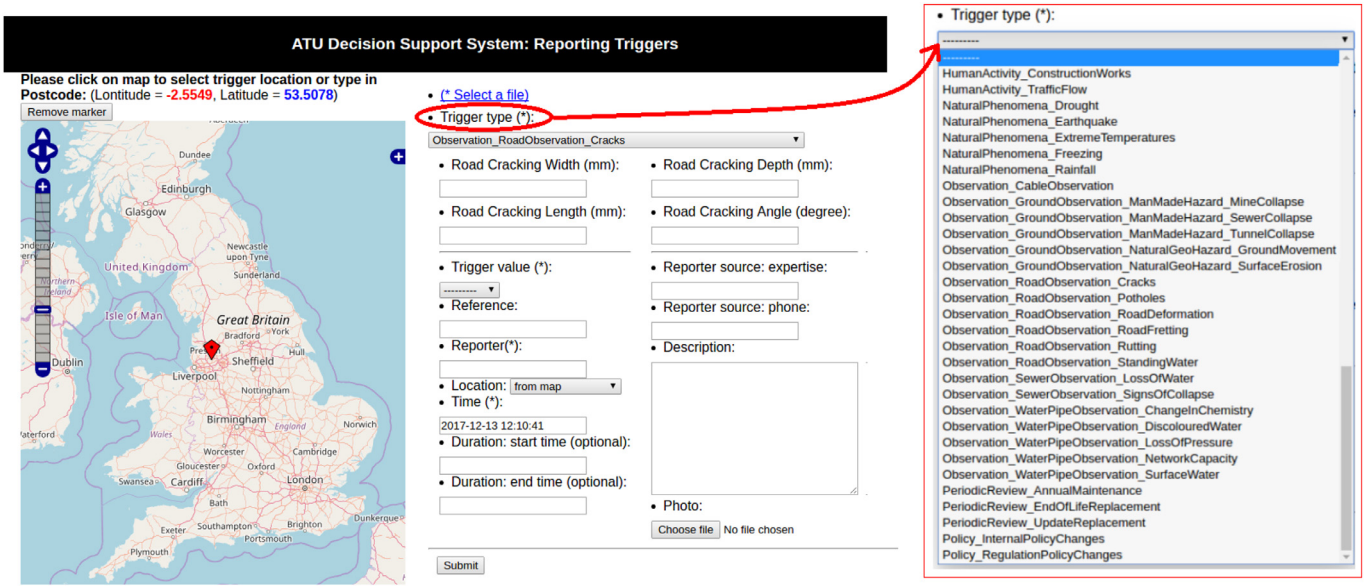

Fig. 10. User interface for reporting new triggers with a drop-down menu of trigger types (©OSM).

Table 5
A predefined correspondence table between data and ontologies.

externalData_name	externalData_column	ontologyConcept
uk_road_network	length	RoadLength
uk_road_network	Function	RoadType
GroundWaterLevels	Value	GroundWaterTableDepth
...	...	...

4.3. ATU-DSS User interface

For ease or use, a web-based user interface has been developed using *Python Django, Geoserver* and *openLayers*. The workflow of utilising ATU-DSS is shown in Fig. 9, including four steps: first users can access the system from standard web browsers to report new triggers by providing the trigger type and properties; then, relevant contextual information is retrieved from local or online databases; after that, the retrieved data is both displayed on the user interface and fed into the rule engine to infer potential consequences with their associated uncertainty, severity levels and antecedents. Finally, the system also gives suggestions on how to get the missing data and supports alternative assessment.

For example, assuming that:

> A member of the public observes a water pipe leakage and phones up the local authority to report it!

In order to estimate the potential consequences of this trigger, the local authority first need to report the *leakage* through the user interface:

- *Reporting new triggers*. Users can either report a trigger by manually typing the information (Fig. 10), or uploading an XML file containing the information of triggers. The second option allows ATU-DSS to be connected with other existing information systems, such as the pothole reporting systems used in a lot of local councils, and to use external data sources as triggers to start the decision support process. When manually reporting a trigger, the trigger type can be selected from a type list (defined in the *Trigger Ontology*) as shown in Fig. 10. The trigger severity level (*High, Medium, Low*), geographic location (e.g. postcode, GPS coordinates, location pinpointed by users on the displayed map, or an uploaded spatial file, Fig. 10) and time should also be provided. Users can also upload multiple photos of a trigger at different times to help analyse/monitor the development of this trigger.
- *Localised Data Retrieval and Automated Reasoning*. Then, the relevant localised contextual data of the reported trigger is automatically retrieved based on its occurrence location and time using different spatial criteria. The retrieved/processed data is displayed on the user interface (Fig. 11). The data provided by users (e.g. trigger information) and data retrieved from the database is written as a fact file (e.g. *GroundPrincipalType (Clay), EnvironmentRainfallDuration (Long)*), then fed into the rule engine for automated reasoning of potential consequences.
- *Identification of Potential consequences*. Once the reasoning process finishes, potential consequences are identified from the inferred facts. Currently we are looking at four types of consequences in ATU-DSS, including consequences on buried utilities (e.g. utility fail), on roads (e.g. road collapse), on ground (e.g. ground collapse) and social/economic/environmental/legal consequences (e.g. traffic disruption, loss of business, loss of utility service, damage to property, injury, and loss of life). Each consequence has five attributes attached: uncertainty level, severity level, given facts {*GF*}, missing information {*MF*} and consequence type (e.g. ground, social/economic)[31]
The uncertainty level, given facts and missing facts of a consequence are propagated/accumulated from the uncertainty reasoning approach, whilst the severity level (*Negligible, Marginal, Critical, Catastrophic*) of a consequence is context dependent and defined in each logical rule. For example, if road deformation happens on a road, the severity of potential traffic disruption is only marginal; but if a pipe burst happens on a road, the potential traffic disruption can be critical. Knowing the likelihood and severity of potential consequences of a reported trigger can help users to prioritise their tasks and take appropriate mitigation measures to reduce the potential risks, especially for those with higher likelihood and (or) higher severity.

[31] Note: as no other causes of pipe leakage is provided by domain experts for the current prototype, the abductive reasoning results are not shown.

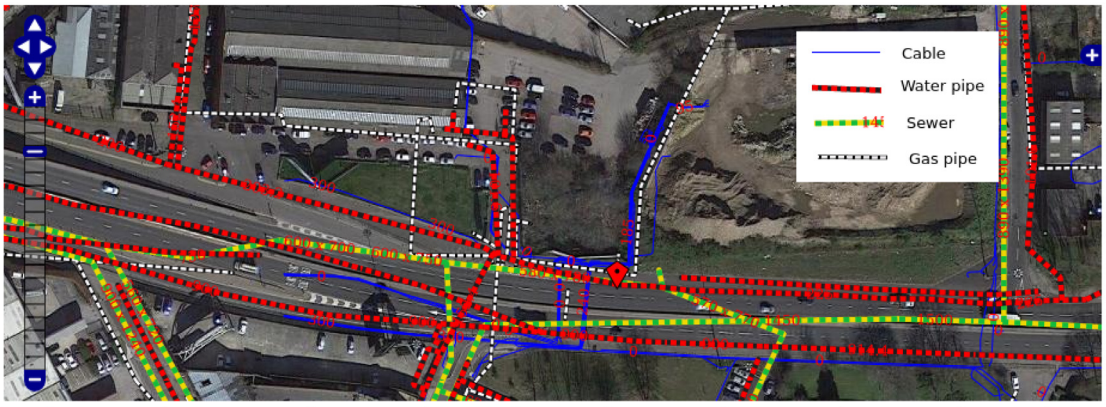

Id	Utility_type	Sub Type	State	Date Installed	Owner	Material	Voltage / Pressure	Depth_(Mm)	Source	Size_(
203	water pipe	Distribution Main - PressurisedMain	operation	1997	united utilities	Polyethylene	—	—	desktop-survey	900
192	water pipe	Distribution Main - PressurisedMain	operation	1932	united utilities	cast iron	—	—	desktop-survey	101.6

(a) Underground utilities around the trigger.

> Nearby services:

The blue region indicates the area which can be accessed in 10 minutes by driving from the trigger. The nearby services within 2km are displayed by coloured marks. Red marks: bank; Yellow marks: hospital; Blue marks: schools.

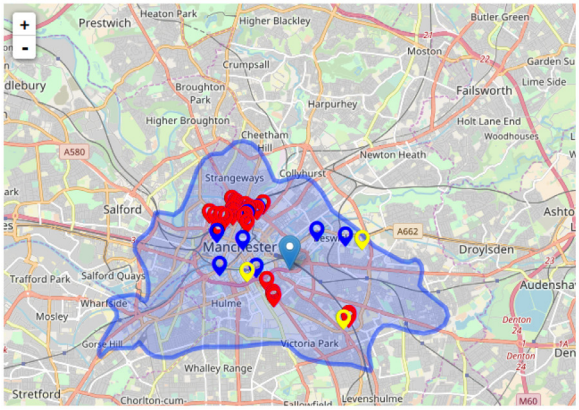

> There are 32 banks around the trigger, the area is close to city center.

- Direct distance: as the crow flies (straight line between two points);
- Routing distance: the shortest routing distance along road by driving.

ID	Service	Name	Direct_Distance (Km)	Routing Distance (Km)
10639	university	Manchester Business School	1.254	2.0916
10675	university	MMU 70 Oxford Street	1.6293	2.2092
10677	university	University of Law	1.59	2.2074
10652	school	—	1.2393	2.5913
10661	school	Instituto Cervantes	2.6441	3.2483
10649	school	Berlitz Language Centre	1.3161	1.9325
10651	school	—	2.3197	3.2785
10660	kindergarten	—	2.5618	2.9568
10641	hospital	Cornerstone Centre	2.9858	3.9
10638	hospital	Boundary Medical Centre	1.5953	2.0113

(b) Nearby services.

