

Automated Reasoning for City Infrastructure Maintenance Decision Support

Lijun Wei¹, Derek R. Magee¹, Vania Dimitrova¹, Barry Clarke²,
 Heshan Du^{1,3}, Quratul-ain Mahesar^{1*}, Kareem Al Ammari¹ and Anthony G. Cohn¹

¹ School of Computing, University of Leeds, United Kingdom

² School of Civil Engineering, University of Leeds, United Kingdom

³ School of Computer Science, University of Nottingham Ningbo China

{L.J.Wei, D.R.Magee, V.G.Dimitrova, B.G.Clarke, A.G.Cohn, K.AlAmmari}@leeds.ac.uk

Abstract

We present an interactive decision support system for assisting city infrastructure inter-asset management. It combines real-time site specific data retrieval, a knowledge base co-created with domain experts and an inference engine capable of predicting potential consequences and risks resulting from the available data and knowledge. The system can give explanations of each consequence, cope with incomplete and uncertain data by making assumptions about what might be the worst case scenario, and making suggestions for further investigation. This demo presents multiple real-world scenarios, and demonstrates how modifying assumptions (parameter values) can lead to different consequences.

1 Introduction

Disasters and damage caused by ageing infrastructure and construction works have been increasingly reported in the news, e.g., giant sinkholes in Rome [AbcNews, 2018], Manchester [Glendinning, 2015] and Tokyo [McCurry, 2016] due to underground leaking pipes/construction. These problems are likely to increase as historical assets design lifetimes are exceeded, and increasing urbanisation creates ever more destructive street works. Complex decision making in city infrastructure management is a challenging task [Rogers *et al.*, 2012] that requires different pieces of information which are time consuming to collect and requires comprehensive multi-sector knowledge for data interpretation and risk estimation.

In this work, we present an interactive decision support system developed within the UK EPSRC funded Assessing the Underworld (ATU) Project¹, exploiting the application of knowledge representation and automated reasoning techniques in urban infrastructure inter-asset management and maintenance, a domain with high economic and societal importance. The system (referred to as **ATU-DSS** hereafter)

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

¹ <http://assessingtheunderworld.org/>

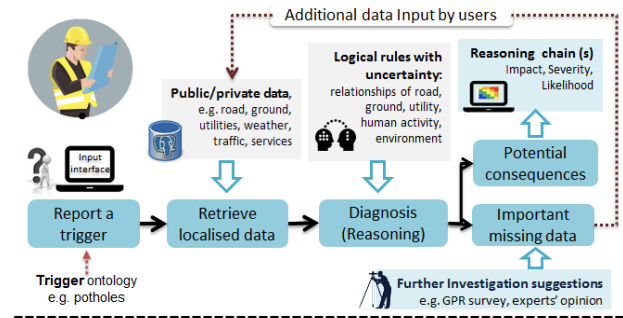


Figure 1: System architecture

can help decision makers (e.g., incident managers, contractors, local authority) answer the questions like [Clarke *et al.*, 2017]: *What is the current condition* of road surface, adjacent pipes and ground at a specific site? *How will this condition change* because of extreme weather, deterioration of assets, and planned streetworks? *Will this change cause any undesirable consequences* (e.g. traffic disruption, loss of services)? ATU-DSS can be used in a wide range of contexts (as suggested by the industry partners), such as prioritisation of asset maintenance, preparation of health/safety files for construction works, or training junior engineers to cope with complexed decision making that requires multi-sector knowledge.

2 ATU-DSS Architecture

The ATU-DSS prototype is a web-based system composed of three components: 1) A *knowledge model* captures the multi-sector knowledge held by different technical experts (e.g. civil engineers, utility management experts) in a computer-processable format, including a family of city infrastructure ontologies capturing the main concepts in roads, buried assets and ground [Du *et al.*, 2016; Du *et al.*, 2017], a trigger ontology specifying the events that require decisions to be made in infrastructure asset management (e.g. planned excavation/reinstatement works, road cracks/pothole, pipe leakage) [Clarke *et al.*, 2017], and an investigation ontology encoding the available geophysical techniques for measuring different asset properties in shallow (0-5m depth) streetworks surveys. These ontologies provide a common vocabulary for data

