How a Conversational Agent Might Help Farmers in the Field

Bevan Koopman[1*], Guido Zuccon[2*]
Scott Chapman[1,2], Yash Dang[2], David Lawrence[3]
bevan.koopman@csiro.au,g.zuccon@uq.edu.au,s.chapman@uq.edu.au
y.dang@uq.edu.au,david.lawrence@daf.qld.gov.au
[1]Commonwealth Scientific and Industrial Research Organisation (CSIRO), Brisbane, Australia
[2]The University of Queensland, Brisbane, Australia
[3]Department of Agriculture and Fisheries, Queensland Government, Brisbane, Australia

ABSTRACT

In this position paper, we look at how conversational search technology might be used to help farmers meet their increasingly demanding information needs. Agriculture has become more data-driven as a plethora of resources are available for farmers in making growing decisions. However, accessing, interpreting, collating and contextualising all these resources is a real impediment for farmers in the field. We posit that conversational search offers an attractive solution to this problem. In this paper, we categorise the unique information needs of farmers and explain some of the problems and challenges that these create from a search perspective. We show why conversational search offers a unique solution to these problems. The key components of a hypothetical conversational agent that meets farmers information needs are presented. Finally, we highlight how the agricultural domain actually offers an interesting and fruitful playing field for research on conversational search agents and encourage further work in this area.

ACM Reference Format:
Bevan Koopman[1*], Guido Zuccon[2] and Scott Chapman[1,2], Yash Dang[2], David Lawrence[3]. 2020. How a Conversational Agent Might Help Farmers in the Field. In *CAIR '20: AInternational Workshop on Conversational Approaches to Information Retrieval, March 18, 2020, Vancouver, British Columbia, Canada.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/1122445.1122456

1 INTRODUCTION

Conjure up the image of a present day farmer. The cliché maybe be an old man in a straw hat who's knowledge and practice has been past down through generations. This image is shattered by the reality of modern 21st century agriculture: increasingly mechanised, data-driven and based on scientific, evidence-based practice [3, 21, 24]. Some might naturally rebut this by pointing to large populations of subsistence farming in developing countries; however, even these sectors increasingly benefit from digital disruption [10, 18].

*Both authors contributed equally to the work.

CAIR '20, March 18, 2020, Vancouver, British Columbia, Canada

ACM ISBN 978-1-4503-XXXX-X/18/06…$15.00
https://doi.org/10.1145/1122445.1122456

While a wealth of potentially valuable resources and data — from research studies, project reports and communications, through to metrological and soil sample data — could be used by farmers, there is significant barrier in them effectively accessing these resources. Much of these resources are currently locked away as free-text documents, scattered across different repositories, not easily discoverable and synthesised. Thus growers are not able to put into practice these valuable insights.

This barrier is not caused by a lack of access to either the Internet or suitable handheld devices — farmers have these and use them readily; Twitter being one popular platform for farmers to keep informed of the latest trends [16, 22]. Instead, the barrier results from a mismatch between the types of questions farmers would like to ask (e.g., "what's the best way to control pest snails in my local wheat crop") and the resulting documents and resources that, in their current form, hide the answer to such questions.

We posit that this domain represents both an interesting and fruitful area to investigate the use of conversational search approaches. Specifically, how can a conversational agents help answer farmers' questions and make better growing decisions? This position paper:

- Surveys the literature for information technology solutions used in agriculture that involve some form of conversational agent;
- Presents a review of the possible and differing information needs of growers;
- Maps the problems and challenges of agricultural information seeking to research directions in information retrieval;
- Highlights why conversational search is good for agriculture, and why agriculture is a good playing field for conversational search research;
- Sketches the solution space of a conversational search solution that meets growers information needs.