Fig. 11. Snapshots of the user interface with a list of retrieved contextual data (background maps (a) Google Satellite Image, (b) OpenStreetMap).

– *Visualisation of the Potential Consequences.* The uncertainty level of potential consequences can be shown to users as a confidence vector or using a textual representation. In this work, the textual representation is obtained by taking the first element in the $\langle VU, U, L, V\rangle$ vector with a non-zero value since we assume an order-of-magnitude semantics for the confidence vectors, i.e. the left-most component is more surprising than the next unless there were a huge number of rules applied to get the first value (Appendix A). For example, for an inferred fact with a confidence vector $\langle 0, 0, 2, 1\rangle$, since 2 on the position of *Likely* is the first non-zero value, this fact is considered as *Likely* to happen. To meet the different needs of users, two views have been designed in ATU-DSS for visualising the potential consequences:

The first option is to display the consequences in an impact matrix table (Fig. 12) according to their severity levels/impact and uncertainty level/likelihood. The number of missing facts (if there are any) used for deducing a certain consequence is also displayed behind each consequence. As we make assumptions of all possible values of missing facts, multiple instances of the same consequence (e.g. utility fail) could be derived from different sets of facts and with different likelihood. To ease the analysis by users, we have added several filtering boxes on the right panel of the matrix view so that users can either filter consequences based on their categories (e.g. social consequences) or combinations of different assumptions. When users hover the mouse on one consequence in the risk table, other duplicate consequences are also highlighted; a tooltip will also appear to show the facts used to derive the hovered consequence (Fig. 12).

The second option is a list view (Fig. 13) in which consequences and their related parameters (e.g. confidence vector, likelihood, severity, number of missing facts, data used) are displayed in a table and can be sorted according to different attributes. It should be noted that by keeping elements in the confidence vectors with the same length, the alphabetically sorting algorithm used in an html

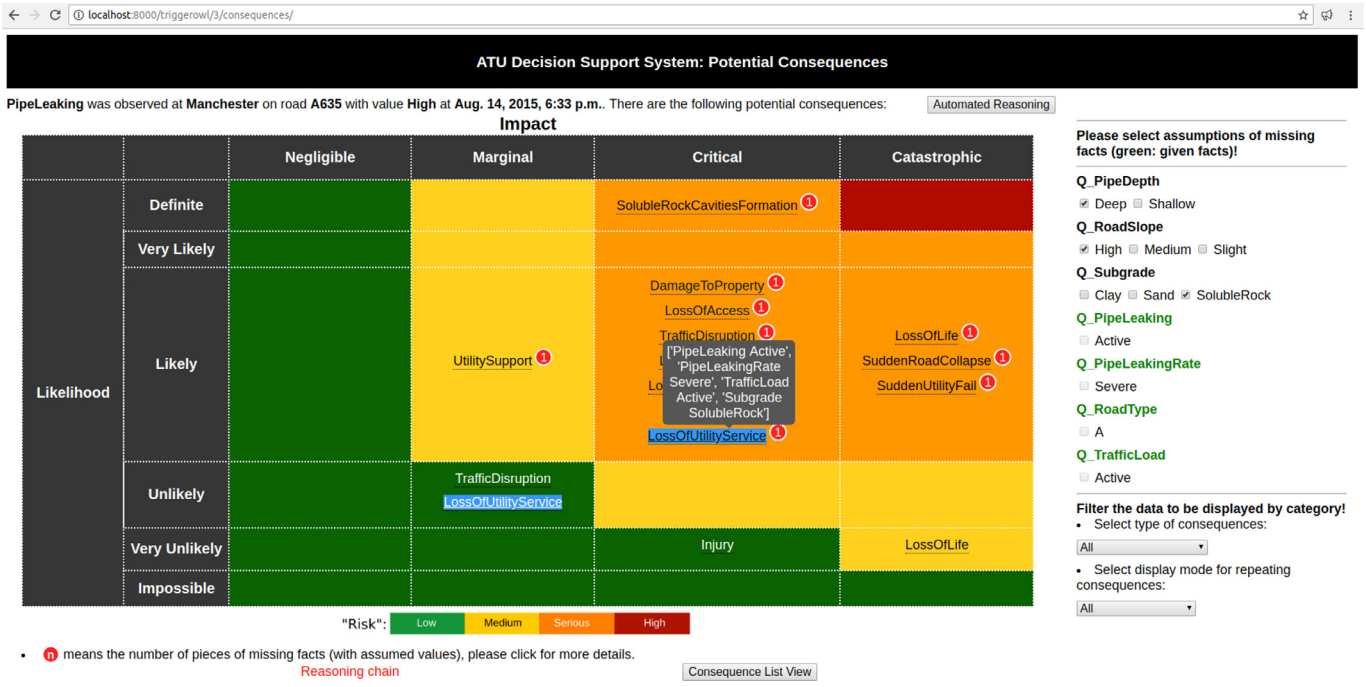

Fig. 12. The impact matrix view for visualisation of the potential consequences of a reported pipe leakage. There are several filtering options on the user interface: (a) users can select different combinations of assumptions of the missing facts from the multi-checkboxes on the right panel and the corresponding consequences will be shown in the impact matrix table (N.B. colours in the table are assigned based on generic heuristics); the facts used to infer a specific consequence are shown with a mouse-over effect in a tooltip; (b) users can also filter the consequences to be displayed by their categories (e.g. road, social/economic consequences); c) users can click on a consequence and more details will be shown in a new page.

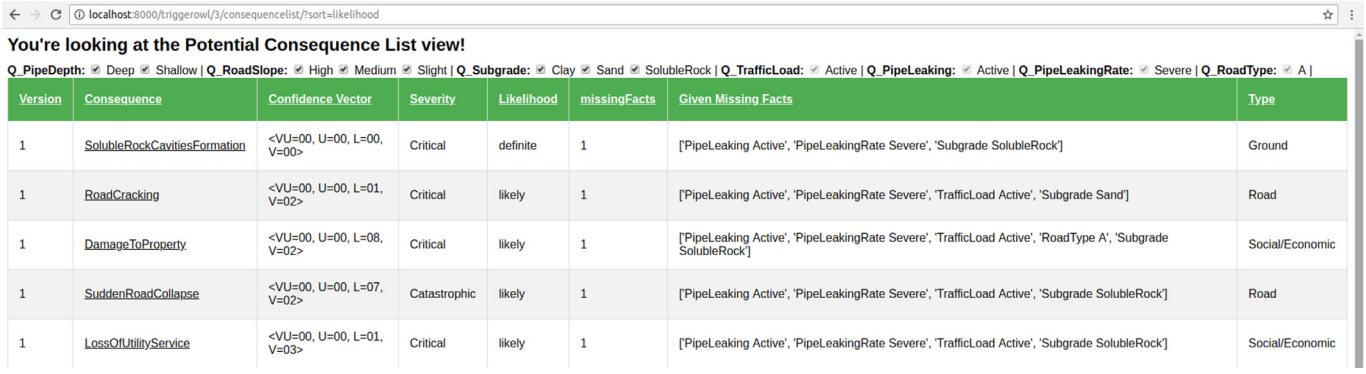

Version	Consequence	Confidence Vector	Severity	Likelihood	missingFacts	Given Missing Facts	Type
1	SolubleRockCavitiesFormation	<VU=00, U=00, L=00, V=00>	Critical	definite	1	['PipeLeaking Active', 'PipeLeakingRate Severe', 'Subgrade SolubleRock']	Ground
1	RoadCracking	<VU=00, U=00, L=01, V=02>	Critical	likely	1	['PipeLeaking Active', 'PipeLeakingRate Severe', 'TrafficLoad Active', 'Subgrade Sand']	Road
1	DamageToProperty	<VU=00, U=00, L=08, V=02>	Critical	likely	1	['PipeLeaking Active', 'PipeLeakingRate Severe', 'TrafficLoad Active', 'RoadType A', 'Subgrade SolubleRock']	Social/Economic
1	SuddenRoadCollapse	<VU=00, U=00, L=07, V=02>	Catastrophic	likely	1	['PipeLeaking Active', 'PipeLeakingRate Severe', 'TrafficLoad Active', 'Subgrade SolubleRock']	Road
1	LossOfUtilityService	<VU=00, U=00, L=01, V=03>	Critical	likely	1	['PipeLeaking Active', 'PipeLeakingRate Severe', 'TrafficLoad Active', 'Subgrade SolubleRock']	Social/Economic

Fig. 13. The list view for visualisation of potential consequences of a given trigger. Consequences in the table can be sorted according to different attributes.

table is the same as the formula for confidence vector comparison defined in Eq. (20).

- *Reasoning Chain of Potential Consequences.* Users can explore the details of each potential consequence by clicking on the links in the risk view or the list view. The system can give explanations of each consequence in the form of a sequence of text descriptions or as a network diagram of nodes and directed arcs. For example, the reasoning chain of a potential *utility support decrease* caused by a reported *pipe leakage* is illustrated in Fig. 14, in which the colour of ellipses indicates whether a fact is given, assumed or inferred; and the arrows indicate the reasoning flow together with the likelihood of each rule. By showing users the reasoning process of arriving at a particular consequence, the system can help users make a more reasonable and confident decision. Users can also manually update (e.g. decrease) the confidence and severity levels of a consequence based on their expert opinion.
- *Investigation Suggestions for Collecting and Updating Missing Data.* As mentioned previously, in cases where real data is missing in the reasoning process of a potential consequence, the system will suggest suitable investigation techniques to get the missing data based on the *ATU Investigation Ontology* (Fig. 16). Users can decide whether to accept the assumed value of a missing fact or to do some investigation and update the data later (Fig. 16). When new data is added/updated, the whole reasoning process automatically re-activates. We designed this function since it is often rare to have complete datasets in an information system. Even if such datasets do exist, they are subject to change or the observations from different people may vary. For example, the system may suggest that the ground type at a specific location is *Sand*, but one user may suggest the ground type as *Clay* based on his/her investigation. In ATU-DSS, the value of a fact can be updated no matter whether it is retrieved from the database or provided by a user. The system records which user makes a modification so that the provenance of all data is recorded. The ability for users to see alternative results by modifying assumptions can help users understand better the impact of a trigger in different context.

