A Case-Based Real-Time Adaptive Engineer Site Support System

Kyle Martin

School of Computing Science and Digital Media,
Robert Gordon University,
Aberdeen AB25 1HG, Scotland, UK
k.martin@rgu.ac.uk

Abstract: Employee experience is a valuable asset for any company. A system which can store, retrieve and adapt these experiences to meet the requirements of new scenarios can play an important role in corporate knowledge management. Case-Based Reasoning provides an excellent mechanism for this because it allows the capture and reuse of past experiences, which in this paper involves the generation of expert-level support in response to questions raised by British Telecommunications field engineers. However, experience capture is hard in dynamic environments involving multi-modal communication and content. This paper examines the context of this research and details the work which has been completed thus far, as well as potential next steps.

Keywords: Case-Based Reasoning · Siamese Neural Networks · Knowledge Management

1 Introduction

This research project is a collaboration between British Telecommunications (BT) and Robert Gordon University to produce a system which uses experiential content gathered from users as a case-base to answer new queries. The query and answer process will take place while engineers are out in the field, giving them access to the support they require within the dynamic environment of their jobs. This would facilitate the exchange of knowledge and experience between employees within the company and assist in the development of a 'corporate memory', which stores the relevant experiences of every engineer in the company, preventing these assets from being lost if an employee were to leave BT.

This paper is structured as follows: section 2 gives an overview of the context of the project and section 3 describes its contributions. Section 4 discusses related work and research. The report concludes with a description of the research which has been completed so far and details future work and next steps the research could take.

2 An Overview of the Project

The goals of this project are to develop a means of capturing experiential knowledge from BT engineers and to produce an application which can learn to provide expert-level support in the field using that information. The system should be able to support engineers by retrieving relevant results and adapting to new situations accordingly, producing results which are applicable in the real-world.

In the field, an engineer is only able to access notes (historical task records) which are part of the current task. These notes are plain text and contain customer and location information, but rarely indicate what piece of equipment the fault may relate to (as this may not be known at the time of task allocation), the context of the task, or similar tasks. An engineer is expected to rely heavily on past experiences and training in order to diagnose and solve problems, but this can fail if the task has novel, unusual or specialist elements. Some of these failures may be avoided if there existed a means of drawing upon the experiences of engineers who had previously encountered similar problems.

To give a real example, a BT power engineer was called out to a 'low-voltage' alarm in a rural exchange. It transpired that the mains power had failed and the back-up generator had not started, so the exchange was beginning to lose power and risked the network going down for nearby customers. The engineer was unsure of what the exact fault with the back-up generator could be and spent some time attempting to diagnose the problem before calling for help. It was only when an engineer with more experience arrived that the fault was eventually diagnosed - the ac/mc contactor within the generator needed replaced. This was a time-dependent and critical fault which could have had important business repercussions and was only solved because a more experienced engineer was able to attend on short notice. If experience and knowledge could be more effectively transferred between engineers, then we could severly reduce the time taken to solve these faults, as well as the manpower required to do so and the risk of failing to complete time-critical faults.

We aim to use the notes and other knowledge assets (including video, photo and audio content) as the case base for an intelligent system which can reply to engineer's queries and propose a solution. This would alleviate pressure upon engineers by giving them access to a 'digital expert' which could draw upon historical experiences from the entire work force to provide support in the field.

3 Contributions

As this project is a collaboration between a representative of industry and a university, it is important that the research provides an academic contribution, but remains viable for use in a commercial setting. There often exists a disconnect between the two, fuelled by the exploratory nature of research and the fact that 'state-of-the-art' measures may involve expensive procedures or equipment which do not have commercial viability. Therefore, this project is a case study of a research project which has both academic importance and good business viability and offers contributions in both areas.

The main contribution of this project towards industry will be to improve information access for engineers in the field and facilitate the exchange of knowledge and expertise by using an intelligent system. By doing so, this system should ultimately improve an engineer's ablity to complete a task and, at an organisational level, increase overall engineer productivity.

The contribution of this project to academia is to develop a dynamic decision support system which can operate on a large scale across multiple engineer domains. This project will ultimately showcase a system which can retrieve relevant results from vast quantities of complex, inter-related multi-media data.

4 Related Works

Due to its very nature of reusing and adapting previous examples, Case-Based Reasoning (CBR) is the branch of machine learning which most closely reflects the goals of the project. Using CBR techniques to facilitate knowledge flow between users is not a new concept, having been pursued with differing levels of success for a number of years. Of these, many projects have specifically targeted domain experts within a pre-identified industry niche, including food quality control [8] and help desk support [5], in order to provide relevant assistance based upon experts' and users' experiences.

In [6] Goker et al developed an adaptive expertise provider dubbed the Pricewaterhouse Cooper (PwC) Connection Machine. The Connection Machine allowed users to enter their queries into a web application and made use of CBR techniques in order to identify experts who may be able to answer. Use of the system facilitated the exchange of knowledge and experience between users and provided a singular forum for accessing all experts within the company. The biggest disadvantage of the system was that it relied upon experts to actively answer queries. Drawing upon this idea, our project aims to allow users to access the sum of all experience of BT engineers in a single place, but remove the need for human experts to explicitly answer queries. The system will return relevant answers based upon its knowledge gained from the input task notes.

