AFRL-AFOSR-JP-TR-2019-0018

Seamless Integration of Knowledge Acquisition for Autonomous Systems by Domain Users with Prudence Capability

Byeong Ho Kang
UNIVERSITY OF TASMANIA

09/28/2018
Final Report

Air Force Research Laboratory
AF Office Of Scientific Research (AFOSR)/ IOA
Arlington, Virginia 22203
Air Force Materiel Command

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*
15-03-2019

2. REPORT TYPE
Final

3. DATES COVERED *(From - To)*
12 Sep 2016 to 11 Sep 2018

4. TITLE AND SUBTITLE
Seamless Integration of Knowledge Acquisition for Autonomous Systems by Domain Users with Prudence Capability

5a. CONTRACT NUMBER

5b. GRANT NUMBER
FA2386-16-1-4045

5c. PROGRAM ELEMENT NUMBER
61102F

6. AUTHOR(S)
Byeong Ho Kang

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
UNIVERSITY OF TASMANIA
2 CHURCHILL AVE
SANDY BAY, 7005 AU

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
AOARD
UNIT 45002
APO AP 96338-5002

10. SPONSOR/MONITOR'S ACRONYM(S)
AFRL/AFOSR IOA

11. SPONSOR/MONITOR'S REPORT NUMBER(S)
AFRL-AFOSR-JP-TR-2019-0018

12. DISTRIBUTION/AVAILABILITY STATEMENT
A DISTRIBUTION UNLIMITED: PB Public Release

13. SUPPLEMENTARY NOTES

14. ABSTRACT
Knowledge-based systems are typically constrained by their ability to acquire new knowledge, thus limiting their applicability to autonomous systems. This work developed an extensive, easily maintainable hierarchical Knowledge-base System (KBS) for Autonomous Systems (AS) technologies trained by Knowledge Domain Experts (KDE) using a Natural Language (NL) interface for communication. The system implements an abstracted architecture, taking a layer-based approach to separate data and hardware, information, and services, each with an associated, contextual knowledge base. The developed process, Contextual MCRDR, improves upon classical Multiple Classification Ripple Down Rules (MCRDR), with constrained natural language conversation systems associated with querying of in-situ databases of pre-existing information. This was then expanded to support Automatic Speech Recognition (ASR). Finally, the work was extended to a semi-autonomous system (Robotis Turtlebot3). The full effort produced three published journal/conference papers, and two additional papers in the review process.

15. SUBJECT TERMS
Knowledge Acquisition, Autonomous System, Knowledge Base System, Natural Language

16. SECURITY CLASSIFICATION OF: a. REPORT	b. ABSTRACT	c. THIS PAGE	17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
Unclassified	Unclassified	Unclassified	SAR		LIN, ALAN
					19b. TELEPHONE NUMBER *(Include area code)* 315-227-7009

Final Report for AOARD Grant FA2386-16-1-4045

"Seamless Integration of Knowledge Acquisition for Autonomous Systems by Domain Users with Prudence Capability"

09/12/2018

Name of Principal Investigators (PI and Co-PIs):

Byeong Ho Kang, David Herbert, Caren Han

- **e-mail address** : bhkang@utas.edu.au, david.herbert@utas.edu.au, soyeon.han@utas.edu.au
- **Institution**: University of Tasmania
- **Mailing Address**: Discipline of ICT, University of Tasmania, Private Bag 87, Hobart, Tasmania 7001
- **Phone**: +61 03 6226 2919
- **Fax**: +61 3 6226 1824

Period of Performance: 09/12/2016 – 09/11/2018

Abstract: Knowledge-based systems are typically constrained by their ability to acquire new knowledge without the intercession of a technical knowledge engineer. This introduces a fundamental disconnect between the system and the domain expert - even if a knowledge acquisition interface is provided, the domain expert is usually highly constrained in their expressiveness and ability to train the system due to technology-specific implementation.

Knowledge-based systems are also currently limited in their applicability to autonomous systems. The domain expert/knowledge acquisition bottleneck in this paradigm also poses a great challenge to effectively train an autonomous system with new or modified behaviours without considerable effort and implementation change. In the autonomous system-operating environment, there is no consistent model of abstraction (of services/behaviour, information and data) that can be leveraged across different systems and domains.