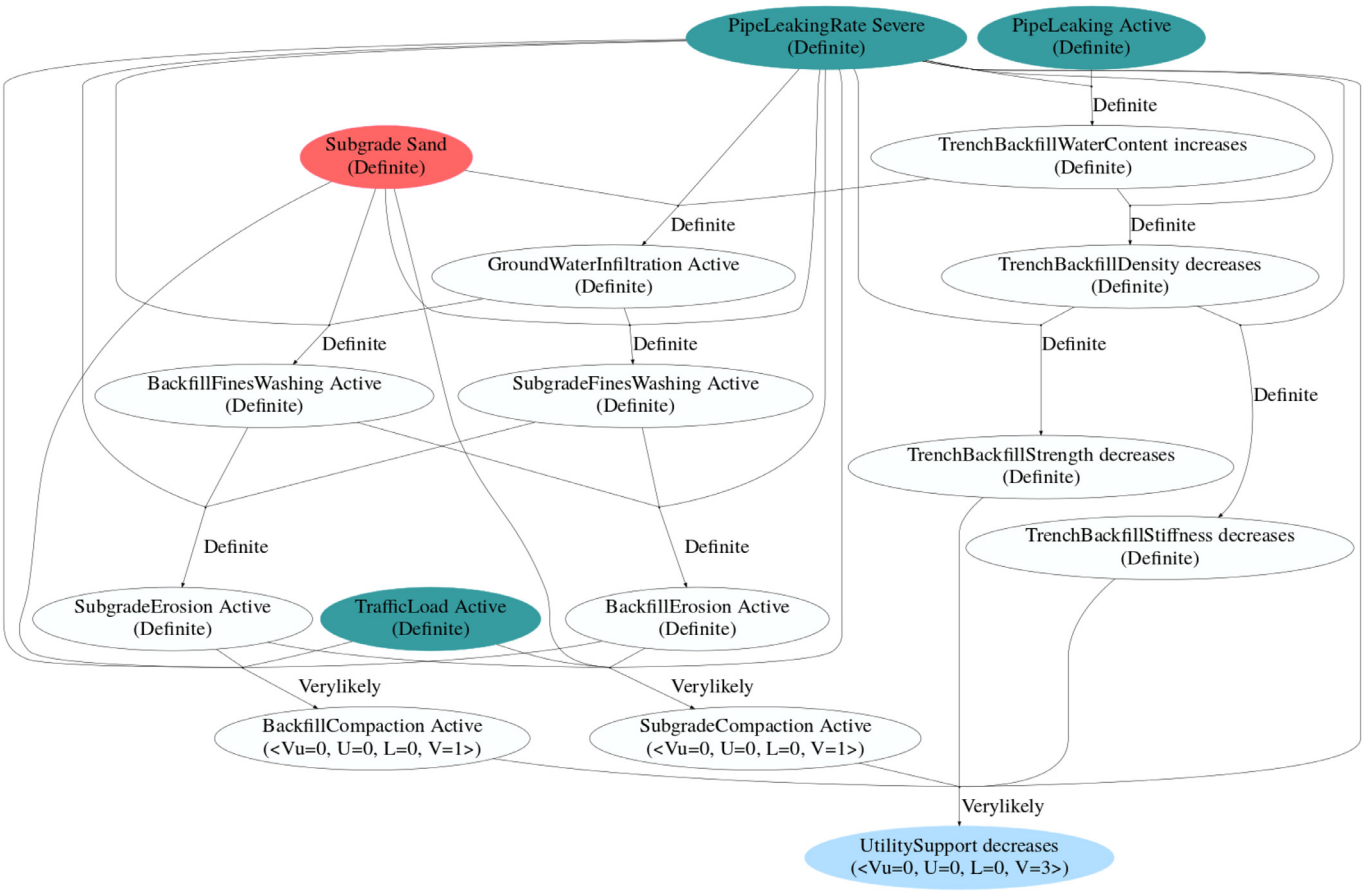

Fig. 14. Example of a consequence reasoning chain. Colour scheme in the figure: turquoise for given facts, pink for assumed facts (missing facts), white for intermediate facts and yellow for the final output. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

– *Modification of Rule Likelihood for Case-based Analysis*. The same as we may have different measurements of a fact, different experts may have different opinions of the content (e.g. likelihood) of a rule. Therefore, our system not only provides a common rule base for all scenarios, it also allows users to modify/personalise these rules. As shown in Fig. 15(a), users can click on the likelihood of a certain rule; then, a separate page will be shown (Fig. 15(b)). On this page, users can select the new rule likelihood from a drop-down list and add their comments or explanations. Once submitted, this rule will replace the original rule in this case study (i.e. this specific trigger) and the whole reasoning process re-activates.[32] The re-reasoning step is essential since the modified rule may have also been used for inferring other potential consequences. User can notice whether a rule has been modified in this scenario by the colour of likelihood on a reasoning graph: the original ones are in blue while the modified ones are shown in red. The modification history of each rule will be recorded and shown on the bottom of the page so users can see others' insights.

4.4. Computational complexity

We used the rule inference engine *Jess* for rule implementation and reasoning in the prototype. As *Jess* uses an improved Rete algorithm (Forgy, 1982) for reasoning, the performance is largely independent of the number of rules/facts, but as we are also dealing with missing facts, the computational time of different scenarios will depend on the number of missing facts in the antecedents of a rule and the possible values of each missing fact, as well as the similar patterns of rules' LHS (left hand side) in the knowledge base. Therefore, it is always a good practice to put the most specific patterns near the top of each rule's LHS. The system currently comprises 377 predictive rules and each rule has 4 antecedents in average; the most complex rule has 12 antecedents and the simplest rule has two antecedents. For the pipe leaking scenario illustrated in Fig. 12, two facts were provided (*PipeLeaking* (Active), *PipeLeakingRate* (Severe)), two facts were retrieved (*RoadType* (A class road), *TrafficLoad* (Active)), and three missing facts were assumed with possible values: *PipeDepth* (Deep) (Shallow), *RoadSlope* (High) (Medium) (Slight), *Subgrade* (Clay) (Sand) (SolubleRock). 235 facts were inferred from different combinations of facts and 65 were identified as potential consequences. The maximum depth of inference is 22 and the minimum depth is only one. The inference, data tidying and figure rendering process took about 4 seconds on a laptop with an Intel 2.7-GHz processor.

4.5. Evaluation

In addition to test the system with several real scenarios (e.g. the pipe leaking scenario presented in previous sections), we have also evaluated users' acceptance of the system by demonstrating the ATU-DSS prototype to a wide range of potential users in two workshops and collecting their feedback. The first workshop was organised in September 2017 to assess the system framework and user interface design. The participants included one invited experienced utility manager, one utility surveyor trainers, and around

[32] It should be noted that when a new trigger is added, only the original rule base will be used.

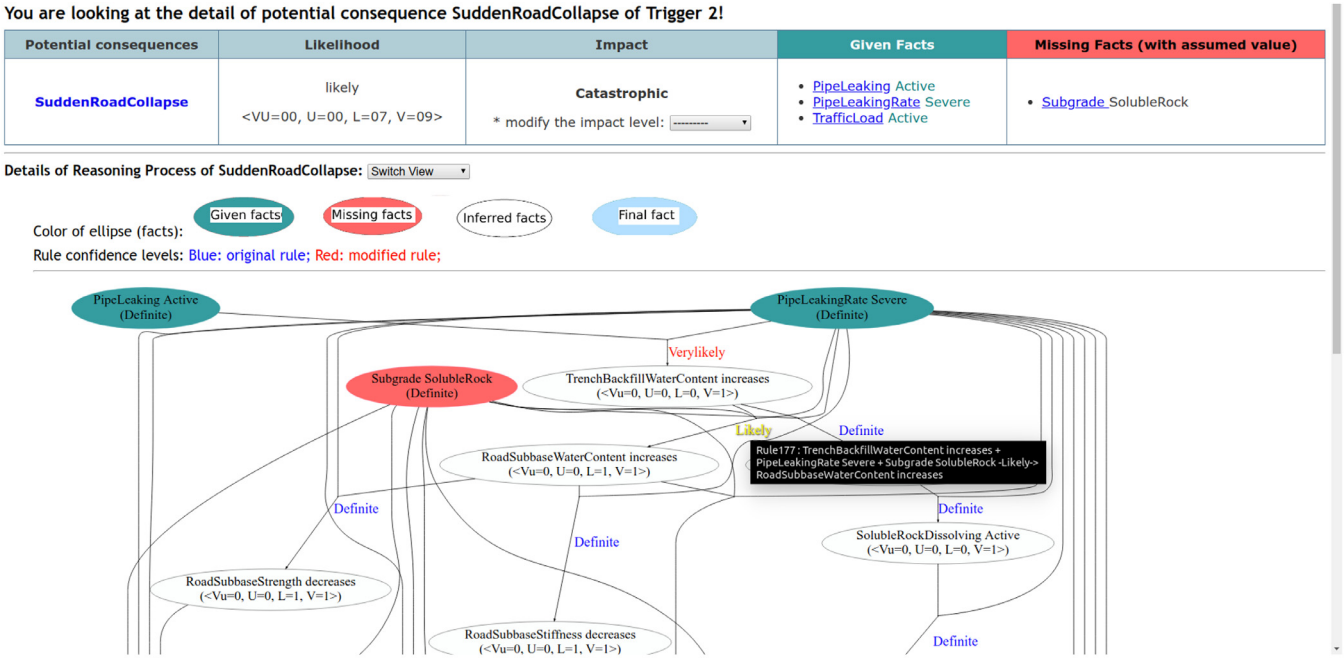

(a) User interface for displaying the reasoning process of a potential consequence. Users can click on a fact to change its value or check its modification history. User can also click on the likelihood of a rule to modify it.