2 CONVERSATIONAL AGENTS IN AGRICULTURE

Despite practitioners having identified conversational agents as a viable means to provide rapid, contextual and personalised agricultural best practice evidence to farmers and growers [3, 21], a limited number of solutions have been proposed and explored. The majority of these solutions have emerged from research and development efforts carried out in India. This attention to the Indian region may be because (1) agriculture represents one of the main production activities in India, (2) farming practitioners in India are often characterised by low literacy, (3) the Indian government has made available agriculture specific data related to queries and conversations from the Kisan Call Center (KCC)[1], a phone helpline service for farmers to consult with agriculture expert advisors about best practice. Examples of conversational agents developed in this context include AgriBot [11] and FarmChat [10]. Agribot was developed to address growers information needs related to weather, market rates, plant protection and government funding opportunities. This conversational agent relies on sentence embeddings (sent2vec [2]) and entity extraction to compute the similarity between a user question and a background of common question-answer pairs. FarmChat [10] is a speech-based conversational system that relies on decision rules and answers manually derived from the KCC data to identify answers, and on the IBM Watson APIs to perform intent identification and dialogue flow management. Much of the attention in FarmChat is on information access in a context of limited literacy and technology expertise, and on information delivery modality (audio vs audio+text). It focuses only on one crop (potato); it does not leverage machine learning for extracting knowledge but instead relies on a manually built knowledge base, which is not scalable and is difficult to maintain and link to information sources.

Besides conversational agents that facilitate the access to agricultural advice, other conversational agents have been developed to aid access to sensor data from devices deployed on a farm. This is the case, for example, of Agronomobot [17], which allows its users to acquire data from sensors deployed on a vineyard via a conversational-style interaction and free-text queries.

Despite the increasing availability of rich data resources for farmers to draw on, there is a dearth of search-based systems that can brings this data together to answer a farmers query. The few examples of actual conversational agents in the agricultural, although limited in scope, showed promise and indicate that a larger effect in this area would be fruitful.

3 INFORMATION NEEDS OF GROWERS

From the literature we summarise the specific information needs of growers. We constrain our analysis to those farmers involved in crop production (i.e., growers) and exclude animal production. While much of the concepts outlined here are relevant to both, animal production includes substantial veterinary content, excluded for the benefit of brevity.

The Kisan Call Center (KCC) provides a large dataset to derive insights into the information needs of growers. Although this data is specific to farmers in India, many of the underlying information needs captured there would be generally applicable in other contexts. An analysis of this dataset showed the top 5 query types were for pest and disease (61%), weather (14%), best practices (7%), fertiliser use (5%), and seeds (4%) [10]. Other surveys of growers show similar rankings of query types: plant protection and disease, marketing, fertiliser and water management, preparation of seedlings and sowing, and harvesting technology [5].

From the literature, [5, 10, 21] some key categories of information needs were identified:

Crop protection

A significant number of grower's questions relate to protecting their crop from diseases or pests, whether for future prevention or because of an existing outbreak. In the latter cases, farmers often describe crop disease via visible symptoms (e.g., "brown spots on the leaves") in order to first identify the diseases and second determine the best course of treatment (e.g., what fungicide to use, including dosage and application instructions). Similarly they may describe pest species (e.g., "2cm black and yellow snail") to determine the relevant pesticide to use. Many queries relate to identifying and eradicating weeds [5]. For all these queries, it is important to point out that the grower's query typically does not contain keywords that match the relevant answer (e.g., the actual pest species name); instead this needs to be inferred from the description of the symptoms / problems.

Best practices

Growers are constantly on the lookout for how they can increase the quantity or quality of their yield as well as reduce their costs or wastage, consequently increasing profitability. Agriculture is constantly evolving with new products and practices; many growers feel that keeping abreast of current best practice is critical [21]. While growers will ask specific questions on a topic when they require information, they also seek out recommendation services that "push" relevant information. For example, the use of Twitter is one common way of keeping abreast of trends [16, 22].

Unbiased Product Recommendations

Growers rely heavily of many agriculture products to run their farms. These can constitute a significant expense for them and as such they would like reliable and trustworthy product recommendations. Recommendations for different types of fertiliser, seed and crop variants and herbicide or pesticide are some commonly sought examples [10].

Markets and Weather/Climate

While the market and weather are factors outside grower's control, they will certainly wish to understand and adapt their practices to changes in both. Because a farm is a business producing agricultural products, it has the same requirements of access to and understanding of markets that all business has. Growers would like to understand and adapt to the market in which they operate [5]. This includes understanding of current and projected prices on products they sell as well as costs of products and services they consume.