To address these shortcomings, we will develop an extensive, easily maintainable hierarchical Knowledge-base System (KBS) for Autonomous Systems (AS) technologies that will be trained by Knowledge Domain Experts (KDE) using a Natural Language (NL) interface for communication. The system will implement an abstracted architecture, taking a layer-based approach to separate data and hardware, information, and services, each with an associated, contextual knowledge base.

Introduction and Methodology:

Ripple Down Rules (RDR) are an artificial intelligence classification technology and methodology for the capture and maintenance of knowledge associated with a knowledge-base (the modern term for expert systems), together with a method for classification of knowledge (inference). The application of RDR alleviates issues associated with expert system-style knowledge acquisition and helps to optimize the knowledge engineering process, by allowing the domain expert to incrementally add knowledge to the system and only having to justify new conclusions in a local context (by differentiating a new case against cases already seen). In domains where large quantities of information are already available (and required for the developed KB system), the equivalent standard RDR (or more precisely) MCRDR knowledge-base, that integrates an appropriate number of rules to cover the domain, will suffer from considerable rule bloat and repetition. This rule bloat is of course in comparison to some alternative, which is the main focus of the work to date. MCRDR (Multiple Classification Ripple Down Rules) is an extension of RDR that allows for multiple, concurrent classifications of knowledge (cases) presented to the system.

Phase 1

We focused on modifying MCRDR to support constrained natural language conversation systems associated with querying of in-situ databases of pre-existing information:

- ***Querying*** - *Standard MCRDR* has no capability of referring to external databases for classification – considerable (unnecessary) replication effort is required to generate rules with **static** conclusions that are the equivalent of querying the stored database knowledge. The extra rule bloat here is simply related to the row count returned by a query that would be associated with a given knowledge quanta, for example, "*who are employees over 65*?" is a simple database query, which might return α rows, but this knowledge represented in Standard MCRDR would require α rules that are satisfied by the attribute *age > 65*, whose conclusions are the employee names.

- ***Conversational Context*** - *Standard MCRDR* is not conducive to the conversation paradigm – every "utterance" in a dialog is treated separately and independently of all other utterances, whereas normally conversations continue to flow around *topics*. We refer to this as *conversational context*.

- ***Data retention*** - *Standard MCRDR* does not retain (or "label") data in an utterance, so factual information cannot be used to generalize rules. For example, "*my name is John*", "*what is my name*?" requires two rules, but if we then introduce "*my name is Jane*", "*what is my name*?", we would require two additional rules to cope with the different factual data (*John* versus *Jane*). These four rules could be generalized to two rules if we could retain the value of *name* between inference requests.

- ***Brittleness*** – *Standard MCRDR* does not address brittleness, which is a failure of the system due to users not understanding the system's content and structure. Conversational samples that the system does not understand is an example of this.

The modified MCRDR approach, ***Contextual MCRDR***, overcomes the *Standard MCRDR* limitations by:

1. **Querying** - Including post-inference "deferred" classification results that are parsed for query references, and bound by context variables;

2. **Conversational context** - Adopting a stack-based structure of previous inference results and modifying the inference engine to start from the satisfied rules in the top stack frame (instead of the default root rule). If the inference results include only the default rule, the next stack frame is used as the starting point (and so on);

3. **Data retention** – pattern matching (via regular expressions and ontological lookup) is used to associate attribute values with *context variables*. Context variables are then maintained across inference requests. They are currently used in classification results (either to bind queries, or as their literal value).

4. **Brittleness** – often a problem with knowledge-based systems, we address brittleness by prompting the user with contextual responses indicating example utterances that are recognised according to the current conversational topic. For example, if the current discussion topic is navigation, and the user's last utterance is not recognised, a prompt might be "*You can ask me to turn in a particular direction*".

Phase 2

The spoken interface to the conversational system was initially constrained to in-browser support (the Google Speech API via the Chrome browser). The second phase included development of interfaces for commodity-based *Intelligent Personal Assistances* (IPAs) such as the *Google Home* and *Amazon Echo*, which allowed such devices to be used as the main spoken interface to the conversation system. An evaluation study of the device's Automatic Speech Recognition (ASR) performance then followed to ascertain the best device to use as a speech-enabling interface.

