Young, A. T. and Aylward, W. and Murray, P. and West, G. M. and McArthur, S. D. J. (2019) Automatic anomaly detection in fuel grab load trace data using a knowledge-based system vs. multiple deep autoencoders. In: 11th Nuclear Plant Instrumentation, Control and Human-Machine Interface Technologies (NPIC&HMIT) 2019, 2019-02-09 - 2019-02-14. ,

This version is available at https://strathprints.strath.ac.uk/66490/

AUTOMATIC ANOMALY DETECTION IN FUEL GRAB LOAD TRACE DATA USING A KNOWLEDGE-BASED SYSTEM VS. MULTIPLE DEEP AUTOENCODERS

A. T. Young, W. Aylward, P. Murray, G. M. West, and S. D. J. McArthur
Institute for Energy and Environment
Department of Electronic and Electrical Engineering
University of Strathclyde,
204 George Street, Glasgow,
G1 1XW
andrew.young.101@strath.ac.uk; william.aylward@strath.ac.uk; paul.murray@strath.ac.uk;
graeme.west@strath.ac.uk; s.mcarthur@strath.ac.uk

ABSTRACT

Of the seven Advanced Gas-cooled Reactor nuclear power stations in the UK, the majority are approaching their planned closure date. As the graphite core of these type of reactors cannot be repaired or replaced, this is one of the main life-limiting factors. The refuelling of a nuclear power station is an ongoing process and refuelling of the reactor occurs typically every 6 to 8 weeks. During this process, data relating to the weight of the fuel assembly is recorded: this data is called fuel grab load trace data and the major contributing factor to this are the frictional forces, with a magnitude related to the channel bore diameter. Through an understanding of this data, it is possible to manually interpret whether there are any defects in the individual brick layers that make up the graphite core but doing so requires significant expertise, experience and understanding.

In this paper, we present a knowledge-based system to automatically detect defects in individual brick layers in the fuel grab load trace data. This is accomplished using a set of rules defined by specialist engineers. Secondly, by splitting up the trace into overlapping regions, the use of multiple deep autoencoders is explored to produce a generative model for a normal response. Using this model, it is possible to detect responses that do not generalise and identify anomalies such as defects in the individual brick layers. Finally, the two approaches are compared, and conclusions are drawn about the applications of these techniques into industry.

Key Words: Fuel grab load trace, machine learning, autoencoders, condition monitoring, advanced gas-cooled reactor (AGR)

1   INTRODUCTION

There are currently eight nuclear power stations in operation within the United Kingdom, only one of these Sizewell B is not of Advanced Gas-cooled Reactor (AGR) design. Each station contains two identical reactors, and the majority of these are approaching their planned closure date. One of the major life-limiting factors of this design is the graphite core as it cannot be repaired or replaced. Assessment of the health of the graphite core is useful for supporting the case for the extension of the lifetime of the reactor. Currently, there are three main approaches to perform this assessment. Two of these approaches use data gathered during routine outages, which typically occur every 12 months to 3 years. The other approach uses monitoring data collected during refuelling events which are far more frequent and typically occur every 6 to 8 weeks.

The core of an AGR consists of thousands of interlocking hollow cylindrical graphite bricks, the arrangement of these graphite bricks within the core is shown in Fig. 1. The typical AGR design consists of eleven of these graphite bricks stacked vertically to make one fuel channel, and each of these fuel channels are arranged in a lattice structure which is also shown in Fig. 1.

Figure 1. The arrangement of graphite bricks with an AGR core [1]

It has been shown in [2] and [3] that due to the irradiation the graphite bricks receive during operation, the bricks begin to shrink, however, in later life they reach a turn-around point where the graphite bricks begin to expand. The shrinking of the bricks causes internal stresses to build up within the bricks and at the graphite turn-around point, the tensile and compressive forces switch, see Fig. 2, this can cause cracks to manifest in the fuel channel walls in several ways.

These cracks during normal operation present no impact on the safe operation of the reactor. However, in the very unlikely event of major seismic activity occurring in close proximity to the reactor, this has the potential to distort the core further and impede the insertion of the boron control rods necessary to shut down the reactor [4]. That being said even with significant amounts of core distortion and a major seismic event, only 12 out of the 81 boron control rods are required to be inserted to safely shut down the reactor [5]. While the impact of these cracks to safe reactor operation is highly unlikely, for the safety case all eventualities have to be considered and therefore it is necessary to detect and analyse these cracks.