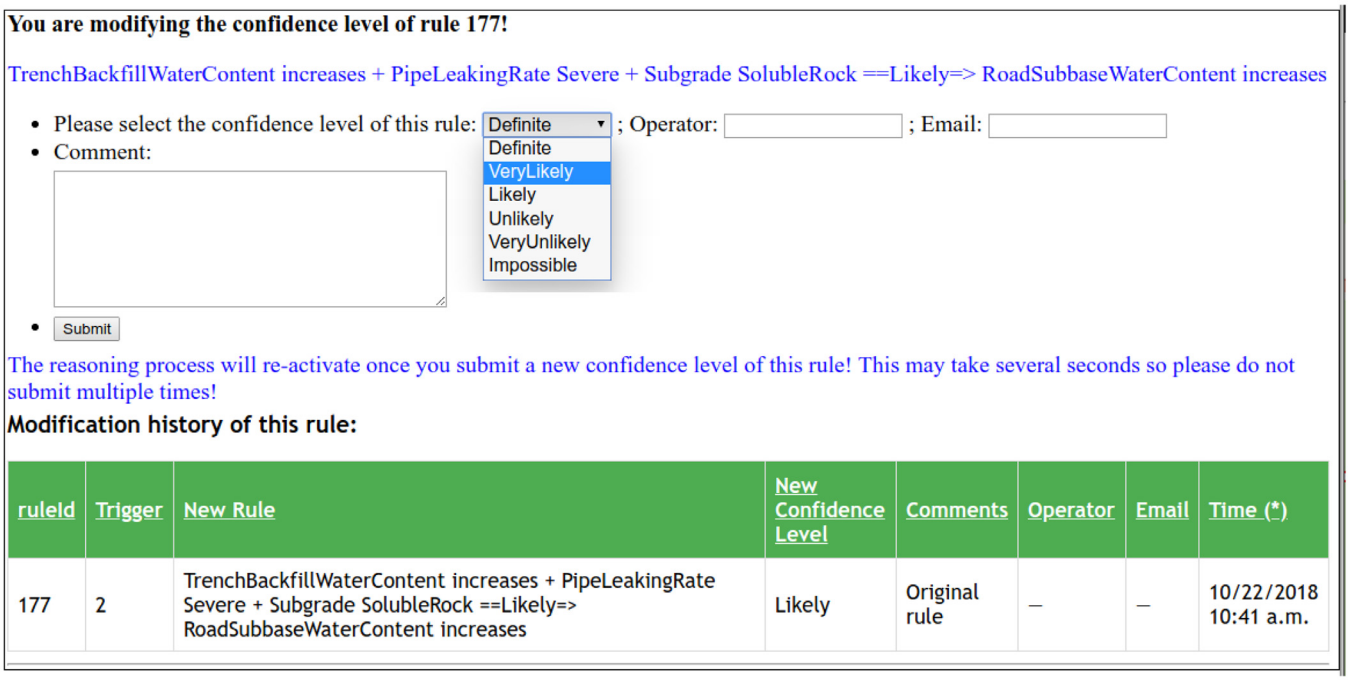

(b) User interface for modifying the likelihood of a rule. The modification history of this rule is also recorded and displayed.

Fig. 15. User interface for displaying and modifying the likelihood of a rule.

twenty academics from different UK universities and institutions (with diverse backgrounds such as civil engineering, geotechnical engineering, geophysics, computer science). An overview of the decision support system (including the framework, data and underlying semantic technologies) was given at the beginning of the workshop, followed by a live demonstration with real data from a historic ground collapse in Manchester which caused major disruption. After that, feedback from participants was acquired via a plenary discussion. The participants showed great interest and generally praised the effectiveness of the system, especially the rich data provided, the transparency of the reasoning module and the informative investigation suggestions. Several suggestions were received regarding the user interface design (e.g. colour in the matrix table, legend).

The second workshop was organised in November 2017 and attracted attendees from various backgrounds, e.g. local authorities, utility companies, survey companies, contractors (utility pipe lining/design), risk managers, sensor manufacturing companies, individual consultants, academics. The workshop followed the same procedure as the first one, except that feedback from participants

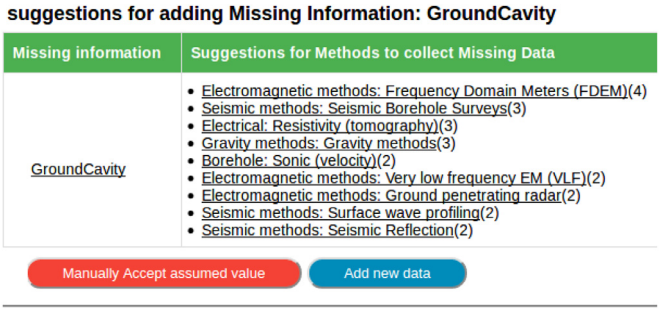

Fig. 16. Interface for suggesting the investigation techniques for obtaining a missing fact (e.g. *Ground Cavity*), where 0 = Not considered applicable; 1 = limited use; 2 = used or could be used, but not best approach or has limitations; 3 = excellent potential but not fully developed; 4 = generally considered an excellent approach, techniques well developed. There are also two buttons linking to two different pages for updating the data value.

was acquired via individual questionnaires as well as a plenary discussion. In the questionnaire (available at http://bit.ly/2oL7MAt),[33] participants were asked whether the prototype addressed key problems in their practice and whether it would fit into their current work. 18 questionnaires were collected - some sample responses from the participants are shown in Table 6. Among these questionnaires, ten participants answered the question "Is this useful for you or somebody you can think of ?" and suggested that this system can be a potentially useful tool for different stakeholders, such as incident managers, survey company developers, constructors, asset owners, local authority. Possible tasks included risk mitigation, prioritisation and justification of asset design and maintenance expenses and activities. The participants also pointed out that the system could be useful for the general public and could have the potential benefits for training novice or junior staff in streetworks management. Eight people did not answer this question but their responses in questions "How is this related to your practice ?" and "Does it address any specific challenges you are facing ?" indicated that the DSS was not directly relevant to them because they were from a company selling sensors, were academics, or in one case because "current practice is so far behind the presented ideas that it is hard to imagine it in use for xx years yet". Regarding the specific functions of the system, the users were particularly interested in the integrated data platform that brought various critical contextual data together. The participants suggested that the automated reasoning module was useful for helping determine the impact of an incident in a short period of time, identify potential consequences from seemingly insignificant triggers and potentially reduce the streetworks disruptions. One participant was particularly pleased with the visualisation of the reasoning chain, which could help users to better understand the cause and effect in relation to understanding spatio-temporal correlation between utility and road problems.

As for future improvements of the system, participants suggested to add additional data sources, such as bus routes, agriculture data, and archaeological data. It was also recommended to develop a smart phone application for easier access to the system. We are encouraged by the number of new stakeholders interested in the system since the workshop and are exploring further case studies (e.g. diagnosing leaking pipes, surveyor training, route planning for street excavations) under a recently started follow-on impact acceleration project.

[33] Accessed: 2020-02-23.

5. Discussion

In this section, we discuss the key challenges we met and lessons we learned from building the above knowledge-driven decision support system for urban infrastructure inter-asset management, as well as the advantages and limitations of the proposed framework.

5.1. Distributed and integrated rule base development

In order to predict the future behaviour of infrastructure assets, in this work, we captured the asset dependencies using rules. However, it is not feasible to create an exhaustive rule base for urban infrastructure management, especially when working with few domain experts whose time is extremely valuable. In this work, we created a rule base with two scenarios to demonstrate the applicability of the proposed reasoning and decision supporting framework. For larger scale applications, rule bases can be created by different experts or organisations (e.g. industry) for different scenarios/applications in a distributed manner. Importantly, the ATU ontologies can provide a common language to facilitate this process. We have discussed the problem of rule base inconsistency in Section 3.5 by assuming that all rules were created with the same confidence levels. However, in practice, different experts may have different understanding of a probability phrase (Wallsten & Budescu, 1995) and these phrases tend to change their meanings in different contexts. Future work is required to investigate how to combine/align the rules created by several domain experts (inter-variances) and the rules created by the same domain experts at different times (intra-variances) or in different context. Section 3.5

5.2. Transferability to other applications

In this work, we provided a framework for developing knowledge-driven DSSs based on ontologies and rules with qualitative confidences. This framework can be adapted for various engineering applications, which require a system-of-systems thinking and qualitative uncertainties. Since users can change data values and compare different alternatives with explanations, the system framework can also be used for training less experienced engineers to perform complex decision making that requires multi-sector knowledge. In addition to the above applications, the system can be used as an evidence base to store all reported triggers, decisions/mitigation made by users, and the actual consequences for sharing experiences and lessons learned.

5.3. System maintenance: balancing rule-driven and data-driven approaches in long term

For building a decision support system for infrastructure management, both the historic and current data is important. As the historic data of infrastructure behaviours is not often available or can be incomplete, it is practical by starting working with domain experts to identify the key data/knowledge and encode the essential process rules (i.e. cause and effect) in representative situations. It should be noted that these hand-crafted rules are not immutable, instead they are only used as the starting points to understand a domain, especially how a domain expert would evaluate the risk/potential consequences from certain observations, such as the level of details of knowledge they used and the order of their inference. As more data is gradually available, especially in the context of Internet-of-Things as more smart sensors are installed to monitor the cities, the hand-crafted rules could be used to guide quantitative/logic rules learning and to validate the existing rules provided by domain experts. Another function our system provides is to let users to specialise a rule by modifying the

Table 6
Users' feedback about the **ATU-DSS** from a user workshop.