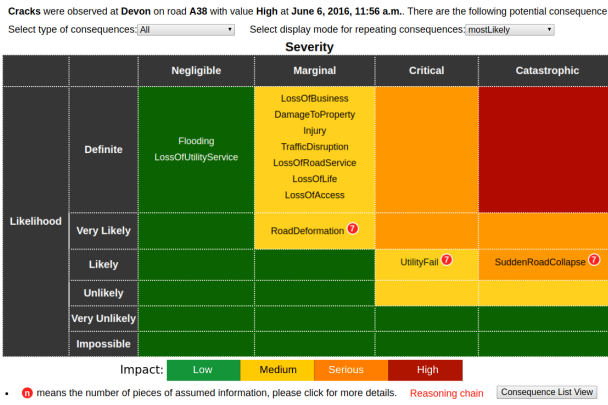


Figure 2: The impact matrix view of potential consequences (N.B. colours in the table are assigned based on generic heuristics).

integration and rule development. 2) An integrated database stores the relevant real-world data of infrastructure assets and contextual information sourced from different data owners, such as the historic weather data sourced from the British Met Office and the ground condition datasets sourced from the British Geological Survey; 3) An inference engine links the knowledge base and the real world data, enabling an integrated approach for inference of the potential consequences of a reported trigger. The work flow (Figure 1) of the ATU-DSS starts with a report of a trigger (e.g. pothole), followed by real-time retrieval of the localised contextual data based on the location and time of the trigger using different spatial criteria. E.g., weather data is interpolated from all nearby weather stations within 10km to the trigger; traffic information is inferred from the traffic counting point on the nearest road segment. All the retrieved data is fed into the inference engine for automated reasoning of the potential consequences.

3 Automated Reasoning and Explanation

The inference engine links the data with a rule base for logical inference [Poole and Mackworth, 2017]. In the ATU-DSS, predictive rules (“If Cause, then Effect”) are defined by domain experts employing the common vocabulary defined in the ATU ontologies; diagnosis rules (“If Effect, then Possible Cause”) are then automatically generated. These rules encode the complex intra-asset and inter-asset relationships among different infrastructure assets and other important factors (e.g. weather and human activities). E.g., a leaking water pipe is likely to change the condition of its surrounding ground:

$$\text{PipeLeaking is Active} + \text{PipeLeakingRate is Low} \xrightarrow{\text{Likely}} \text{TrenchBackfillWaterContent increases}$$

In practice, these rules are not always certain but require a hedging with a confidence level. Since the experts we consulted found it difficult to formulate rules with numeric probabilities [Heckerman and Shortliffe, 1992], qualitative confidences were assigned to the predictive rules by domain experts by selecting from {Impossible, Very Unlikely, Unlikely, Likely, Very Likely, Definite}, which can be written in the form of a confidence vector $\langle VU, U, L, V \rangle$ (extended from [Mahesar et al., 2017]). *D* (Definite) is implicit in the definition when

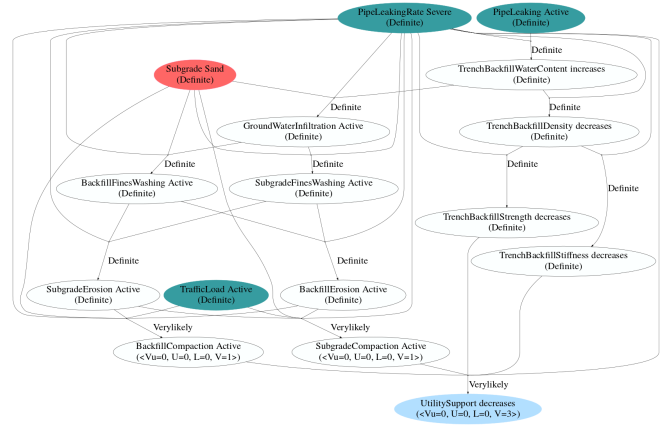


Figure 3: Example of a simple reasoning chain. Colour scheme: turquoise for given facts, pink for assumed facts (missing facts), white for intermediately inferred facts and light blue for the final output.

VU, U, L, V are all zero; and I (Impossible) is implicit when any of the vector elements is a very large positive integer (e.g. $VU = 1000$ is assigned to a rule marked as *Impossible* in this application). For diagnosis rules, an extra element A is attached to the rules to record the accumulated steps of abductive reasoning. Qualitative confidences are also attached to human supplied data. The inference engine is able to propagate the qualitative confidences of data and rules in the reasoning chain. Our inference method also provides a mechanism to handle incomplete data, i.e. if any premises (facts) of an inference rule are missing², the rule can still be fired with an assumed value (effectively assuming the worst case scenario, since the rules are formulated so as to predict problems that may occur).