Figure 2. Stress reversal at graphite turn-around diagram [1]

This paper discusses the initial work undertaken for the purpose of automatic crack detection in refuelling data using both a knowledge-based approach and an approach using multiple deep autoencoders (AEs). For each approach, a system for automatic crack detection was designed and these are then compared using performance metrics, i.e. precision, recall and F1 score. Finally, the explicability of both techniques is discussed in regards to providing decision support for an industrial application.

2   INSPECTION AND MONITORING

There are three main approaches used to find and analyse the cracks that can form in the fuel channel bricks. The first two approaches are performed during planned outages, which occur every 12 months to 3 years. During these outages, after the fuel has been removed, specialist inspection equipment is inserted into the fuel channels to physically measure the diameter of the fuel channel bricks at several orientations. This is performed by using either a Channel Bore Inspection Unit (CBIU) or a New In-Core Inspection Equipment Mark 2 [6]. Using these accurate measurements of the brick dimensions it is possible for

engineers to determine if there are any cracks present within the fuel channels as these manifest as dimensional changes within the bricks.

Also, during these physical inspections, video footage of the inside of the channel is taken at each of the orientations, and it has been shown in [7] that it is possible to produce a chanorama (fuel channel panorama) for each fuel channel inspected.

Both types of data can only be gathered when the station is offline and is therefore collected periodically, rather than continuously. However, additional information that can be retrieved during normal operation is advantageous and can be used to support the safety case. The final approach uses data collected during refuelling events to determine if there are any cracks present in the fuel channels. This data is gathered every 6 to 8 weeks when the fuel stringers of selected channels are removed and replaced with new fuel stringers. Throughout the refuelling process, the mechanical load of the fuel stringer is measured, and the resulting data is known as a Fuel Grab Load Trace (FGLT). Originally this data was used to inform the reactor protection system, ensuring that refuelling stopped in the event of a restriction in the channel or the failure of the tie-bar mechanism. However, closer inspection can reveal further information about the dimensional changes of the channel [8][9].

The shape of the FGLT is driven by the mass of the stringer, the frictional interaction of the stabilising brushes with the fuel channel walls (see Fig. 3a and 3b), and, for onload refuelling, from the effects of coolant gas flowing within the core. During refuelling, the fuel stringer is first removed from the fuel channel. The data collected at this time is called the discharge FGLT. Following this, when the new fuel is inserted into the channel the data gathered is called a charge FGLT. Fig. 3c shows an example of both a charge and discharge FGLT. The regions highlighted indicate the individual BLs, with the trough before the 3-peak spike indicating the brick interface for the discharge data, and the peak before the three trough dip for the charge data. The 3 peaks or troughs in either set of data relate to each set of brushes in the nose assembly passing over the brick interface.

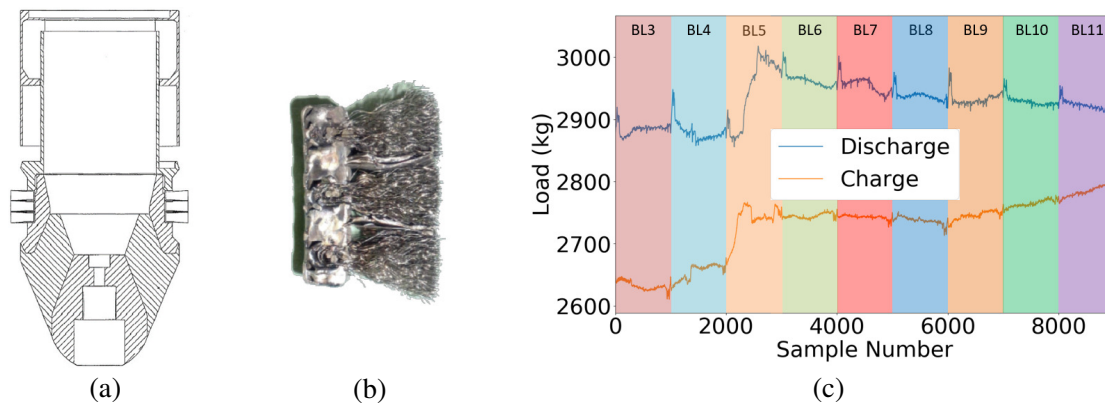### 3   FUEL GRAB LOAD TRACE ANALYSIS

Figure 3. (a) Fuel stringer nose assembly [9], (b) Lower stabilising brush [9], and (c) Example of FGLT data

The analysis of fuel grab load trace data is an important research area, as it has the potential to provide information about the health of the graphite core in an AGR without the requirement of turning off the reactor. As previously explained, the FGLT is affected by frictional effects of the fuel assembly brushes, which are designed to form a close fit with the fuel channel walls.