Questions	Users' response	
	Category	Sample responses from workshop participants
1. What excites you about the *ATU-DSS*?	Integrated data platform	"The pool of information on a single platform, from the output his seems limitless which is very exciting" "A one step location for critical information" "by having a system with the potential to show and provide all records and local specific information is a huge benefit"
	Help estimate potential consequences through automated reasoning	"Ability to determine the impact of an incident in a short period of time and consider vulnerability of all assets" <ddq>'Bring expert analysis across a consistent analysis / reasoned approach'</ddq> 'Help to better visualise cause and effect in relation to understanding temporal spatial correlation between utility and road problems" "Identifying real and potential future problems from what may appear an insignificant trigger" "The opportunity for decisions on interventions on different utilities / highway authorities assets based on rational reasoning"
	Smart city/ Minimise disruption	"Potential tool to help reduce streetworks disruption by improving the decisions made to arrive at appropriate responses and associated consequences (vulnerabilities)" "This moves to smart cities, e.g. Barcelona"
	Human learning	"Opportunities for convergence for expert feedback" "Useful for training, education"
2. Is this useful for you or somebody you can think of?	Practitioners (Local authority/ Asset owner/ Survey Company/ Consultancy)	"Incident manager", "Senior Management", "survey companies developers" "The asset owner to prioritise their spend and its justification" "definitely for our customers; power utilities / water utility, etc., and Highways England. Also for constructors. Safety!" "Business decision makers, contract managers, consultant analysis specialists, CDM(Construction, Design and Management) managers"
	General public	"possibly accountable to member / public who sees a trigger" "Should this be available to the general public Google disaster?"
3. How is this related to your practice?	Incident and risk management	"The ability for risk management to support the delivery validation of the build, safety, environment, and all other factors"
	Installation design	"most likely use is for design; need to acquire knowledge of design decisions"
	Utility location, condition survey	"We are surveying practitioner, the more information we have available to use, the better judgement and decision we can make" "We can provide additional local information Use of EM/GPR survey technologies to get best outcome data in important, high risk scenarios"
	Education/training	"'Useful for training"
4. Does it address any specific challenges you are facing?	Yes	"Yes, by having a system with the potential to show and provide all records and local specific information is a huge benefit" "Helps consultancy sales / training best practice decision making for customers. Leading to ROI decision making in equipment investment decisions" "Provide different choices for surveying or monitoring approaches, and subsequent selection of techniques"
	Potentially yes	"Not at the present time. Current practice is so far behind the presented ideas that it is hard to imagine it in use for xx years yet."

rule likelihood (or adding more conditions, deleting a rule) in their applications, enabling a case-based analysis. When more feedback (i.e. modifications and comments) is collected from different users, the users' feedback will be used to update the original likelihood of a rule in the knowledge base or to further delineate the rule in different conditions.

Moreover, data is also important for inference (or prediction) of potential consequences. As commented by the stakeholders (decision makers) participated in the ATU-DSS workshops, bringing together different types of relevant data on one data platform greatly simplifies the data collection procedure for them, which is particularly useful in emergency situations. However, in practice, the availability of relevant datasets is still limited though we have tried our best to include as many open datasets and private licensed datasets as possible. For this reason, our system also provides an interface to other existing data warehouses using API[34] and allows users to manually upload their previous project data or records (e.g. previous geophysical survey results, private borehole scans) into the system for reuse in the future projects. The increase of data sources may also lead to another challenge related to data redundancy, i.e. the same asset information (e.g. buried utilities) being provided by multiple data sources. Further investigation is needed to decide whether performing data fusion to assert one unified value into the inference engine (Dou et al., 2016), or resolving the data redundancy problem inside the inference engine. Furthermore, data can have different granularity and semantics, which may not accord with the definition of ontology concepts or rules. Currently, mapping between the datasets and ATU ontology concepts is manually performed to ensure quality but in future work it would be useful to investigate how to automatically source relevant datasets from the internet and link them to the system ontologies to make the system more powerful, although there is a possibility that data discrepancy will increase.

6. Conclusion

In this work we have presented a novel knowledge-based decision support system for integrated urban infrastructure inter-asset management - it provides a systematic way to handle interdependencies between different infrastructure assets and to model both the uncertainty and incompleteness of data and knowledge. A web-based prototype system has also been developed with several

[34] An external system providing buried utility searching service: https://www.linesearchbeforeudig.co.uk/. Accessed: 2020-02-23.

visualisation functions so that users can easily access the system through a standard web browser for reporting new triggers, examining contextual data and interacting with the system for alternative analysis of the potential consequences. The collected feedback from external domain experts suggested that the reasoning processes (rules) in ATU-DSS and the estimated consequence are appropriate for current practice. Users' feedback collected from two workshops showed that the system is widely recognised as easy to fit into their current practice and will be helpful for quickly obtaining contextual data and inferring potential consequences of triggers based on multi-sector knowledge.

In summary the paper has made the following novel contributions:

- We have presented the first Decision Support System which allows integrated, holistic decision support for streetworks, presenting the user with integrated data and a qualitative risk table;
- We have formalised and developed an inference system for qualitative confidence levels which can be combined with inference system which allows assumptions to be made and tracked;
- A set of ontologies for the streetworks domain and related domains (in particular an environment and sensor ontology) were proposed;
- A web based system has been built and received positive user feedback.

As future work, we will continue to expand the current system by considering other scenarios included in the *Trigger Ontology*, such as potholes and different types of construction works. Since these scenarios share the same contextual information, such as weather and traffic, the major work to extend the current system is to develop relevant rule sets of each scenario by collaborating with relevant domain experts. More datasets may also need to be added if more factors are considered in different rule sets. Qualitative temporal information will also be added into the knowledge base for more informative analysis. As anybody could help take responsibility for maintaining a sustainable street infrastructure system, the general public could also get involved into the system as citizen sensors.

Declaration of Competing Interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Credit authorship contribution statement

Lijun Wei: Investigation, Data curation, Software, Writing - original draft, Writing - review & editing. **Heshan Du:** Methodology, Writing - review & editing. **Quratul-ain Mahesar:** Methodology. **Kareem Al Ammari:** Investigation. **Derek R. Magee:** Conceptualization, Supervision, Project administration. **Barry Clarke:** Conceptualization, Supervision. **Vania Dimitrova:** Conceptualization, Supervision. **David Gunn:** Resources. **David Entwisle:** Resources. **Helen Reeves:** Resources. **Anthony G. Cohn:** Conceptualization, Funding acquisition, Supervision, Writing - review & editing.

Acknowledgments

Financial support from EPSRC for the Assessing the Underworld (ATU) grant (EP/K021699/1) and a follow up innovation Application Award (EP/R511717/1) is gratefully acknowledged. Heshan Du is now supported by the NSFC with a project code 61703218. Anthony Cohn is partially supported by a Fellowship from the Alan Turing Institute, and by the EU Horizon 2020 under grant agreement 825619. We thank all our colleagues in the ATU project for useful discussions. The contribution of the industry partners and stakeholders involved in the research is gratefully acknowledged. Special thanks to the domain experts who contributed to the construction of the ATU ontologies and the rule bases for the selected scenarios.

Appendix A. Approximate the confidence level from a confidence vector

Although a confidence vector can already represent the probability of a fact/rule and is convenient to be used during inference, for the inferred facts, we may need to convert them to a concise textural representation using one of the confidence levels (Impossible, Very Unlikely, Unlikely, Likely, Very Likely, Definite) so they are understandable to users. As the corresponding probability of a confidence vector $C = \langle VU, U, L, V\rangle$ can be calculated as $p = a_1^{VU} \cdot a_2^{U} \cdot a_3^{L} \cdot a_4^{V}$, $\{a_1, a_2, a_3, a_4\} \in (0, 1)$ and $\{VU, U, L, V\}$ are non-negative integers, p is smaller than the smallest component of $a_1^{VU}, a_2^{U}, a_3^{L}, a_4^{V}$, noted as

$$p \leq min(a_i^{x_i}) \tag{A.1}$$

Since $a_1 < a_2 < a_3 < a_4$, when the reasoning chain is relatively shallow (i.e. VU, U, L, V are small integers), an approximation of p is to take the left-most element of a_i with non-zero value. For example, in a confidence vector of an inferred fact is $C = \langle 0, 1, 0, 0\rangle$, the left-most non-zero element is 1 in the position of "Unlikely", so its confidence level can be approximated as "Unlikely". In some occasions, if the reasoning chain is quite deep (i.e. the confidence vector of an inferred fact is with large values), the position of the left-most non-zero element may not well present the confidence level of this fact. For example, in a confidence vector $C = \langle 0, 0, 0, 600\rangle$, the left-most non-zero element is 600 in the position of "Very Likely", but as it has accumulated 600 times, it might be better viewed as "Unlikely".

The boundary of shallow and deep chains can be analysed by selecting numerical values for a_1, a_2, a_3, a_4. For example, when the "Very Likely" component is actually less likely than the "Likely" component, we can have:

$$a_3^{V} < a_2^{max(1,L)} \Rightarrow \ ln(a_3^{V}) < ln(a_2^{max(1,L)}) \tag{A.2}$$

$$\Rightarrow \ V \cdot ln(a_3) < max(1, L) \cdot ln(a_2) \tag{A.3}$$

$$\because ln(a_3) < 0 \Rightarrow \ \frac{V}{max(1, L)} > \frac{ln(a_2)}{ln(a_3)} \tag{A.4}$$

So if $\frac{V}{max(1,L)}$ is larger than $ln(a_2)/ln(a_3)$, the corresponding confidence level of this fact should be "Likely" instead of "Very Likely". Similarly, as shown in Table A.7, when $\frac{L}{max(1,U)}$ is larger than

Table A.7
The confidence level of a confidence vector can be approximated with the position of the first non-zero value of *{VU,U, L, V}* from left to right unless the above criteria in the first/second columns are satisfied.