Growers would like to take into account the past, current and future weather and climate. Planting, for example, is often tied specifically to periods of rainfall. Similarly, pest outbreaks often relate to weather and climate. Thus growers would want any information returned to be tailored to the recent weather. Similarly, upcoming weather impacts grower's decisions so information should be tailored to weather forecasts. Longer term climate information — both historic and projected — is also important to growers and needs to inform what information is presented to them.

The case for contextualisation

A key requirement of all the information needs outlined so far is that they are contextualised to the specific grower. While personalisation

[1] texthttps://data.gov.in/dataset-group-name/kisan-call-centre

is important in Information Retrieval in general, in this domain it is critical. Few sources of contextualisation include:

Weather & Climate: Information should be tailored to recent weather, forecasted and longer term climatic predictions (e.g., if the farmer is located in a drought predicted area then recommendations for drought resistant crops would be important).

Location: The grower's region strongly informs their information need. The growing conditions, access to markets, infrastructure (e.g., rail or irrigation networks), historical crop yields and many other factors can be inferred from location. Thus, growers would like information that is location-aware.

Markets: Contextualisation to the specific market that the grower operates in, including price, trends and changing customer demands/preferences.

Literacy/Interpretably: Evidence-based agriculture involves making decisions based on scientific evidence and sources. While growers may recognise the value of this, they do not necessarily want to delve into detailed scientific information, or have the expertise to do so. Instead, they would like outcomes of the scientific literature to be provided to them in an understandable, concise and digestible form. Furthermore, grower's expertise varies considerably — some may have detailed technical expertise in certain areas and thus would like to see associated technical details; others may have no technical expertise in the area and require a lay overview. Information should be tailored to different grower's literacy and expertise.

4 PROBLEMS & CHALLENGES OF AGRICULTURE INFORMATION SEEKING

From the information needs identified in the previous section we map these to the problems and challenges from a conversational search perspective.

Scattered Resources

Growers access a wide and varied assortment of information sources [7]. Valuable agriculture data, best practice recommendations and R&D output is currently locked away into large and heterogeneous datasets, including soil and weather data, research project reports, communications and scientific publications [15]. Some are structured data while large amounts are still in the form of natural language. In Australia, for example, key evidence-based agricultural information for growers is disseminated across the Grains Research and Development Corporation (GRDC) project reports, National Variety Trial data[2], GrowNotes[3] and GrowCovers resources (professional magazines), and international scientific publications. This text-based information is not easily discoverable and synthesised. No federated service is in place that offers growers a single entrypoint to search this information. Once these disparate resources are assembled it is not simply sufficient to provide them directly to the grower as an answer — instead, they need to be synthesised into a concise and easily digestible form. There is little value returning a 200 page scientific report, even if it's highly relevant to a grower's query — the relevant conclusions and summary need to be extracted and provided as the result.

[2]https://grdc.com.au/resources-and-publications/apps/national-variety-trials
[3]https://grdc.com.au/resources-and-publications/grownotes

Vocabulary mismatch

The language in grower's queries may differ considerably from that of relevant information they wish to find — this is a vocabulary mismatch problem. Section 3 highlighted how growers may describe, for example, a pest via visible symptoms, with the aim of first identification and finally of deciding upon a course of action (treatment). The issue here is how to automatically infer the semantic association between the language of the query and the language of the relevant information. The level of technical expertise of the growers varies too: some growers may pose queries using specific technical terms that match those of scientific sources they are seeking (e.g., use of Latin species names), while others will use lay or even colloquial terms. A conversational agent needs to handle both.

Contextualisation and Personalisation

Information needs to be provided contextualised to the weather, climate and location, individual preferences, etc. As such, meteorological and climate information needs to be utilised to augment the grower's query. Similarly, location information needs to be used to tailor the answer to the relevant geographical region. When a relevant resource is found, this may not be easily interpretable: it may require a level of understanding beyond that of the grower, or it may require examining a long article, when instead the key information could be summarised into a short answer. Thus information should be tailored to the literacy and technical expertise of the grower. This is particularly important where you may have low literacy users posing questions.

Evaluation

Evaluation of conversational agents is an active and challenging area of research [8]. The agricultural setting has some unique challenges. One is that the user may not know the relevance of the information a system provides until that information is actioned; e.g., they do not know at the time if the system-recommended insecticide works until two weeks after it has been applied. The question then arises of whether we evaluate the outcome of the recommendation not the relevance of the information. A real system might be best actually asking the grower some time later about whether a recommendation worked or not.