Phase 3

- (*Ongoing*) A semi-autonomous system was purchased (*Robotis' Turtlebot3*). Turtlebot3 is a two-wheeled differential drive type platform that is small, extensible, programmable, ROS-based mobile robot for use in education and research. It contains a robust embedded system for control of its servo motors, a general-purpose single board computer

(SBC), (Raspberry Pi), and a 360-degree laser-distance sensor (which is used for SLAM (Simultaneous Localization And Mapping) and navigation). SLAM is used to build a map of the environment that is then later referred to by navigational tasks.

- Turtlebot3 supports ROS (Robot Operating System) – this is a meta-operating system (or conceptually a middleware layer) that leverages community-written, open source drivers and applications for robotic platforms. This significantly assists in developing a coupled interface between the phase 1 and 2 conversation system and the robotic platform, allowing the hierarchical KBS to drive navigational tasks in both an immediate mode, as well as via learned, hierarchically-stored behaviours. This aspect of phase 3 is the focus of current activity associated with a doctoral thesis.

Results and Discussion:

Phase 1

Contextual MCRDR (C-MCRDR) forms the basis of the user conversational interface being developed for integration with autonomous systems. As part of a doctoral thesis, this component was developed and then evaluated – we chose an educational domain where access to pedagogical databases was readily available and a chatbot conversational system would be a useful tool. The criteria used in this selection were the following:

1. a sizeable aspect of the domain data was readily accessible in a database form;
2. the nature of the domain could be expressed in a question and answer paradigm;
3. a usability study could be conducted to ascertain that the application of the technology was feasible in the domain;
4. the resulting developed system would prove to be a useful tool in the target domain; and
5. scope exists for extensions to the system in the target domain to make it attractive for relevant stakeholders

The developed system, ICS – *Intelligent Conversation System* was written with domain-independence as part of the system philosophy, but the system was evaluated against a domain that reports data associated with *unit outlines*, pedagogical information related to units delivered by the School of Engineering and ICT at the University of Tasmania. Student can use the system (using constrained natural language, via speech or text) to ask questions associated with units they are studying, for example, who is the unit coordinator, what are the learning outcomes, what is the teaching pattern, when are assessment items due and so on. The system was evaluated in the form of a qualitative integrated feedback system that enabled undergraduate participants to rank the system's responses to their queries, as well as rank the usefulness of the entire system as a tool.

The system architecture is shown Figure 1.