The rule engine Jess³ is used for rule implementation and reasoning using an improved Rete algorithm [Forgy, 1982] so the performance is largely independent of the number of rules/facts. The inference engine can determine the potential consequences on assets (e.g. road collapse) and the societal and economic consequences (e.g., traffic disruption), determine whether additional information is needed, and suggest ways of obtaining such information (e.g. geophysical surveys). The identified potential consequences are shown to users in an impact matrix table according to their likelihood and severity levels (Figure 2). The system can also give explanations of each consequence by providing text description or displaying the reasoning chain (Figure 3) which can help users make more reasonable and informed decisions. For the missing data with assumed values, users can add/update the missing data and the whole reasoning process automatically re-activates. A video⁴ shows the complete workflow.

Acknowledgments

Financial support from EPSRC Assessing the Underworld grant (EP/K021699/1) is gratefully acknowledged. We thank all our project partners and stakeholders for useful discussions.

²At least one premise of a rule must be not assumed.

³<http://www.jessrules.com/jess/docs/71/>

⁴<http://bit.ly/2pV7k3Z>

References

- [AbcNews, 2018] AbcNews. Rome sinkhole swallows several cars in residential street. <https://ab.co/2GHHSJD>, 2018. Accessed: 2018-05-04.
- [Clarke *et al.*, 2017] Barry G. Clarke, Derek Magee, Vania Dimitrova, Anthony G. Cohn, Heshan Du, Quratul ain Mahesar, Ali M. Sadeghioon, C.D.F. Rogers, David Gunn, David Entwisle, Helen Reeves, Ross Stirling, and Stephanie Glendinning. A decision support system to proactively manage subsurface utilities. In *International Symposium for Next Generation Infrastructure*, Institution of Civil Engineers, London, 11-13 Sep 2017.
- [Du *et al.*, 2016] Heshan Du, Vania Dimitrova, Derek Magee, Ross Stirling, Giulio Curioni, Helen Reeves, Barry Clarke, and A.G. Cohn. An ontology of soil properties and processes. In *Proceedings of the 15th International Semantic Web Conference (ISWC)*, page 30–37, 2016.
- [Du *et al.*, 2017] Heshan Du, Barry Clarke, David Entwisle, Mehran Eskandari Torbaghan, Richard Collins, Ross Stirling, Giulio Curioni, David Gunn, Helen Reeves, Vania Dimitrova, and Anthony Cohn. Ontologies for describing properties and processes of city infrastructure assets: the ground, roads and pipes. In *University of Leeds. [Dataset]* <https://doi.org/10.5518/190>, 2017.
- [Forgy, 1982] Charles L. Forgy. Rete: A fast algorithm for the many pattern/many object pattern match problem. *Artificial Intelligence*, 19(1):17 – 37, 1982.
- [Glendinning, 2015] Amy Glendinning. Mancunian way sinkhole: Collapsed road 'becoming tourist attraction' as visitors flock to see massive abyss. <http://bit.ly/2IDYjTt>, 2015. Accessed: 2018-05-04.
- [Heckerman and Shortliffe, 1992] David E. Heckerman and Edward H. Shortliffe. From certainty factors to belief networks. *Artif. Intell. Med.*, 4(1):35–52, February 1992.
- [Mahesar *et al.*, 2017] Quratul-ain Mahesar, Vania Dimitrova, Derek Magee, and A. G. Cohn. Uncertainty management for rule-based decision support systems. In *IEEE International Conference on Tools with Artificial Intelligence (ICTAI)*, Boston, USA, 06-08 Nov 2017.
- [McCurry, 2016] Justin McCurry. Huge sinkhole opens up in japanese city. <http://bit.ly/2GKX0ll>, 2016. Accessed: 2018-05-04.
- [Poole and Mackworth, 2017] David L. Poole and Alan K. Mackworth. *Artificial Intelligence: Foundations of Computational Agents*. Cambridge University Press, New York, NY, USA, 2nd edition, 2017.
- [Rogers *et al.*, 2012] C.D.F. Rogers, T. Hao, S.B. Costello, M.P.N. Burrow, N. Metje, D.N. Chapman, J. Parker, R.J. Armitage, J.H. Anspach, J.M. Muggleton, K.Y. Foo, P. Wang, S.R. Pennock, P.R. Atkins, S.G. Swingler, A.G. Cohn, K. Goddard, P.L. Lewin, G. Orlando, M.A. Redfern, A.C.D. Royal, and A.J. Saul. Condition assessment of the surface and buried infrastructure – A proposal for integration. *Tunnelling and Underground Space Technology*, 28:202–211, 2012.