As the diameter of the channel bore reduces due to the shrinkage of the graphite, an effect called "barrelling" occurs whereby the channel diameter at the top and bottom of the bricks reduces in comparison to the centre. For a discharge FGLT, a narrower bore is a restriction in the movement of the fuel and

therefore translates to an increase in apparent weight of the fuel stringer. For a charge FGLT, any restriction in movement supports some of the weight of the fuel stringer and so corresponds to a decrease in apparent weight.

For normal bricks without any defects, the FGLTs generated are uniform in shape, with brick interfaces and the barrelling affects across the BL visible. Fig. 4 shows one such example. The top and bottom edges of the bricks are chamfered, and the interfaces appear on the FGLT as a large peak or trough. This peak is usually made up of three smaller peaks, corresponding to the three rows of brushes shown in Fig. 3a and Fig. 3b which make up the brush assembly.


Figure 4. FGLT for brick layers 6 to 8, all bricks are uncracked

When defects do occur, cracks may appear circumferentially - around the brick diameter, or axially - along the entire length of the brick. Circumferential cracks lead to barrelling of both brick sections, which leads to a reduced diameter at the crack location. This results in a small peak or trough in the FGLT at this location depending on whether the trace refers to fuel removal or insertion, for the reasons described above. Fig. 5a shows an example of a circumferential crack at approximately sample number 1500. Axial cracks may cause the diameter of the entire brick to increase slightly. Therefore, there is no spike in the FGLT but there is a difference in the perceived average load across the entire brick as compared to neighbouring bricks, or previous traces from the same brick. The prominent brick interface spike described above may also reduce in magnitude. These features are shown in Fig. 5b. When inspecting FGLTs, engineers may refer to the presence or absence of many of these features in order to gain a better understanding of the condition of the fuel channel in question.


Figure 5. FGLT when brick layers 6 and 8 are uncracked, and 7 is cracked (a) Circumferential (b) Axial

4   METHODOLOGY

This section discusses the two approaches taken for the classification of axial cracks in FGLT data. Circumferential cracks do not present an issue for the normal operation of the nuclear reactors and are also easier to detect. Axial cracks present a complex problem with very subtle features, and their detection is also of current use for the industrial application. As the dataset available is heavily imbalanced, containing a large proportion of normal data and only a few instances of abnormal data, it is challenging to train a standard classifier to detect both cracked and uncracked BLs. Therefore, the approach adopted is to design a technique that learns what the range of normal responses could be based on all historical data, then classify whether or not any new data falls within this range and is likely uncracked or out with this range and is likely cracked.

First, a knowledge-based system is proposed, that implements a set of rules based on the information provided by engineers. Secondly, an approach using multiple deep AEs is proposed that splits up the trace into overlapping regions so that defects can be detected in specific BLs.

4.1 Knowledge-Based System

Knowledge-based systems (KBS) are typically used to simulate human decision making. They attempt to digitize the knowledge and experience of engineers into a set of rules that can be applied to a given dataset, with the aim of assessing a problem just as a human would. A key part in the implementation of many of these KBS is the knowledge acquisition step [10], based on the assumption that the knowledge already exists and only needs to be collected and implemented into the system.

4.1.1   Step Change in Load

The simplest of the three rules to implement was the comparison between the average load value of the BL under analysis, with the average load of the BL before and after. Fig. 5b shows an example of how an axial crack will appear in an FGLT when the surrounding bricks are both uncracked. As the brick interface location is at the beginning of the BL for a discharge trace and at the end of the BL for a charge trace the region selected for calculating the average load must also change. For all discharge data, 10% to 95% will be included, and for the charge data, 5% to 90%, which avoids the brick interface being included in the calculation. The difference in the average load before and after can then be compared with the change in load for historically uncracked bricks to produce two rules in the KBS.

4.1.2   Change in Brick Interface Height

The second rule to be implemented was the change in brick interface height. Unlike with the previous rule the same approach can be adopted for both charge and discharge data. Fig. 5b shows an example of a brick interface within an FGLT. By developing a peak and trough detection algorithm, it is possible to calculate the height of the brick interface which can then be compared with the height for historically uncracked bricks to produce the third rule in the KBS.