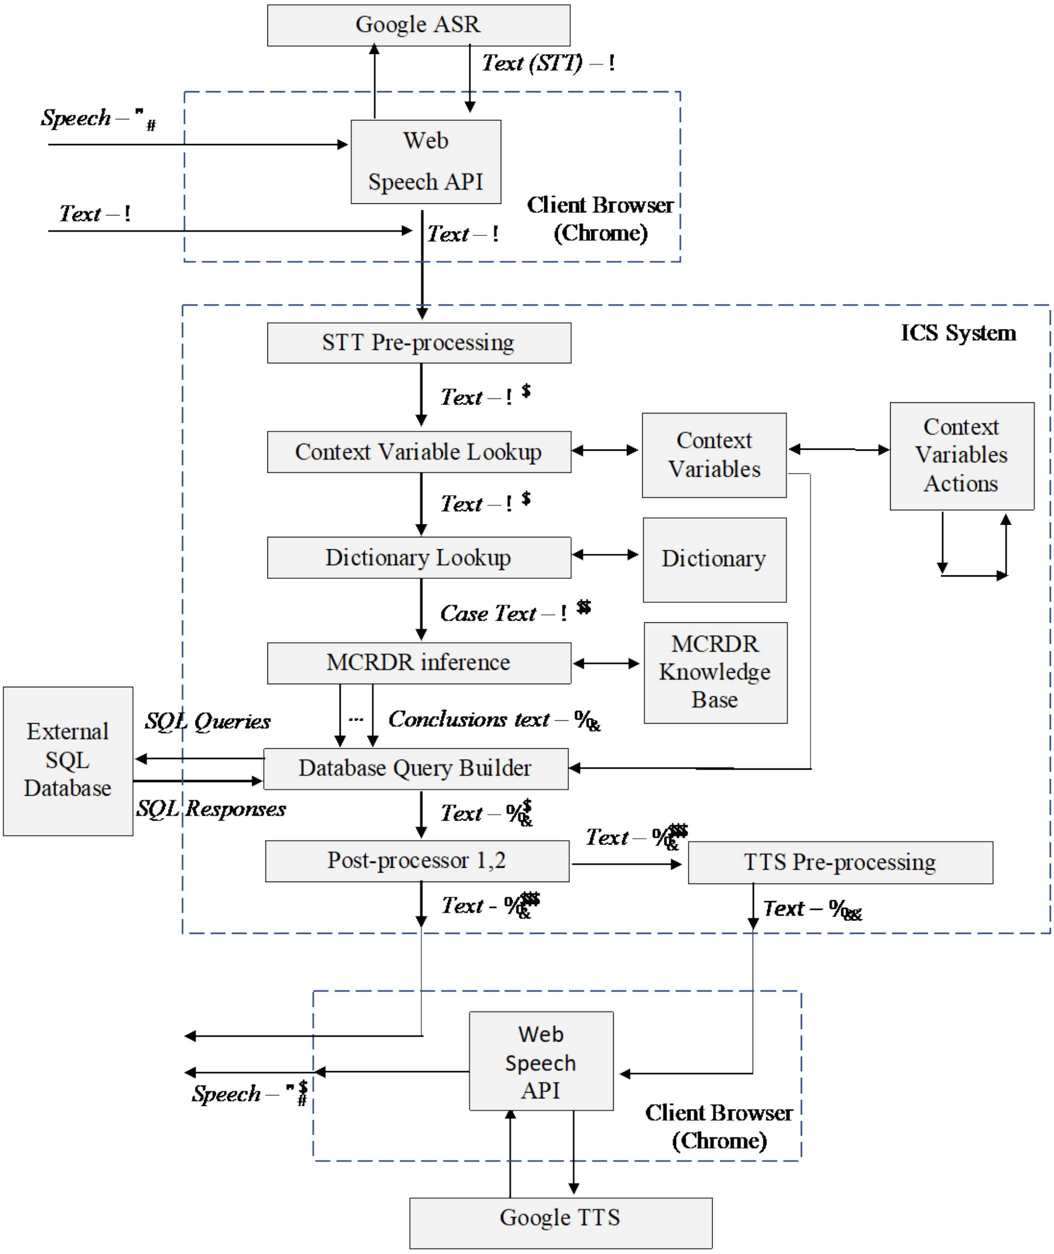

Figure 1 - System Architecture

Evaluation

The system evaluation method utilized two types of feedback – overall system feedback and individual response (rule) feedback. Feedback (along with log file data) was gathered during the first 6 weeks of semester 1, 2017 and analysed. The resulting system feedback was very positive – 61.9% of respondents ranked it at level 4 or 5 (I am satisfied (28.6%) or I am very satisfied (33.3%)) on a 5-point scale. 24% ranked it at an ambivalent level (level 3, I am neither satisfied nor dissatisfied) and the remaining ~14% ranked it at levels 1 or 2 (I am very dissatisfied, or I am dissatisfied) (see Figure 2).