5 WHY CONVERSATIONAL SEARCH IS GOOD FOR AGRICULTURE

A conversational search approach has some unique characteristics that lends itself to the outlined agricultural task. Radlinski and Craswell [19] provide a theoretical framework for conversational search and outline a number of properties of conversational systems; these are highlighted in Table 1. Below, we map these to the information needs and challenges in the agricultural domain. To aid the reader, Figure 1 provides a high-level view on how a conversational search agent might help answer a farmer's question.

User revealment helps users express their information need, which in turn helps tackle the vocabulary mismatch problem by improving grower's queries. In fact, user revealment is key advantage of the conversational search approach. Revealment is achieved by both interactions with the system and by the system explicitly

Table 1: Key properties of conversational search agents as defined by Radlinski and Craswell [19].

Property	Description
User Revealment	The system helps the user express (potentially discover) their true information need, and long-term preferences.
System Revealment	The system reveals to the user its capabilities and corpus, building the user's expectations of what it can and can't do.
Mixed Initiative	The system and user both can take initiative as appropriate, support both push and pull models.
Memory	The user can reference past statements, which implicitly also remain true unless contradicted.
Set Retrieval	The system can reason about the utility of sets of complementary items.

asking clarifying questions of the user. Specific techniques for asking clarifying questions exist [1] and align well with the problem of capturing the complex information need of farmers. Finally, interactions with conversational search systems help to reveal user literacy, expertise and other preferences that can all be used to adapt the system to the specific user.

System revealment aids with the issue of scattered resources by providing the user with a sense of what information is out there and how it might be relevant to them.

Growers use of Twitter and other services that exploit recommendation capabilities, showing they favour a mixed push/pull interaction model. This aligns well with the *mixed initiative* property of conversational search. In addition, they may well only know if a system recommendation worked months later — as such a 'prompt' from the system at this stage may be required.

The complex information needs of growers means an interactive system — whereby the user poses a query/question, reads responses, and then refines their question —is the most appropriate. Sometimes these interactions occur over long sessions and time spans. For such cases, the system needs to retain information on past interactions and use these to inform new ones. Thus, the *memory* property of conversational search is desirable. In addition, previous interactions should help personalise the system to heterogeneous users and their preferences. This helps tackle the challenge of contextualisation and personalisation outlined in the previous section.

The *set retrieval* property of conversational search helps tackle the issue of scattered resources in agriculture. Relevant data will come from multiple sources and in many cases be complementary (e.g., pest management and rainfall data for a specific region). To leverage such data, a system needs to reason about what set of data best meets the information need. Some of this reasoning may come through contextualisation.

Conversational search has also been promoted to help with a number of specific types of complex search tasks: facet elicitation, multi-item elicitation, multi-item facet elicitation and bounding choices [19]. We posit these complex tasks epitomise those of growers and outline why below.

In facet elicitation, a user is searching for an item with rich attributes by identifying different facets rather than directly describing the item (because they do not know how to describe it or do not have the technical knowledge to do so). The conversational