4.1.3   Comparison with Internal Normal Response Model

The final rule to implement was to try and capture what the "normal response" of three consecutive uncracked bricks would be. Due to the variation in the FGLT, a separate "normal response" is required to be produced for each set of three BLs. A moving average filter was then used to remove the noise and an example of the "normal response" for BL 6, 7 and 8 is shown in Fig. 6. The root-mean-square error (RMSE), is calculated for the given data vs. the "normal response". This error can then be compared with the errors produced for historically uncracked bricks to produce the final rule in the KBS.

Figure 6. Normal response for brick layers 6 to 8

4.2 Deep Autoencoders

AEs are a specific type of neural network that are symmetrical and attempt to reproduce their input data as their output data, see Fig. 8. The first half of the AE, the encoder, reduces the original feature space to only the essential features required to reconstruct the data. The second half, or the decoder, attempts to reconstruct the original input data from the reduced feature space. Training the network is a process that attempts to minimize the reconstruction error by altering the weights between the connected


Figure 7. Architecture of individual deep autoencoder

layers. Deep AEs, when the network has more than one hidden layer, have been shown in [11] to be more effective at learning complex features as each layer will learn more complex features than the last.

The second proposed approach using multiple deep AEs for the detection of anomalies in FGLT data uses a bank of seven AEs, see Fig. 9, each trained for an individual BL. This allows for anomalies to be found in specific regions (BLs) of the data, as opposed to the traditional output from a single AE of cracked or uncracked for the entire trace. By splitting up the data into regions of three BLs, one before and after the BL being analysed, and only supply it with uncracked data, it is possible for the AE to learn the essential features required to reconstruct a normal response. When supplying any of the individually trained AEs with an uncracked FGLT, the reconstruction error will be relatively small as it will generalise to the features the AE has learned. If supplied with a dataset containing a crack the reconstruction error will be relatively large and by finding the correct threshold it is possible to distinguish traces generated from cracked bricks.

Each individual AE has the standard bottleneck architecture [12] [13], with 3000 input and output neurons, then "64 – 8 – 64" for the three hidden layers. The input and outputs neurons are constrained by the data, each BL is interpolated to 1000 points, where the neurons for the 3 hidden layers were selected using a brute-force approach to determine the optimal values for the training set. Sigmoid activation functions were used for each of the layers, as this is the most commonly used activation

function in autoencoders for non-linear mapping and Adam gradient descent [14] was used to estimate the parameters in the network with an initial learning rate of 10^{-4}.
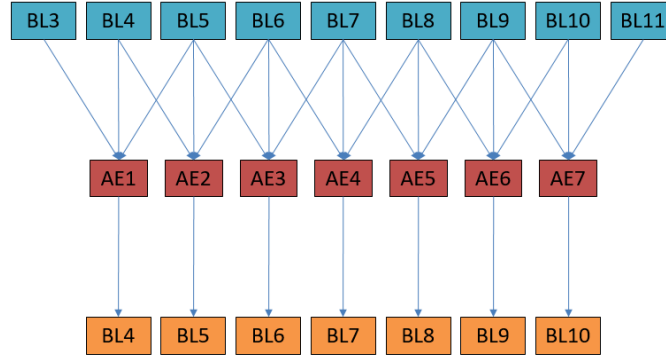

Figure 8. Architecture of multiple deep autoeconders

5   EVALUATION

5.1 Analysis

As mentioned in the previous section the FGLT data available is heavily imbalanced, which presents a problem for conventional classification techniques. By adding the constraint that for each run of three BLs the upper and lower BL must be uncracked there is on average a 9:1 ratio of uncracked to cracked FGLTs in the subset of data used for this study.

For this application the penalty for missing a crack is extremely high, so the most important metric to be considered is the recall, see Equation 1. However, because a classifier that predicts a crack for every brick it encounters would return a recall of 100%, this alone is not the only metric used. The precision, see Equation 2, is also calculated as this considers the number of false positives. Finally, these two metrics are combined to produce the F1 Score, see Equation 3 [15], where, TP is the number of True Positives, FP is the number of False Positives, and FN is the number of False Negatives. This gives a single value for each approach for each BL that can be easily compared.

$$Recall = \frac{TP}{TP + FN} \tag{1}$$

$$Precision = \frac{TP}{TP + FP} \tag{2}$$

$$F1\ Score = 2 \times \frac{Precision \times Recall}{Precision + Recall} \tag{3}$$

5.2 Results

5.2.1   Knowledge-Based System

Using all the uncracked data for each BL it is possible to produce a value for each of the rules selected previously. Having analysed the data, and removed the outliers, the results showed a normal distribution, so a threshold of two standard deviations above and below the mean was selected. By supplying the system with the cracked data, using the defined thresholds, it was possible to classify all the data into cracked or uncracked. The results for the KBS are shown in Table I.