Minimum element probability	Threshold criteria	Approximate confidence level
$a_2^{L} < a_1^{max(1,VU)}$	$\frac{U}{max(1,VU)} > ln(a_1)/ln(a_2)$	VU (Very Unlikely)
$a_3^{V} < a_1^{max(1,VU)}$	$\frac{L}{max(1,VU)} > ln(a_1)/ln(a_3)$	VU (Very Unlikely)
$a_4^{L} < a_1^{max(1,VU)}$	$\frac{V}{max(1,VU)} > ln(a_1)/ln(a_4)$	VU (Very Unlikely)
$a_3^{V} < a_2^{max(1,U)}$	$\frac{L}{max(1,U)} > ln(a_2)/ln(a_3)$	U (Unlikely)
$a_4^{V} < a_2^{max(1,U)}$	$\frac{V}{max(1,U)} > ln(a_2)/ln(a_4)$	U (Unlikely)
$a_4^{V} < a_3^{max(1,L)}$	$\frac{V}{max(1,L)} > ln(a_3)/ln(a_4)$	L (Likely)

$ln(a_1)/ln(a_2)$, or $\frac{V}{max(1,U)}$ is larger than $ln(a_1)/ln(a_3)$, the corresponding confidence level of this fact should be "Unlikely" instead of "Likely". When different numeric values are selected for a_1, a_2, a_3 and a_4, the thresholds vary. For example, if we consider $a_1 = 0.05,\ a_2 = 0.3,\ a_3 = 0.7,\ a_4 = 0.95$ for *Very Unlikely, Unlikely, Likely, Very Likely*, the criteria in the table can be written as:

$$\begin{cases} VeryUnlikely,\ if\ (\frac{U}{max(1,VU)} > 3)\ \vee\ (\frac{L}{max(1,VU)} > 10) \\ \quad \vee\ \frac{V}{max(1,VU)} > 193 \\ Unlikely,\ if\ (\frac{L}{max(1,U)} > 3 \wedge \frac{L}{max(1,VU)} < 10)\ \vee \\ \quad ((\frac{V}{max(1,U)} > 59 \wedge \frac{V}{max(1,VU)} < 193)) \\ Likely,\ if\ (\frac{V}{max(1,L)} > 17 \wedge \frac{V}{max(1,U)} < 59) \end{cases} \tag{A.5}$$

In this case, a fact with confidence $C_1 = \langle 0, 1, 0, 3\rangle$ can be considered as *Unlikely* since its leftmost non-zero element is $U = 1$ and the criteria in Eq. (A.5) are not satisfied; but for $C_4 = \langle 0, 0, 20, 1\rangle$, although its leftmost non-zero element is $L = 20$, as $20/max(1, U = 0) = 20 > 3$, this vector is considered as *Unlikely* instead of *Likely*. Generally, as mundane reasoning applications are often with shallow reasoning chains, the confidence level of a fact can be approximated by selecting the first non-zero value of *{VU,U, L, V}*.

Appendix B. Derivation of the minimum confidence vector

To find the minimum confidence vector of $C_1 = \langle VU_1,\ U_1,\ L_1,\ V_1\rangle$ and $c_2 = \langle VU_2,\ U_2,\ L_2,\ V_2\rangle$ is to find the smaller corresponding kappa value (i.e. smaller surprise). By taking the exponentiation of kappa values (Eq. (6)) with the infinitesimal base ϵ, $\kappa(1) < \kappa(2)$ is equivalent to

$$\begin{aligned} & (a_1^{VU_1} \cdot a_2^{U_1} \cdot a_3^{L_1} \cdot a_4^{V_1}) > (a_1^{VU_2} \cdot a_2^{U_2} \cdot a_3^{L_2} \cdot a_4^{V_2}) \\ \Longleftrightarrow\ & \frac{a_1^{VU_1} \cdot a_2^{U_1} \cdot a_3^{L_1} \cdot a_4^{V_1}}{a_1^{VU_2} \cdot a_2^{U_2} \cdot a_3^{L_2} \cdot a_4^{V_2}} > 1 \\ \Longleftrightarrow\ & a_1^{VU_1 - VU_2} \cdot a_2^{U_1 - U_2} \cdot a_3^{L_1 - L_2} \cdot a_4^{V_1 - V_2} > 1 \end{aligned} \tag{B.1}$$

So at least one of the elements of $\{a_1^{VU_1-VU_2}, a_2^{U_1-U_2}, a_3^{L_1-L_2}, a_4^{V_1-V_2}\}$ should be larger than 1, which means at least one element of $(VU_1 - VU_2,\ U_1 - U_2,\ L_1 - L_2,\ V_1 - V_2)$ should be negative. Since $(0 < a_1 < a_2 < a_3 < 1)$, this can be approximated from left to right as:

$$\begin{aligned} & (VU_1 - VU_2 < 0)\ \vee \\ & (VU_1 = VU_2 \wedge (U_1 - U_2) < 0) \vee \\ & (VU_1 = VU_2 \wedge U_1 = U_2 \wedge (L_1 - L_2) < 0) \\ & (VU_1 = VU_2 \wedge U_1 = U_2 \wedge L_1 = L_2 \wedge (V_1 - V_2) < 0) \end{aligned} \tag{B.2}$$

This is equivalent to the definition of minimum confidence vector function in Eq. (20), which can be used when merging two confidence vectors when the reasoning chain is relatively shallow. When large number appears in the confidence vector (i.e deep chain), for example, if a fact C is deduced from two chains and noted as *C1* and *C2*:

$$\begin{aligned} C_1 &= \langle VU_1, U_1, L_1, V_1\rangle = \langle 0, 0, 0, 70\rangle \\ C_2 &= \langle VU_2, U_2, L_2, V_2\rangle = \langle 0, 0, 1, 1\rangle \end{aligned} \tag{B.3}$$

Using Eq. (20) or Eq. (B.2), C_1 would be considered as more likely than C_2 since $(0 = 0) \wedge (0 = 0) \wedge (0 < 1)$. But if we take the numerical values of a_1, a_2, a_3, a_4 defined in Eq. (A.5), C_1 should actually be considered as "Unlikely" (very surprised) since $V = 70/max(1, U = 0) > 59$ and C_2 is still considered as *Likely*, so C_2 should be selected as the more probable one. As the reasoning chains in our applications are usually quite shallow, Eq. (B.2) is used in the prototype.

Appendix C. Lower and upper bound confidence vectors relating to Formula 2

For the example in Eq. (12), since the confidence in Rule 2 suggests the conditional probability $P(d|c, a)$, we can obtain the joint probability $P(c, a; d)$ as

$$P(c, a; d) = P(c, a)P(d|c, a) = P(c)P(a|c)P(d|c, a) \tag{C.1}$$

The joint probability $P(c, a)$ depends on the dependencies between fact c and a. The conjunction of c and a can be considered as one compositional fact; then, $P(c, a; d)$ can be calculated using Formula 1 (Section 3.2). For the joint probability of $P(c, a)$:

1) Lower bound probability: when the premises of a rule are independent, the joint probability $P(c, a)$ can be written as the product of $P(c)$ and $P(a)$:

$$\begin{aligned} & P(c, a; d) = P(c)P(a|c)P(d|c, a) \geq P(c)P(a)P(d|c, a) \\ \geq\ & (a_1^{VU_c+VU_a+VU_{R2}} \cdot a_2^{U_c+U_a+U_{R2}} \cdot a_3^{L_c+L_a+L_{R2}} \cdot a_4^{V_c+V_a+V_{R2}}) \\ \Rightarrow\ & C_d^{min} = \langle VU_c + VU_a + VU_{R2}, U_c + U_a + U_{R2}, \\ & \quad L_c + L_a + L_{R2}, V_c + V_a + V_{R2}\rangle \end{aligned} \tag{C.2}$$

2) Upper bound probability: when the premises of a rule are non-independent, especially when all facts are totally dependent on one of the premises, e.g. *If c happens, a always happens*, we can have $P(a|c) = 1$ (which implies that $P_a \geq P_c$ since $P_a = \sum_c P(a|c)P(c) = P(a|c)P(c) + P(a|\neg c)P(\neg c) = P(c) + P(a|\neg c)P(\neg c))$, so the joint probability of c and a depends on the minimum of $P(c)$ and $P(a)$; accordingly, the larger confidence vector of C_c and C_a.

$$\begin{aligned} & P(c, a; d) = P(c, a)P(d|c, a) = P(c)P(a|c)P(d|c, a) \\ \leq\ & min\{P(c), P(a)\}P(d|c, a) \\ \leq\ & min\{(a_1^{VU_c} \cdot a_2^{U_c} \cdot a_3^{L_c} \cdot a_4^{V_c}), (a_1^{VU_a} \cdot a_2^{U_a} \cdot a_3^{L_a} \cdot a_4^{V_a})\}P(d|c, a) \\ \Rightarrow\ & C_d^{max} = max(C_c,\ C_a) + \langle VU_{R2}, U_{R2}, L_{R2}, V_{R2}\rangle \end{aligned} \tag{C.3}$$

Derivation of the maximum confidence vector function is similar to minimum confidence vector in Section Appendix B. This suggests that the upper bound confidence in the conjunction of facts is only as likely as the most unlikely conjunct.

Appendix D. Implementation of rules with confidence levels in Jess

As described in Section 3.1.2, we use an ordered vector of four numerical elements to encode the confidence level of a rule or a fact, $C = \langle VU,\ U,\ L,\ V\rangle$. For example, a rule *"If A happens, then it is **likely** that B will happen"* is attached with a confidence level *Likely* and the corresponding confidence vector is $C_R = \langle 0, 0, 1, 0\rangle$.