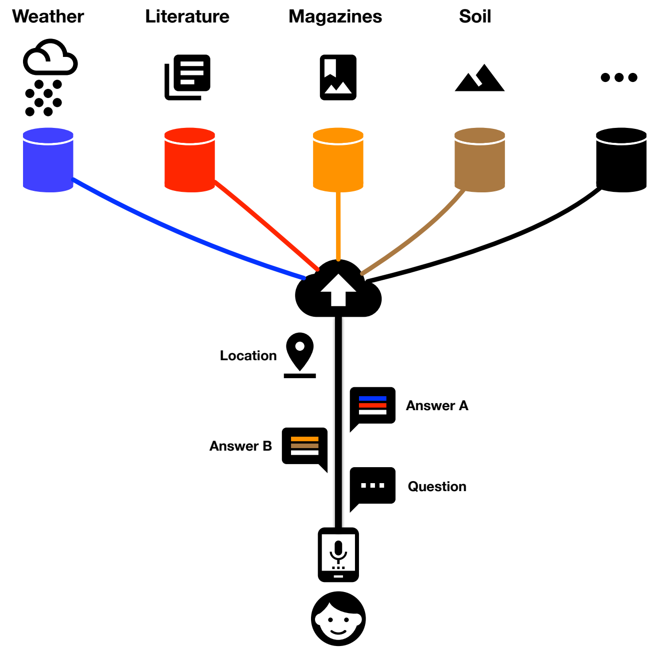

Figure 1: Multiple sources of information need to be collated when providing solutions to farmers. Each solution may rely on a different mix of information sources, e.g. Answer A relies on Weather data and scientific Literature publications, while Answer B on Soil data and Magazines publications.

agent aids the user by helping them understand the different facets, allowing them to iteratively hone in on the relevant item.

In multi-item elicitation, the relevant item can also be found/provided via reference to other items; this is the case when collating multiple agricultural datasets (e.g., rain and pest data for a specific region; or fertiliser recommendation based on soil samples).

Multi-item elicitation represents the complex task of combining the above two tasks: here the system needs to elicit the right information from the user, decide what *set* of items are relevant and return these to the user.

Bounded choices makes it easier for a user to clarify their needs given a set of precise choices rather than expecting them to come up with the particular terms. This may help growers deal with complex agro-scientific resources and distilling them to sets of actions that they may take. Bounded choices also help where the system is uncertain or does not have access to context-specific information to return a single choice (e.g., the system does not know the growers cash flow situation so may suggest three fertilisers at different price points).

Finally, conversational search agents have been a very successful approach for low literacy (and even illiterate) users [10]. Spoken interfaces can provide a mean for low literacy users to pose complex questions and receive detailed responses that would not be possible with textual interfaces.

In summary, the agricultural setting has complex information needs, scattered yet dependent resources and a need for contextualisation. Conversational search offers a number of unique benefits

— from better understanding the user and their query, to collating scattered resources — to tackle these problems.

6 WHY AGRICULTURE IS GOOD FOR CONVERSATIONAL SEARCH

The agricultural domain outlined in this paper, we argue, presents an interesting and fruitful playing field for research in conversational search. It presents a real-world use case of where a conversational agent can be deployed and have a real impact on an industry with wide societal importance. Some of the problems and challenges of the agricultural domain are common across other domain-specific search settings, e.g., vocabulary mismatch and scattered resources. Thus, it provides the opportunity to test and generalise methods across domains, e.g., methods developed in health search to deal with semantic mismatch could be adapted and generalised to agriculture. However, it also presents some exemplary challenges ripe for new research: we detail some of these below.

The outlined information needs and challenges highlight the importance of contextualisation and personalisation. While this is desirable for all domains, it is critical for the agricultural domain. Furthermore, the contextualisation is complex, pulling together many disparate sources (weather, location, soil, past interaction, literacy, etc.). From a conversational Information Retrieval research perspective, there is an opportunity here to develop far richer models that support such contextualisation.

A conversational agent in the agricultural domain needs to draw upon many, multimodal sources of data: geographic images, weather time series, chemical/soil reports, publications in free-text, images of diseases or pests. Some of this data may not necessarily be searched directly, but instead be essential for the contextualisation of the grower's queries, e.g., weather data. This presents an excellent opportunity for new research into multimodal conversational search, including search across heterogeneous datasets but also multimodal contextual models.

Evaluation of conversational agents is an open and challenging area of research [8, 13]. The user- and context-specific nature of intelligent assistants makes it difficult to objectively define a gold standard output for a given situation [13]. Research has called for more focus on inferring user satisfaction using a combination of signals derived from visual data (e.g., facial expressions), gesture-based interactions (e.g., touch paths, touch density, and touch velocity), and voice signals. In this domain, the decision made from the information (e.g., apply a particular fertiliser) along with the outcome that decision produced (e.g., an increase in yield after a certain amount of time) may be the ultimate assessment of success. There is scope for new research on developing novel methods to capture both implicit and explicit user feedback on system recommendations, the associated decisions and their impact or outcome. Conversational search provides an ideal setting for this to take place because the agent could pro-actively follow up on both the decision on the recommendation and the outcome adopting that recommendation has produced over time [19].

