Appendix 1

Industrial Applications of Ripple-Down Rules

This appendix provides brief summaries of the various applications of RDR that we are aware of. It only covers applications where RDR is or has been used on some sort of routine basis. A range of other applications based on RDR have been developed, but only reported in research papers. These are presented in the next appendix.

The information available about these industrial applications of RDR varies greatly. Some companies do not provide information about the technology they use and often the information has come from personal communication.

It is interesting to note that in all the industrial applications we know of, except for YAWL, there was a key person promoting RDR who had had direct experience of RDR previously. The reason this book has been written is precisely to give potential users of RDR some direct experience of RDR.

PEIRS (1991-1995)

PEIRS (Pathology Expert Interpretative Reporting System) was in routine use in the Department of Chemical Pathology St Vincent's Hospital Sydney for about four years. It was used to provide interpretative comments for chemical pathology laboratory results. Before PEIRS the only automated comments possible at St. Vincent's were very generic comment produced by simple triggers such as an analyte result being out the normal range. For more specific comments a pathologist would have to select from a comment bank or write a comment. With PEIRS much more specific comments were produced automatically, but were checked by a pathologist (Edwards, Compton et al. 1993) (Preston, Edwards et al. 1994)..

PEIRS used Single Classification RDR and ended up with about 2,000 rules built by a chemical pathology registrar. About 200 rules where developed off-line before the system went into routine use . A report could contain results on any of up to 200 analytes, but normally contained no more that 20 results. Reports contained five columns of data representing the last five sets of results for a patients. There were a number of simple operators such as minimum, maximum, average, current, nett change that could be used in rule conditions. The laboratory issued about 500 reports a day and comments were provided for about 100 of these.

PEIRS went out of use when a new Hospital Information System was deployed which was not capable of referring cases to a system like PEIRS and receiving back clinical interpretations. Although PEIRS was in routine clinical use, it was not commercialised.

Origins

Research on RDR first started Garvan Institute of Medical Research at St. Vincent's Hospital, to deal with the maintenance problems of GARVAN-ES1 (Horn, Compton et al. 1985,

Compton, Horn et al. 1989) PEIRS was initiated by Garvan Institute researchers who had moved to Chemical Pathology at St. Vincent's and also the University of New South Wales.

Pacific Knowledge Systems

PKS (http://pks.com.au) was set up following the success of PEIRS. PKS started out developing a generic RDR tool, but with an emphasis on interpreting laboratory data. It's business focus is now almost entirely on interpreting laboratory data and auditing data entry to ensure requests are appropriate. Various documents and case studies on the PKS website indicate:

- PKS technology is used at over 90 sites world wide
- over 800 user-developed knowledge bases are in use.
- Over 28 million patient reports generated annually
- Across PKS customers over 80% of patient-specific reports are issued without needing human validation
- Thermo Fisher provides interpretative comments to over 3,000 users covering 70% of the world's immunology testing
- In one large laboratory request errors have reduced by 73%
- RDR-based comments improve detection of familial hypercholesterolaemia (Bender, Edwards et al. 2016)
- Real-time RDR advice for patients at cardiac risk, reduced bed days, readmission, testing and saved money.
- PKS has partnerships with major laboratory the technology companies Abbott, Philips and Thermo Fisher

These knowledge bases are developed by chemical pathologists themselves after a few days training from PKS staff on how to use the technology. In the training period PKS staff may develop a "seed" knowledge base, to which the pathologist will continue to add rules. The largest knowledge base in use has about 16,000 rules added incrementally over years. (Compton 2013) shows the evolution of a 3,000 rule knowledge over 8 years. In the first year about 28 hours were spent adding rules, with about 10 hours or less in later years. The median time for a chemical pathologist to add a rule is about a minute or two http://pks.com.au/technology/resources/. The reason the time to add a rule is so small is because when a pathologist wants to add a rule they have already identified some feature(s) in the data as the reason to add a rule, so their only task is selecting the feature(s) using the interface provided and if requested also selecting some other discriminating feature(s). See also (Compton, Peters et al. 2006, Compton, Peters et al. 2011, Compton 2013).

PKS RDR is based on Multiple-Classification RDR (MCRDR) (Kang, Compton et al. 1995, Kang 1996, Kang, Compton et al. 1996) to enable multiple independent conclusions for a case rather than single conclusions. It provides a language to enable users to describe and extract features to be then used in rule conditions. It also provides a resource to allow users to set validation levels http://pks.com.au/technology/resources/. A major issue in any PKS deployment is interfacing between PKS technology and the laboratory information system.

Origins

PKS was founded by members of the group who had initiated and worked on PEIRS and other associates[1]. It was developed through venture capital funding and is now majority owned by a private investment company.