A Case Retrieval Net (CRN) is a CBR framework which facilitates the return of a small number of cases in a large case base. CRNs use a memory structure that stores both the contents of the case base and similarity knowledge between cases [9] using Information Entities (IEs). An IE is any specific piece of information pertaining to a case (such as an attribute-value pair). Results are returned using 'spreading activation'; the most relevant IE to a query is activated and nearby IEs receive diminishing activation the less similar they are to the identified IE. The case nodes associated with the activated IEs are then collected and returned. CRNs have demonstrated promising results in reuse of textual cases within large medical databases [1]. The return of a small number of relevant results from a massive case base and the use of similarity knowledge to facilitate case adaptation are both vital components of the project. However, this requires a method of generating the extensive similarity knowledge required for spreading activation of the net. This may be achieved using the object-to-object similarity generated by a Siamese Neural Network.

A Siamese Neural Network (SNN) architecture consists of two neural networks that share identical weights and are joined at one or more layers. Introduced in [2] as a method of signature verification, SNNs are trained and tested on pairs of examples to develop similarity knowledge at a case-to-case level. Desirable pairs are dubbed as 'genuine' during training, while undesirable pairs are 'impostors', so that the network develops vectors representative of case features. At test time, the SNN measures the distance between the queried vectors to determine whether they are 'genuine' or 'impostor' based on a threshold.

Recent research has demonstrated that SNNs are able to generate object-to-object similarities after being trained with relatively few examples or in datasets where a vast number of classes exist [7]. This could be particularly useful in the current project, where the broad domain could mean that there are a massive number of classes within the case base. SNNs have been applied with success in areas like sketch-based retrieval [10], and speaker recognition [3]. In [4], an SNN is applied to the task of similar question retrieval, and outperforms the state-of-the-art. In the same way, this project would aim to return similar cases to the situation described in the query, but unlike in the examples above it would also attempt to adapt these cases to better suit the described situation.

5 Current and Future Work

Much of our recent work has been gathering data to determine the industrial and academic context of this project. In particular, we reviewed literature featuring industry examples of experience capture and knowledge transfer systems to see how others have dealt with similar problems. In addition, we gathered data from within BT to establish the specific business context of the project. This involved determining the available information sources for use by engineers, how they are used on specifc tasks and in what areas they are lacking.

One of the key aspects of this project is the development of a large and dynamic case base which can be used and updated in real-time throughout the day. Often, retrieval from a huge case base can be extremely costly. We are examining methods of reducing this cost without sacrificing case base coverage or retrieval accuracy through similarity-based retrieval in a CRN, but learning similarity knowledge in a huge system can be expensive. To this end we are performing experiments to learn similarity between cases in a quick and inexpensive manner. Recently we have examined generating case-to-case similarity knowledge by using an SNN and have demonstrated that this is capable of developing case-to-case similarity knowledge suitable for similarity-based retrieval.

In future work, we would like to experiment with training SNN on limited data to ascertain whether they can successfully learn similarity knowledge. Also, we would like to examine populating the IEs of a CRN using the values developed by the output of an SNN to see whether we can return improved results with spreading activation.

References

[1] Adeyanju, I. et al. 2009. Case Retrieval Reuse Net (CR2N): An Architecture for Reuse of Textual Soutions. In: L. McGinty, D. C. Wilson (eds) Case-Based Reasoning Research and Development. ICCBR 2009. Lecture Notes in Computer Science, vol. 5650. Berlin: Springer-Verlag

[2] Bromley, J. et al. 1993. Signature Verification Using a"Siamese" Time Delay Neural Network. International Journal of Pattern Recognition and Artificial Intelligence, 7(4) pp. 669 - 688.

[3] Chen, K and Salman, A. 2011. Extracting Speaker-Specific Information with a Regularized Siamese Deep Network. In: J. Shawe-Taylor et al (eds) Advances in Neural Information Processing Systems 24 (NIPPS 2011). Granad, Spain, December 12 17, 2011.

[4] Das, A. et al. 2016. Together We Stand: Siamese Networks for Similar Question Retrieval. In: Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (ACL 2016). Berlin, Germany, August 7 - 12, 2016. Stroudsburg, PA, USA: ACL

[5] Goker, M.H. and Roth-Berghofer, T. 1999. Development and Utilization of a Case-Based Help-Desk Support System in a Corporate Environment. In: K-D. Altho et al. (eds) Case- Based Reasoning Research and Development: Third International Conference on Case-Based Reasoning, ICCBR-99. Seeon Monastery, Germany. July 27 - 30. Berlin: Springer-Verlag

[6] Goker, M.H. et al. 2006. The PwC Connection Machine: An Adaptive Expertise Provider. IN: Thomas R. Roth-Berghofer et al. (eds) Advances in Case Based Reasoning 8th European Conference, ECCBR 2006. Fethiye, Turkey, September 4 - 7, 2006. Berlin: Springer-Verlag.

[7] Koch, G. et al. 2015. Siamese Neural Networks for One-Shot Learning. In: ICML 2015. Lille Grand Palais, Lille, 6 11 July, 2015.

[8] Lao, S. I. et al. 2012. Achieving Quality Assurance Functionality in the Food Industry Using a Hybrid Case Based Reasoning and Fuzzy Logic Approach. Expert Systems with Applications. Vol 39. pp 5251 - 5261

[9] Lenz, M. and Burkhard, H. D. 1996. Case Retrieval Nets: Foundtions, Properties, Implementation and Results. [Technical Report] Berlin, Germany: Humboldt University

[10] Qi, Y. et al. 2016. Sketch-Based Image Retrieval via Siamese Convolutional Neural Network. In: 2016 IEEE ICIP. Phoenix Convention Centre, Phoenix, USA. September 25 28, 2016. Red Hook, NY, USA: Curran Associates