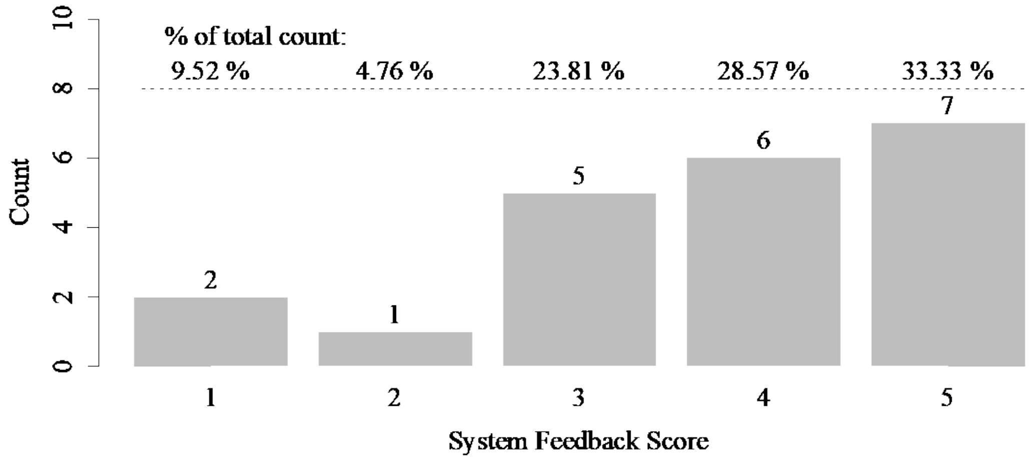

Figure 2 – Overall System Feedback

Individual inference feedback results were grouped into Appropriate System Response and Inappropriate System Response – see Figure 3:

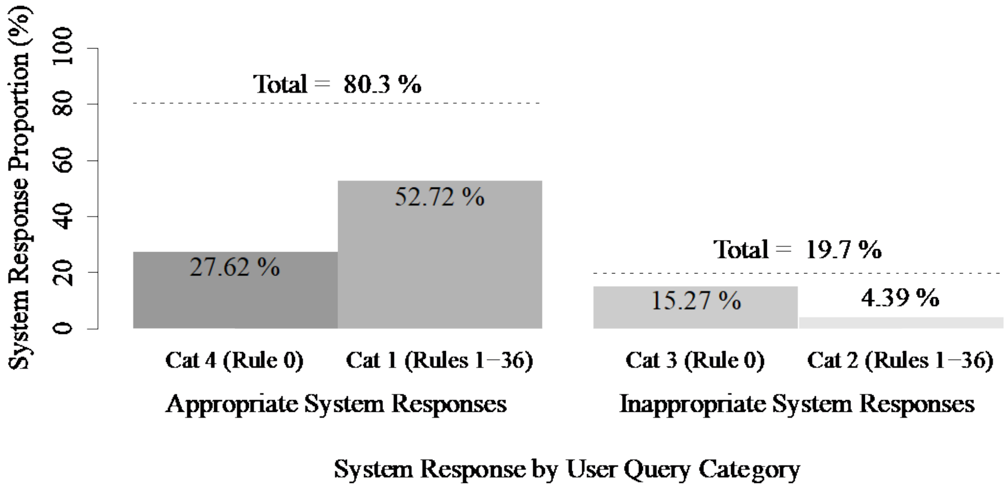

Figure 3 – System responses grouped by category

- ***Appropriate System Response***: the system's inference results are appropriate for the user's question; either the user has posited a valid question and received a valid, non-default rule response, or the user has asked nonsense question and received the default rule response (see Figure 4).

- Positive feedback – 70.22% [score 5, 57.45%; score 4, 12.77%]
- Ambivalent feedback – 11.7% [score 3]
- Negative feedback – 18.09% [score 1, 13.83%; score 2, 4.26%]