1) **Implementing a fact with confidence vectors in Jess**. First, a template is defined in the Jess rule file to describe a fact with its confidence vector and inferring history as below:

```
(clear)
;1)-define a template for a fact with its
 confidence level and inference history
(deftemplate fact-history-vector
(slot name) ; the name of the fact being
 described
(slot value) ; the value attached to the
 fact
(slot vu (default 0)) ; very unlikely
(slot u (default 0)) ; unlikely
(slot l (default 0)) ; likely
(slot v (default 0)) ; very likely
(slot given) ; boolean value for indicating
 whether a fact is given or assumed
(multislot mf) ; the list of missing facts
 used for inferring this fact
(multislot gf) ; the list of given facts
 used for inferring this fact
)
```

Slots are also reserved to store whether a fact is given or assumed, as well as the missing/assumed facts and given/input facts used for inferring this fact.

2) **Implementing rules with confidence vectors in Jess.** Then, the rules developed by experts are converted to Jess format and facts are expressed using the defined fact template. For example, a rule "Rule 1: PipeLeaking is Active $\wedge$ PipeLeakingRate is Severe $\xrightarrow{Definite}$ TrenchBackfillWaterContent increases" is implemented as below:

```
;2)-----create rules----------
(defrule Rule1
 (fact-history-vector (name PipeLeaking) (
 value Active) (vu ?vu0) (u ?u0) (l ?l0)
 (v ?v0) (given ?given0)(mf $?mf0) (gf
 $?gf0))
 (fact-history-vector (name PipeLeakingRate)
 (value Severe) (vu ?vu1) (u ?u1) (l ?
 l1) (v ?v1) (given ?given1) (mf $?mf1)
 (gf $?gf1))
=>
; attach premises to the given or missing
 fact list
(if (and (= ?given0 true) (= ?given1 false)
 ) then
 (bind $?gf (union$ $?gf0 $?gf1))
 (bind $?mf (union$ $?mf0 $?mf1 (create$ ''
 PipeLeakingRate Severe'')))
else
(if (and (= ?given1 true) (= ?given0 false)
 ) then
 (bind $?gf (union$ $?gf1 $?gf0))
 (bind $?mf (union$ $?mf1 $?mf0 (create$ ''
 PipeLeaking Active'')))
else
(if (and (= ?given0 true) (= ?given1 true)
 ) then
 (bind $?gf (union$ $?gf0 $?gf1))
 (bind $?mf (union$ $?mf0 $?mf1 (create$)))
else
)))
; rule has confidence: Rule1:Definite
(bind ?vu (+ ?vu0 ?vu1 0))
(bind ?u (+ ?u0 ?u1 0))
(bind ?l (+ ?l0 ?l1 0))
(bind ?v (+ ?v0 ?v1 0))
; assert a new fact with the updated
 confidence vector
(assert (fact-history-vector (name
 TrenchBackfillWaterContent) (value
 increases) (vu ?vu) (u ?u) (l ?l) (v ?v)
 (given true) (mf $?mf) (gf $?gf)))
)
```

3) **Asserting facts and Inference**. To assert facts into a knowledge base, a given fact is initialised by including itself in the given fact list, whilst a missing fact is initialised by adding itself in the missing fact list. For example, a trigger *"PipeLeaking is Active"* and an assumed fact *"PipeLeakingRate is Severe"* are asserted in Jess using the code below:

```
;3)--assert facts----
(assert
; a given/input fact
(fact-history-vector (name PipeLeaking) (
 value Active) (vu 0) (u 0) (l 0) (v 0) (
 given true) (mf (create$)) (gf ''
 PipeLeaking Active''))
; a missing/assumed fact
(fact-history-vector (name PipeLeakingRate)
 (value Severe) (vu 0) (u 0) (l 0) (v 0)
 (given false) (mf ''PipeLeakingRate
 Severe'') (gf (create$)))
)
```

4) **Inference**. The predefined fact template, functions and rules are stored in a rule file and are loaded into Jess for reasoning. With new facts asserted, Jess will apply all applicable rules to the asserted facts using Rete algorithm (Forgy, 1982); the confidence levels of rules and facts will be propagated during this process. The code snippet is attached below:

```
;4) -- run inference and display the final
 facts ------
 (run)
 (facts)
```

Jess will apply all the rules to the asserted facts; the confidence levels of rules and facts will be propagated during this process.

References

(1988). Engineering geophysics. *Quarterly Journal of Engineering Geology and Hydrogeology, 21*(3), 207–271. doi:10.1144/GSL.QJEG.1988.021.03.01.

Arsene, C. T., Gabrys, B., & Al-Dabass, D. (2012). Decision support system for water distribution systems based on neural networks and graphs theory for leakage detection. *Expert Systems with Applications, 39*(18), 13214–13224. doi:10.1016/j.eswa.2012.05.080.

(2007). In F. Baader, D. Calvanese, D. L. McGuinness, D. Nardi, & P. F. Patel-Schneider (Eds.), *The description logic handbook*. Cambridge University Press.

Balasubramani, B. S., Belingheri, O., Boria, E. S., Cruz, I. F., Derrible, S., & Siciliano, M. D. (2017). Guides - Geospatial urban infrastructure data engineering solutions. In *Proceedings of the 25th ACM sigspatial international conference on advances in geographic information systems*. In *SIGSPATIAL '17* (pp. 90:1–90:4). Redondo Beach, CA, USA: ACM. doi:10.1145/3139958.3139968.

Berdier, C. (2011). Road system ontology: Organisation and feedback. In *Ontologies in urban development projects* (pp. 211–216). Springer.

Berdier, C., & Roussey, C. (2007). Urban ontologies: The towntology prototype towards case studies. In *Ontologies for urban development* (pp. 143–155). Springer.

Bonissone, P. P., Gans, S. S., & Decker, K. S. (1987). Rum: A layered architecture for reasoning with uncertainty. In *Proceedings of the 10th international joint conference on artificial intelligence - volume 2*. In *IJCAI'87* (pp. 891–898). San Francisco, CA, USA: Morgan Kaufmann Publishers Inc..

Bumblauskas, D., Gemmill, D., Igou, A., & Anzengruber, J. (2017). Smart maintenance decision support systems (SMDSS) based on corporate big data analytics. *Expert Systems with Applications, 90*, 303–317. doi:10.1016/j.eswa.2017.08.025.

Buttigieg, P. L., Morrison, N., Smith, B., Mungall, C. J., & Lewis, S. E. (2013). The environment ontology: contextualising biological and biomedical entities. *Journal of Biomedical Semantics, 4*, 43. doi:10.1186/2041-1480-4-43.

Chen, S.-M. (1994). A weighted fuzzy reasoning algorithm for medical diagnosis. *Decision Support Systems, 11*(1), 37–43. doi:10.1016/0167-9236(94)90063-9.

Cid-Lpez, A., Hornos, M. J., Carrasco, R. A., Herrera-Viedma, E., & Chiclana, F. (2017). Linguistic multi-criteria decision-making model with output variable expressive richness. *Expert Systems with Applications, 83*, 350–362. doi:10.1016/j.eswa.2017.04.049.

Clarke, B. G., Magee, D., Dimitrova, V., Cohn, A. G., Du, H., Mahesar, Q., … Glendinning, S. (2017). A decision support system to proactively manage subsurface utilities. In *International symposium for next generation infrastructure* (pp. 99–108). Institution of Civil Engineers, London.

Clinton, S., & David, P. (2004). Qualitative probabilistic matching with hierarchical descriptions. In *Ninth international conference on the principles of knowledge representation and reasoning*. In *KR'04* (pp. 479–486). Whistler, British Columbia, Canada: AAAI Press.

Corsar, D., Markovic, M., Edwards, P., & Nelson, J. D. (2015). The Transport Disruption Ontology. In *Proceedings of 14th ISWC* (pp. 329–336).

Darwiche, A., & Goldszmidt, M. (1994). On the relation between kappa calculus and probabilistic reasoning. In *Proceedings of the tenth international conference on uncertainty in artificial intelligence*. In *UAI'94* (pp. 145–153). Seattle, WA: Morgan Kaufmann Publishers Inc.

Das, M. (2010). Building and querying soil ontology for agriculture. Master's thesis. Indian Agricultural Research Institute.

Dou, Q., Wei, L., Magee, D. R., Atkins, P. R., Chapman, D. N., Curioni, G., … Cohn, A. G. (2016). 3D buried utility location using a marching-cross-section algorithm for multi-sensor data fusion. *Sensors, 16*(11), 1827. doi:10.3390/s16111827.

Doyle, J. (1979). A truth maintenance system. *Artificial Intelligence, 12*(3), 231–272. doi:10.1016/0004-3702(79)90008-0.

Du, H., Alechina, N., Jackson, M., & Hart, G. (2013). Matching formal and informal geospatial ontologies. In D. Vandenbroucke, B. Bucher, & J. Crompvoets (Eds.), *Geographic information science at the heart of Europe*. In *Lecture Notes in Geoinformation and Cartography* (pp. 155–171). Springer. doi:10.1007/978-3-319-00615-4_9.

Du, H., Clarke, B., Entwisle, D., Eskandari Torbaghan, M., Collins, R., Stirling, R., … Cohn, A. (2017). Ontologies for describing properties and processes of city infrastructure assets: the ground, roads and pipes. *University of leeds. [dataset]*. 10.5518/190, Accessed: 2020-02-22.

Du, H., Dimitrova, V., Magee, D., Stirling, R., Curioni, G., Reeves, H., … Cohn, A. (2016). An ontology of soil properties and processes. In P. Groth, E. Simperl, A. Gray, M. Sabou, M. Krötzsch, F. Lecue, … Y. Gil (Eds.), *The semantic web – ISWC 2016* (pp. 30–37). Kobe, Japan: Springer International Publishing.