7 A PROPOSED SOLUTION

From the analysis of growers information needs and the challenges of the domain outlined so far, we now briefly sketch the key components of an hypothetical conversational search solution. The system, dubbed AgAsk, would provide personalised access to valuable agricultural best-practice information to drive better, data-driven growing decisions. Implemented as a machine learning driven question-answering system, AgAsk would elicit and understand grower's information needs and preferences, providing contextualised insights in agriculture R&D to flow directly to growers, something that is not possible at large scale with current practices. AgAsk would also collect and analyse insightful information about growers, their pressing needs and what they access, giving insights into grower's learning preferences and needs, uptake of specific resources, what worked for them and what did not, decision drivers and barriers to adoption.

AgAsk would require a natural language understanding component capable of (1) extracting and representing information from scattered resources; (2) understanding and matching questions with relevant information (i.e., ranking), (3) formulating an understandable, evidence-based answer to a grower's question; and (4) accounting for user feedback and context to personalise answers and interactions. A number of key components of the proposed system are detailed below.

Knowledge Graph Construction and Reasoning

Agricultural information resources would be mined from textual information and converted into a knowledge graph capturing key agricultural concepts and relations. These knowledge graphs would be akin to, but domain-specific, to those used in general web search [6], and other domain-specific applications such as health search [14]. For example, by mining the scientific report titled "*Ciliate Protozoa In Baits For The Control Of Grain Pest Molluscs*" the algorithms would identify the entities `protozoa` and `pest molluscs` and the relation `control_of` between the first and the second entity (e.g., `protozoa —control_of→ pest molluscs`). AgAsk would use this knowledge graph to formulate contextualised and interpretable answers to a grower's question, e.g., the question "how to deal with slugs in Darling Downs wheat crop?".

Query Understanding and Matching

A query understanding model would take a user's question and convert it into a structured, machine readable query. Given this structured query, the Query Matching module would find suitable related entities from the knowledge graph. Learning to rank models could be employed to find a series of different answers for the grower's questions. The provision of methods for the conversational agent to ask clarification questions, e.g. [1], would further disambiguate the grower's requests and improve query understanding. This would provide a means of returning to the grower a range of different possible answers, important when there is no clear answer or conflicting evidence. Search results diversification techniques could also be relied upon in these situations [20].

Answer Construction

Search and machine learning algorithms would be constructed to augment the candidate entities identified by the matching component with human digestible answers, including identifying the relevant information sources to present as evidence, e.g., [25]. The answer would include references to the source evidence to aid explainability and allowing the grower to dive deeper into material relevant to the question. Where multiple or conflicting answers exist for the question, these multiple answers would be delivered in an effective and understandable format to the user; e.g., in a form similar to entity cards which are used in web search to summarise the key information about an entity in a human readable format [4, 12, 23].

Personalisation and Feedback

The conversational agent would use the knowledge graph to formulate contextualised and interpretable answers to a grower's question. Thus the system would take into account user location, crop variety, seasonal variants (rainfall, temperature, etc.) and previous interactions/questions. By soliciting users feedback on valuable answers, the agent would adopt an online learning methodology [9], thus continuously improving and adapting to an ever changing agricultural environment and grower's needs.

Analysis Platform

AgAsk would collect and analyse, at scale, insightful information about growers, their pressing needs, what they access, what worked for them and what did not. This provides real-time feedback on grower's learning preferences and needs, uptake of specific agriculture resources, decision drivers and barriers to adoption.

Agriculture-specific conversational agents are limited. The proposed solution offers both novel avenues of research and a means of meeting the unfulfilled information needs of growers.

8 CONCLUSION

A survey of literature highlighted that there are good opportunities to apply conversational agents in the agricultural domain. Only few existing systems have been developed despite research showing the potential.

Research has, however, provided extensive insight into grower's information needs. We summarise the different information needs according to the high level categories of crop protection, best practices, markets and climate. From this, the importance of contextualisation and personalisation becomes apparent.

From the information needs, a number of problems and challenges from a search perspective are derived. These included harmonising scattered resources, the vocabulary mismatch problem between a grower's query and the available resources, contextualisation/personalisation and how to perform evaluation. Conversational search offers a number of solutions to these problems, as we show by analysing specific properties of a theoretical framework for conversational search [19].