Ivis

Ivis (https://www.ivisgroup.com) provides technology for multi (now omni)-channel retailing. Tesco, famous for its use of multi-channel retailing is one of Ivis's customers. Ivis Sonneto technology is based on an integration of Ripple-Down Rules and Conceptual Graphs (Ellis 2005, Sarraf and Ellis 2006) and http://jtonedm.com/2008/12/29/first-look-sonetto-retail/. There is no mention of RDR on the current website, but the references above make clear the basis of the technology.

Basically the rules were used to customise the user experience, for example in suggesting products that may be of interest. If a British user was looking at a holiday in Spain, they may well be interested in holiday on a Pacific island that was on special. This type of recommendation cannot be made by techniques such as collaborative filtering, as it requires knowledge that a British person wanting a holiday in Spain probably wants a holiday in the sun. By 2006 apparently 100s of 1000s of rules had been written (personal communication)

Origins

Ivis developed their RDR system after Gerard Ellis, PKS software engineer, took up a senior position at Ivis

Erudine Pty Ltd

Erudine no longer exists, although apparently at one stage had over 70 staff, but is included here because of their particular application of RDR. It is claimed that Erudine failed not because of the technology, but because very large debtors did not pay their bills https://www.bloorresearch.com/2014/01/simulating-behaviour-to-replace-legacy/. The technology reappeared in another company, Erudine Financial, but this has disappeared. There was no acknowledgement that Erudine's Behaviour Engine was based on Ripple-Down Rules, despite describing an identical approach; however, the company was previously known RippleDown Solutions and can be found on the Wayback Machine (www.rippledownsolutions.com) and there is Linkedin page (https://www.linkedin.com/in/martin-rice-834ba7/) indicating that the CEO of Erudine was also the owner of RippleDown Solutions.

Erudine's main focus was re-engineering legacy software systems. As a legacy system processed cases, each case together with the legacy system outcome for that case would be passed to a business analyst who would write rules to reproduce the outcome for the case using a Ripple-Down Rule approach. Although Erudine apparently failed because of debtors, perhaps a challenge for the technology was that some systems to be re-engineered required very complex outcomes, and constructing these was a major task for the business analyst beyond the simple task of identifying differentiating features to build rule. But this is speculation

[1] Paul Compton has a small shareholding in PKS

Origins

Erudine was founded by a software engineer who had worked for Ivis

Ripple-Down Rules at IBM

IBM Research carried out a research project on cleansing Indian street address data using a commercial tool, a decision treat and a conditional random fields method and RDR (Dani, Faruquie et al. 2010). The methods all worked similarly on test data from the same area of India as the training data, but the RDR-base method worked much better when applied to all-India data presumably because the RDR rules reflected an understanding of what an address should be, whereas the learning methods were statistical.

The RDR method used a different single classification RDR for each field in an address. Initial rules were also added independent of cases. The root node of each tree specifies a dictionary and there can be cases which are fixed with a dictionary update rather than a rule and each rule refers to a dictionary against which a token is matched. We point out these adaptations of RDR to highlight that RDR is more a strategy and approach than a particular technology which then needs to be adapted to the target application.

In 2010 this research project was recognised as an IBM "Research Accomplishment". For a research work to be recognized as a Research Accomplishment it must drive new business worth at least USD $10M million and the three team members were recognized with Outstanding Technical Achievement Awards. Apparently the resulting technology has been sold to numerous customers.

Origins

Ashwin Srinivasan who as post-doctoral implemented the original PEIRS system introduced RDR to IBM Research many years later

YAWL

YAWL (http://www.yawlfoundation.org) is an open source workflow language. It allows for complex data transformations, and integration with other organizational resources and Web Services. It uses RDR to be able to make specific decisions about different parts of the workflow. Single classification RDR is used and there can be many RDR associated with different decision points in the workflow (Adams, ter Hofstede et al. 2006). There have been recent extensions to how RDR is used in YAWL (Adams 2013).

Although YAWL is open source it has had significant industry uptake internationally. (http://www.yawlfoundation.org/pages/impact/uptake.html). RDR is YAWL seems to have been of significant research interest with many papers referring to this context-specific aspect of YAWL

Origins

YAWL is the one industrial application of RDR we know of where does not seem to have been direct involvement in earlier RDR development

Medscope

Medscope (https://www.medscope.com.au) provides advice to pharmacists about drug interactions using RDR. Since commercialisation in 2009, over 1200 individual pharmacists have used MRM and the system has made 800,000 recommendations (verbal advice from Medscope). Pharmacists use this system voluntarily, so this repeat use clearly suggests the value of its advice. The systems also finds more potential drug problems than pharmacists (Curtain, Bindoff et al. 2013).