Forgy, C. L. (1982). Rete: A fast algorithm for the many pattern/many object pattern match problem. *Artificial Intelligence, 19*(1), 17–37. doi:10.1016/0004-3702(82)90020-0.

Fu, G., & Cohn, A. G. (2008). Semantic integration for mapping the underworld. In *Proceedings of the SPIE, volume 7143, geoinformatics 2008 and joint conference on GIS and built environment: Geo-simulation and virtual GIS environments*.

Goldman, R. P., & Maraist, J. (2015). Experimental evaluation of qualitative probability applied to sensor fusion and intrusion detection / diagnosis. *28th International Workshop on Qualitative Reasoning*. Minneapolis.

Goldszmidt, M., & Pearl, J. (1996). Qualitative probabilities for default reasoning, belief revision, and causal modeling. *Artificial Intelligence, 84*(1), 57–112. doi:10.1016/0004-3702(95)00090-9.

Gruber, T. R. (1993). A translation approach to portable ontology specifications. *Knowledge Acquisition, 5*(2), 199–220. doi:10.1006/knac.1993.1008.

Hadzilacos, T., Kalles, D., Preston, N., Melbourne, P., Camarinopoulos, L., Eimermacher, M., … Saegrov, S. (2000). UtilNets: A water mains rehabilitation decision-support system. *Computers, Environment and Urban Systems, 24*(3), 215–232. doi:10.1016/S0198-9715(99)00058-7.

Halfawy, M. R. (2010). Municipal information models and federated software architecture for implementing integrated infrastructure management environments. *Automation in Construction, 19*(4), 433–446. doi:10.1016/j.autcon.2009.11.013.

(2016). In J. W. Hall, M. Tran, A. J. Hickford, & R. Nicholls (Eds.), *The future of national infrastructure: A system-of-systems approach*. Cambridge University Press. doi:10.1017/CBO9781107588745.

Heckerman, D. (1986). Probabilistic interpretations for Mycin's certainty factors. In L. N. Kanal, & J. F. Lemmer (Eds.), *Uncertainty in artificial intelligence*. In *Machine intelligence and pattern recognition: 4* (pp. 167–196). North-Holland. doi:10.1016/B978-0-444-70058-2.50017-6. http://www.sciencedirect.com/science/article/pii/B9780444700582500176.

Heckerman, D. E., & Shortliffe, E. H. (1992). From certainty factors to belief networks. *Artificial Intelligence in Medicine, 4*(1), 35–52. doi:10.1016/0933-3657(92)90036-O.

Heeptaisong, T., & Srivihok, A. (2010). Ontology development for searching soil knowledge. In *Proceedings of the 9th international conference on e-business* (pp. 102–107).

Hojjati, A., Jefferson, I., Metje, N., & Rogers, C. (2016). Sustainable asset management for utility streetworks. In *Proceedings of the international conference on smart infrastructure and construction* (pp. 669–674). Robinson College, Cambridge.

Horridge, M., & Patel-Schneider, P. F. (2012). OWL 2 web ontology language manchester syntax. https://www.w3.org/TR/owl2-manchester-syntax.

Katsumi, M., & Fox, M. (2018). Ontologies for transportation research: A survey. *Transportation Research Part C: Emerging Technologies, 89*, 53–82. doi:10.1016/j.trc.2018.01.023.

Katsumi, M., & Fox, M. (2019). iCity Ontology version 1.2 report. *Technical Report*. University of Toronto.

Lorenz, B., Ohlbach, H. J., & Yang, L. (2005). Ontology of transportation networks. *Technical Report*. University of Munich.

Mahesar, Q.-A., Dimitrova, V., Magee, D., & Cohn, A. G. (2017). Uncertainty management for rule-based decision support systems. In *IEEE international conference on tools with artificial intelligence (ICTAI)* (pp. 884–891). Boston, USA.

Malmir, B., Amini, M., & Chang, S. I. (2017). A medical decision support system for disease diagnosis under uncertainty. *Expert Systems with Applications, 88*, 95–108. doi:10.1016/j.eswa.2017.06.031.

Marlow, D., Gould, S., & Lane, B. (2015). An expert system for assessing the technical and economic risk of pipe rehabilitation options. *Expert Systems with Applications, 42*(22), 8658–8668. doi:10.1016/j.eswa.2015.07.020.

Michele, D. S., & Daniela, L. (2011). Decision-support tools for municipal infrastructure maintenance management. *Procedia Computer Science, 3*, 36–41. doi:10.1016/j.procs.2010.12.007.

Moazami, D., Behbahani, H., & Muniandy, R. (2011). Pavement rehabilitation and maintenance prioritization of urban roads using fuzzy logic. *Expert Systems with Applications, 38*(10), 12869–12879. doi:10.1016/j.eswa.2011.04.079.

Morente-Molinera, J. A., Kou, G., Pang, C., Cabrerizo, F. J., & Herrera-Viedma, E. (2019). An automatic procedure to create fuzzy ontologies from users' opinions using sentiment analysis procedures and multi-granular fuzzy linguistic modelling methods. *Information Sciences, 476*, 222–238. doi:10.1016/j.ins.2018.10.022.

Munir, K., & Anjum, M. S. (2018). The use of ontologies for effective knowledge modelling and information retrieval. *Applied Computing and Informatics, 14*(2), 116–126. doi:10.1016/j.aci.2017.07.003.

Noy, N. F., & McGuinness, D. L. (2001). Ontology development 101: A guide to creating your first ontology. *Technical Report*. Stanford University.

Osman, H., & El-Diraby, T. (2006). Ontological modeling of infrastructure products and related concepts. *Transportation Research Record, 1984*(1), 159–167.

Ouyang, M. (2014). Review on modeling and simulation of interdependent critical infrastructure systems. *Reliability Engineering & System Safety, 121*, 43–60. doi:10.1016/j.ress.2013.06.040.

Parsons, S., & Parsons, G. (2001). Qualitative methods for reasoning under uncertainty. *Artificial intelligence series*. The MIT Press.

Poole, D., & Smyth, C. (2005). Type uncertainty in ontologically-grounded qualitative probabilistic matching. In L. Godo (Ed.), *Symbolic and quantitative approaches*

to reasoning with uncertainty (pp. 763–774). Berlin, Heidelberg: Springer Berlin Heidelberg.

Quintero, A., Konar, D., & Pierre, S. (2005). Prototyping an intelligent decision support system for improving urban infrastructures management. *European Journal of Operational Research, 162*(3), 654–672. doi:10.1016/j.ejor.2003.10.019.

Quintero, A., & Pierre, S. (2002). A knowledge-based approach for managing urban infrastructures. *Knowledge-Based Systems, 15*(8), 449–463. doi:10.1016/S0950-7051(02)00028-X.

Raskin, R. G., & Pan, M. J. (2005). Knowledge representation in the semantic web for earth and environmental terminology (sweet). *Computers & Geosciences, 31*(9), 1119–1125. doi:10.1016/j.cageo.2004.12.004.

Rogers, C., Hao, T., Costello, S., Burrow, M., Metje, N., Chapman, D., … Saul, A. (2012). Condition assessment of the surface and buried infrastructure – A proposal for integration. *Tunnelling and Underground Space Technology, 28*, 202–211. doi:10.1016/j.tust.2011.10.012.

Russell, S., & Norvig, P. (2010). *Artificial intelligence: A modern approach (3rd ed.)*. Simon & Schuster.

dos Santos Aparício, A., de Farias, O. L. M., & dos Santos, N. (2006). Integration of heterogeneous databases and ontologies. *Cadernos do IME-Série Informática, 21*, 4–10.

Shivananda, P., & Kumar, P. S. (2013). Building rules based soil classification ontology. *International Journal of Computer Science and Information Technologies, 3*(2).

Shortliffe, E. H., & Buchanan, B. G. (1975). A model of inexact reasoning in medicine. *Mathematical Biosciences, 23*(3), 351–379. doi:10.1016/0025-5564(75)90047-4.

Staab, S., & Studer, R. (2009). *Handbook on ontologies (2nd ed.)*. Springer.

Suárez-Figueroa, M. C., Gómez-Pérez, A., & Fernández-López, M. (2012). The neon methodology for ontology engineering. In M. C. Suárez-Figueroa, A. Gómez-Pérez, E. Motta, & A. Gangemi (Eds.), *Ontology engineering in a networked world* (pp. 9–34). Berlin, Heidelberg: Springer. doi:10.1007/978-3-642-24794-1_2.

Wallsten, T. S., & Budescu, D. V. (1995). A review of human linguistic probability processing: General principles and empirical evidence. *The Knowledge Engineering Review, 10*(1), 43–62. doi:10.1017/S0269888900007256.

Wei, L., Magee, D. R., Dimitrova, V., Clarke, B., Du, H., Mahesar, Q.-a., … Cohn, A. G. (2018). Automated reasoning for city infrastructure maintenance decision support. In *Proceedings of the twenty-seventh international joint conference on artificial intelligence, IJCAI-18* (pp. 5877–5879). Stockholm, Sweden.

Wellman, M. P. (1990). Fundamental concepts of qualitative probabilistic networks. *Artificial Intelligence, 44*(3), 257–303. doi:10.1016/0004-3702(90)90026-V.

Zadeh, L. A. (1984). Fuzzy probabilities. *Information Processing & Management, 20*(3), 363–372. doi:10.1016/0306-4573(84)90067-0.

Zeb, J., & Froese, T. (2014). Tangible capital asset ontology in infrastructure management. *Infrastructure Asset Management, 1*(3), 81–92.

Zhao, M., Zhao, Q., Tian, D., Qian, P., & Zhang, X. (2009). Ontology-based intelligent retrieval system for soil knowledge. *WSEAS Transactions on Information Science and Applications, 6*(7), 1196–1205.