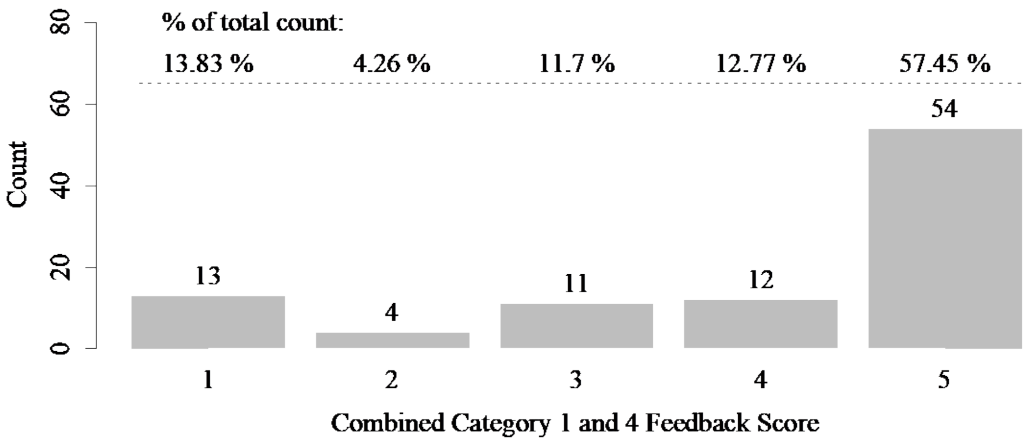

Figure 4 – Feedback scores for appropriate system responses

- ***Inappropriate System Response***: the system's inference results are inappropriate, either the user's question has been misinterpreted (and a non-default rule response is returned), or the user has asked a seemingly valid question and the default rule response is returned. Questions of the former category indicate the user's actual intent have been misconstrued (indicating an insufficient number of rules), and the latter category is mostly due to questions being asked that are completely out of scope (see Figure 5).

    - Positive feedback – 6.06% [score 5, 0%; score 4, 6.06%]
    - Ambivalent feedback – 36.36% [score 3]
    - Negative feedback – 57.57% [score 1, 24.24%; score 2, 33.33%]

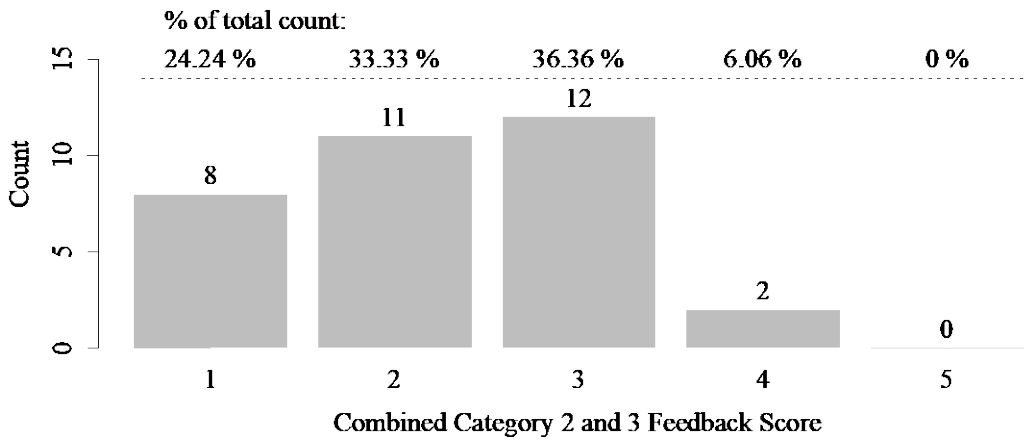

Figure 5 – Feedback scores for inappropriate system responses

- Figure 6 shows the ideal mode of usage – category 1 responses, when the user has asked a valid question and the system has responded with a non-default response. The plot shows category 1 responses against the number of requests made.

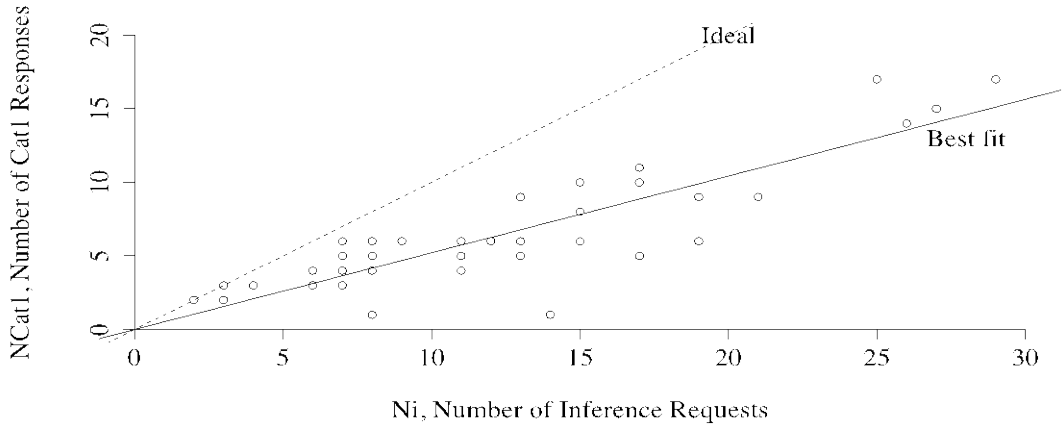

Figure 6 – ideal system response rate with actual rate

Phase 2

A post-ASR (Automatic Speech Recognition) processing scheme to improve the recognition performance of Intelligent Personal Assistants (IPA) was evaluated. We assessed the recognition performance Google Home and Amazon Echo, and their respective virtual agents (Google Assistant and Amazon Alexa). We evaluated the effects of different word and sentence attributes and then improved recognition rates by applying a method for misinterpretation correction when used with the KBS.