Table I. Precision, recall and F1 score for knowledge-based system approach for both discharge and charge datasets

Brick Layer	Discharge			Charge		
	Precision	Recall	F1 Score	Precision	Recall	F1 Score
BL4	23.81	61.54	34.33	10.53	60.71	17.94
BL5	30.00	63.64	40.78	11.76	55.56	19.42
BL6	11.11	64.00	18.93	14.29	57.14	22.86
BL7	18.52	55.00	27.71	38.89	40.74	39.79
BL8	21.74	54.55	31.09	35.71	39.13	37.34
BL9	7.69	68.57	13.83	16.67	57.69	25.86
BL10	12.50	80.77	21.65	4.35	81.48	8.26

It should be noted that the recall for BL10 is above average, however, this is due to there being a smaller number of cracks present in this BL. Despite, the recall being relatively high both the precision and F1 score are relatively low, this is due to a large number of false positives. After analysing the false positives, some of the confirmed uncracked bricks appear to have some of the same features as the cracked bricks, there are also cases where it would not be possible to manually classify a cracked brick, so it was expected that the KBS would not be able to classify these bricks either.

5.2.2 Deep Autoencoders

The same approach was adopted for the deep AEs: first, the range of reconstruction error for the uncracked bricks was calculated. Again, after removing the outliers, this was shown to be a normal distribution, so two standard deviations from the mean was again selected as the threshold. This was however only applied to an error greater than the mean reconstruction error. As before the entire dataset for each BL was input into the AE based system and the results are shown in Table II.

For the AE based approach, the results are far more promising, using this system it is possible to correctly classify all except one crack in BL8 in the charge dataset. Considering the excellent recall, the precision and F1 score are also relatively high due to on average only misclassifying 4 uncracked bricks as cracked. The only downside to this approach is that the precision and F1 score for BL10 is poor compared with the other BLs for the charge dataset, with more than double the average number of false positives. After further inspection of this BL, it is believed that this is due to the unusual shape, and consistently smaller features in this BL, these are the reasons the detection of cracks in this BL is more difficult.

Table II. Precision, recall and F1 score for multiple autoencoder approach for both discharge and charge datasets

Brick Layer	Discharge			Charge		
	Precision	Recall	F1 Score	Precision	Recall	F1 Score
BL4	71.43	100.00	83.33	68.75	100.00	81.48
BL5	66.67	100.00	80.00	80.00	100.00	88.89
BL6	75.00	100.00	85.71	64.29	100.00	78.26
BL7	81.82	100.00	90.00	80.00	100.00	88.89
BL8	75.00	100.00	85.71	76.47	92.86	83.87
BL9	78.57	100.00	88.00	78.57	100.00	88.00
BL10	62.50	100.00	76.92	35.71	100.00	52.63

5.2.3 Comparison

Based on the results it is clear that the AE based approach outperforms the KBS which is likely due to the variability of the data and the subtlety of the features being analysed. Given that engineers are correctly classifying defects manually with a greater accuracy than the KBS, and a similar accuracy to the AE based approach, there are perhaps improvements to be made in the knowledge acquisition step. For data with complex features such as this, it is often difficult to transfer the knowledge from the expert to an algorithm without extensive knowledge acquisition [16]. Even with the relatively small dataset available the AE approach is able to learn the features required to reconstruct the uncracked data accurately. By exploiting the fact that the cracked data will not generalise to these features it is possible to separate the cracked and uncracked data. However, as this approach is a black box technique the explicability of this approach could present a problem when trying to deploy this in industry. By looking at the individual rules it is possible to understand why the system has classified a specific trace as cracked or uncracked, which is not possible with the AE approach. Given the excellent recall for the AE based approach, it is possible that this could be used to produce a list of data to the engineers that require further analysis and be used to provide decision support while still keeping the human in the loop.