Origins

Medscope's use of RDR resulted from a PhD project at the University of Tasmania supervised by Byeong Kang

---

Seegene

SeeGene (https://www.seegenemedical.com) is a large laboratory that recently started providing RDR-based diagnostic reports for about 200 small hospitals as well as GPs in Korea.

Origins

Seegene' use of RDR resulted from a collaborative project again with Byeong Kang at the University of Tasmania

---

IPMS

IPMS (http://www.stable.co.kr) started in 2010, provides RDR-based diagnostic advice about system alerts and at the time of writing has 100 customers

Origins

IPMS's use of RDR resulted from a collaborative project with Byeong Kang at the University of Tasmania

---

Tapacross

Tapacross (https://www.tapacross.co.kr) provides social media trend prediction services. RDR are used to classify very large number of documents.

Origins

IPMS's use of RDR resulted from a collaborative project with Byeong Kang at the University of Tasmania

---

In Korea, RDR have also been used in some of the internal processes of a very large company, but this is under confidentiality agreements. There may also be other industry use of RDR that we are not aware of, because similar to Erudine the companies do not chose to identify that they use RDR.

Bibliography

Adams, M. (2013). Usability Extension for the Worklet Service. YAWL Symposium.

Adams, M., A. ter Hofstede, D. Edmund and W. van der Aalst (2006). Worklets: a service-oriented implementation of dynamic flexibility in workflows. Proceedings of the 14th International Conference on Cooperative Information Systems (CoopIS'06), Montpellier, Springer-Verlag.

Bender, R., G. Edwards, J. McMahon, A. J. Hooper, G. F. Watts, J. R. Burnett and D. A. J. P. Bell (2016). "Interpretative comments specifically suggesting specialist referral increase the detection of familial hypercholesterolaemia." **48**(5): 463-466.

Compton, P. (2013). "Situated cognition and knowledge acquisition research." International Journal of Human-Computer Studies **71**: 184-190.

Compton, P., R. Horn, R. Quinlan and L. Lazarus (1989). Maintaining an expert system. Applications of Expert Systems. J. R. Quinlan. London, Addison Wesley. **2:** 366-385.

Compton, P., L. Peters, G. Edwards and T. G. Lavers (2006). "Experience with Ripple-Down Rules." Knowledge-Based System Journal **19**(5): 356-362.

Compton, P., L. Peters, T. Lavers and Y.-S. Kim (2011). Experience with long-term knowledge acquisition. Proceedings of the sixth International Conference on Knowledge Capture, KCAP 2011. Banff, Alberta, Canada, ACM**:** 49-56.

Curtain, C., I. Bindoff, J. Westbury and G. Peterson (2013). "An investigation into drug-related problems identifiable by commercial medication review software." The Australasian Medical Journal **6**(4): 183-188.

Dani, M. N., T. A. Faruquie, R. Garg, G. Kothari, M. K. Mohania, K. H. Prasad, L. V. Subramaniam and V. N. Swamy (2010). Knowledge Acquisition Method for Improving Data Quality in Services Engagements. IEEE International Conference on Services Computer (SCC), Miami, IEEE.

Edwards, G., P. Compton, R. Malor, A. Srinivasan and L. Lazarus (1993). "PEIRS: a pathologist maintained expert system for the interpretation of chemical pathology reports." Pathology **25**: 27-34.

Ellis, G. R. (2005). Improved Search Engine, Ivis.

Horn, K., P. J. Compton, L. Lazarus and J. R. Quinlan (1985). "An expert system for the interpretation of thyroid assays in a clinical laboratory." Australian Computer Journal **17**(1): 7-11.

Kang (1996). Multiple Classification Ripple Down Rules (PhD thesis), University of New South Wales.

Kang, B., P. Compton and P. Preston (1995). Multiple Classification Ripple Down Rules : Evaluation and Possibilities. Proceedings of the 9th AAAI-Sponsored Banff Knowledge Acquisition for Knowledge-Based Systems Workshop, Banff, Canada, University of Calgary.

Kang, B., P. Compton and P. Preston (1996). Validating incremental knowledge acquisition for multiple classification. Proceedings of the third world congress on expert systems, Seoul, Cognizant Communications.

Preston, P., G. Edwards and P. Compton (1994). A 2000 Rule Expert System Without a Knowledge Engineer. Proceedings of the 8th AAAI-Sponsored Banff Knowledge Acquisition for Knowledge-Based Systems Workshop, Banff, Canada.

Sarraf, Q. and G. Ellis (2006). "Business Rules in Retail: The Tesco.com Story." Business Rules Journal **7**(6): http://www.BRCommunity.com/a2006/n2014.html.