Google Home performed appreciably better compared to Amazon Echo in two categories of tests prior to correction:

1. Isolated word evaluation recognition rates:
   Human speaker: Google Home: 91.80%, Amazon Echo: 72.84% (Figure 7)
   Computer speaker: Google Home: 84.52%, Amazon Echo: 49.78% (Figure 9)

2. Phrasal (sentence) evaluation recognition rates:
   Human speaker: Google Home: 69.0%, Amazon Echo: 30.0% (Figure 8)

Google Assistant coupled with a rule-based correction method achieved recognition rates of up to 100% (after a maximum of five rounds of defining correction rules) in the isolated word categories.

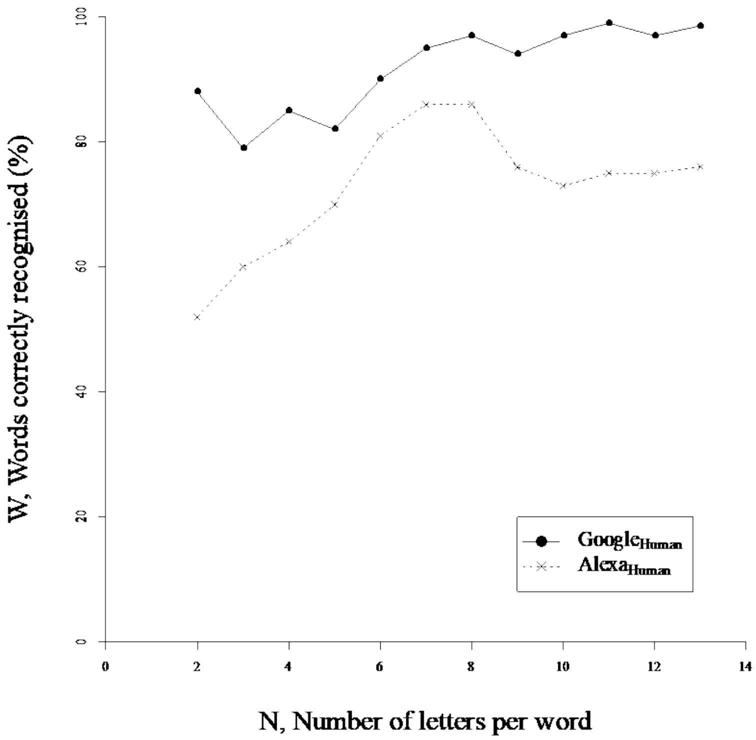

Figure 7- isolated word performance (Human speaker)