6   CONCLUSIONS

This paper has described the initial work undertaken into the application of a knowledge-based system and an approach using multiple deep AEs to detect anomalies in fuel grab load trace data. It was shown that for the problem of automatic FGLT analysis there are several key difficulties in designing a KBS which mainly involve the knowledge acquisition and transfer steps. This is particularly difficult when it is not possible to adopt a formalized approach to the knowledge acquisition and processing, e.g. KADS (Knowledge Acquisition and Documentation Structuring) [17]. This work also showed that the AE approach has great potential to provide decision support to engineers who currently classify defects manually. The main drawback of this approach is that for an industrial application there is often a requirement to provide supporting evidence to explain why a specific brick layer has been classified as cracked. Due to the nature of AEs it is extremely difficult to ascertain the reason behind this, however, it would be more likely that this technique could be used to guide the engineers to any potentially cracked bricks with the human still in the loop to make the final decision on why the brick layer has been classified as cracked.

7   ACKNOWLEDGMENTS

8   REFERENCES

1. A. G. Steer, "AGR core design, operation and safety functions," *Management of Ageing Processes in Graphite Reactor Cores*, pp.11-18 (2007).

2. J. E. Brocklehurst, B. T. Kelly, "Analysis of the Dimensional Changes and Structural-Changes in Polycrystalline Graphite Under Fast-Neutron Irradiation," *Carbon*, **31**, pp.155-178 (1993).

3. B. J. Marsden, G. N. Hall, O. Wouters, J. A. Vreeling, J. Van Der Laan, "Dimensional and material property changes to irradiated Gilsocarbon graphite irradiated between 650 and 750 ℃," *Journal of Nuclear Materials,* **381**, pp.62-67 (2008)

4. "Explaining Graphite Fact Sheet", EDF Energy, https://www.edfenergy.com/sites/default/files/graphite_overview_a4_v4.pdf (2018)

5. "Graphite core ageing", Office for Nuclear Regulation, http://www.onr.org.uk/civil-nuclear-reactors/graphite-core-ageing.htm (2018)

6. A. Cole-Baker, J. Reed, "Measurement of AGR Graphite Fuel Brick Shrinkage and Channel Distortion," *Management of Aging Processes in Graphite Reactor Cores*, pp. 201-208 (2006)

7. G. M. West, P. Murray, S. D. J. McArthur, "Improved visual inspection of advanced gas-cooled reactor fuel channels," *International Journal of Prognostics and Health Management,* **6** (2015).

8. C. Berry, G. M. West, S. D. J. McArthur, "Semi-supervised learning approach for crack detection and identification in advanced gas-cooled reactor graphite bricks," *10th International Topical Meeting on Nuclear Plant Instrumentation, Control and Human Machine Interface Technologies,* San Francisco, CA, June 11-15 (2017)

9. G. M. West, C. Wallace, S. D. J. McArthur, "Combining models of behaviour with operational data to provide enhanced condition monitoring of AGR cores," *Nuclear Engineering and Design*, **272**, pp.11-18 (2014)

10. R. Studer, V. R. Benjamins, D. Fensel, "Knowledge engineering: Principles and methods," *Data & Knowledge Engineering*, **25**, pp.161-197 (1998)

11. H. Larochelle, Y. Bengio, J. Louradour, P. Lamblin, "Exploring strategies for training deep neural networks, *Journal of Machine Learning Research*, **10,** pp.1-40 (2009)

12. Y. Bengio, "Learning deep architectures for AI," *Foundations and Trends in Machine Learning,* **2**, pp. 1–127 (2009)

13. Y. Bengio, A. Courville, and P. Vincent, "Representation learning: A review and new perspectives," *IEEE Transactions on Pattern Analysis Machine Intelligence*, **35**, pp. 1798–1828 (2013)

14. D. Kingma, J. Ba, "Adam: A method for stochastic optimization," *International Conference for Learning Representations*, (2015)

15. C. J. van Rijsbergen, "*Information Retrieval,*" Butterworths, London, UK (1979)

16. J. H. Boose, B. R. Gaines, "Knowledge acquisition for knowledge-based systems: Notes on the state-of-the-art," *Knowledge Acquisition: Selected Research and Commentary*. Springer, Boston, MA, (1989)

17. B. J. Wielinga, A. Th. Schreiber, J.A. Breuker, "KADS: a modelling approach to knowledge engineering," *Knowledge Acquisition,* **4**, pp.5-53 (1992)