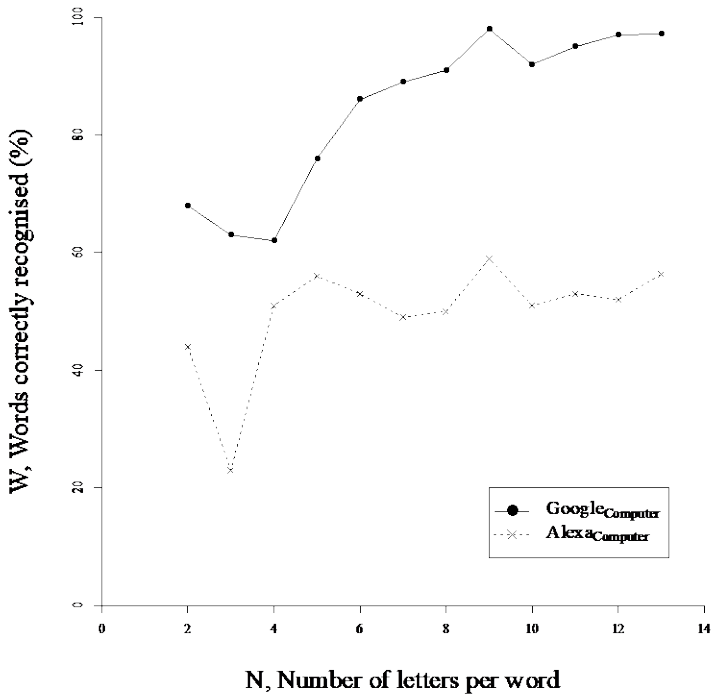

Figure 8 - Sentence recognition rates

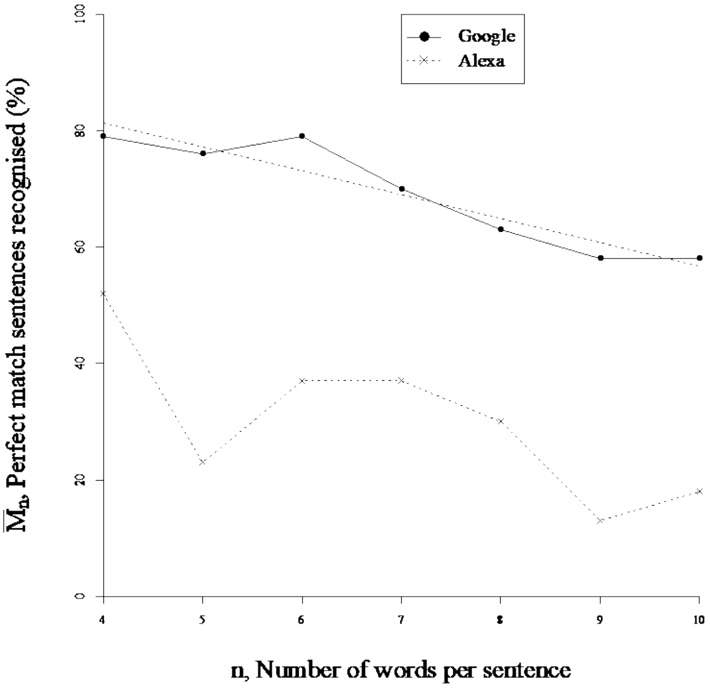

Figure 9 - isolated word performance (computer speaker)

The significant performance separation between the devices indicates the Google Home should be chosen as the speech-enabling interface for the conversation system developed in phase 1.

Phase 3

This is an active phase of the research; software development is currently being undertaken for integration between the Turtlebot3 hardware platform, the ROS operating system, IPA devices (Google Home) and the Enterprise Java-based C-MCRDR conversation system. There will be additional modification of the C-MCRDR inference mechanism to support hierarchical references (for example, deferred classification which refers to further, repeated inference requests for the resolution of hierarchically-based behavioural descriptions that drive the autonomous system).

Research outcomes for FA2386-16-1-4045: papers published

- Herbert, D. and Kang, B.H., 2018. Intelligent conversation system using multiple classification ripple down rules and conversational context. *Expert Systems with Applications*, Vol 112, Pages 342-352, https://doi.org/10.1016/j.eswa.2018.06.049
- Ameen, S., Han, S.C., Lin, Y., Lah, M. and Kang, B.H., 2017. Intelligent medical case based e-learning system. In *Proceedings from the Australasian Conference on Information Systems*. Pages 1-12.
- Lin, Y., Han, S. C., & Kang, B. H. 2018. Machine Learning for the Peer Assessment Credibility. In *Companion of The Web Conference 2018*, Pages 117-118. International World Wide Web Conferences Steering Committee.

Unpublished Manuscripts

- Herbert, D. and Kang, B.H., 2018. Intelligent Personal Agent evaluation and utterance correction coupled to a Knowledge-based Conversational Agent. Submitted to journal: *Expert Systems with Applications* (Elsevier)
- Herbert, D. and Kang, B.H., 2018. Comparative Analysis of Intelligent Personal Agent Performance
  Submitted to AI 2018: The 31th Australasian Joint Conference on Artificial Intelligence