AgAsk, a hypothetical conversational search system, is detailed, including the major working components, which meet the requirements for grower's information needs.

Finally, we highlight how the agricultural domain offers an interesting test bed for research on conversational agents, with key focus on better contextualisation models, multimodal search, and evaluation of conversational agents. It is our aim to both foster more research in this area and to translate research into real-world systems deployed in the field.

ACKNOWLEDGMENTS

This work was funded by a Grains Research & Development Corporation grant.

REFERENCES

[1] Mohammad Aliannejadi, Hamed Zamani, Fabio Crestani, and W Bruce Croft. 2019. Asking clarifying questions in open-domain information-seeking conversations. In *SIGIR*. 475–484.
[2] Sanjeev Arora, Yingyu Liang, and Tengyu Ma. 2017. A simple but tough-to-beat baseline for sentence embeddings. In *ICLR*.
[3] Manlio Bacco, Paolo Barsocchi, Erina Ferro, Alberto Gotta, and Massimiliano Ruggeri. 2019. The Digitisation of Agriculture: a Survey of Research Activities on Smart Farming. *Array* 3 (2019), 100009.
[4] Horatiu Bota, Ke Zhou, and Joemon M Jose. 2016. Playing your cards right: The effect of entity cards on search behaviour and workload. In *CHIIR*. 131–140.
[5] Nikulsinh M Chauhan et al. 2012. Information hungers of the rice growers. *Agriculture update* 7, 1/2 (2012), 72–75.
[6] Xin Dong, Evgeniy Gabrilovich, Geremy Heitz, Wilko Horn, Ni Lao, Kevin Murphy, Thomas Strohmann, Shaohua Sun, and Wei Zhang. 2014. Knowledge vault: A web-scale approach to probabilistic knowledge fusion. In *SIGKDD*. 601–610.
[7] Pravin C Gedam and RN Padaria. 2016. Information Needs of Orange Growers of Maharashtra. *Indian Research Journal of Extension Education* 14, 1 (2016), 99–101.
[8] Marco Guerini, Sara Falcone, and Bernardo Magnini. 2018. A Methodology for Evaluating Interaction Strategies of Task-Oriented Conversational Agents. In *EMNLP Workshop SCAI*. 24–32.
[9] Katja Hofmann, Shimon Whiteson, and Maarten de Rijke. 2013. Balancing exploration and exploitation in listwise and pairwise online learning to rank for information retrieval. *Information Retrieval* 16, 1 (2013), 63–90.
[10] Mohit Jain, Pratyush Kumar, Ishita Bhansali, Q. Vera Liao, Khai Truong, and Shwetak Patel. 2018. FarmChat: A Conversational Agent to Answer Farmer Queries. *ACM Interact. Mob. Wearable Ubiquitous Technol.* 2, 4, Article 170 (2018).
[11] Naman Jain, Pranjali Jain, Pratik Kayal, Jayakrishna Sahit, Soham Pachpande, Jayesh Choudhari, et al. 2019. AgriBot: agriculture-specific question answer system. *IndiaRxiv* (2019).
[12] Jimmy, Guido Zuccon, Gianluca Demartini, and Bevan Koopman. 2019. Health Cards to Assist Decision Making in Consumer Health Search. In *AMIA*. Washington D.C., USA.
[13] Julia Kiseleva and Maarten de Rijke. 2017. Evaluating personal assistants on mobile devices. In *SIGIR CAIR*.
[14] Bevan Koopman, Guido Zuccon, Peter Bruza, Laurianne Sitbon, and Michael Lawley. 2015. Information Retrieval as Semantic Inference: A Graph Inference Model applied to Medical Search. *Information Retrieval* 19, 1 (2015), 6–37.
[15] Rob Lokers, Rob Knapen, Sander Janssen, Yke van Randen, and Jacques Jansen. 2016. Analysis of Big Data technologies for use in agro-environmental science. *Environmental modelling & software* 84 (2016), 494–504.
[16] Jane Mills, Matthew Reed, Kamilla Skaalsveen, and Julie Ingram. 2019. The use of Twitter for knowledge exchange on sustainable soil management. *Soil Use and Management* 35, 1 (2019), 195–203.
[17] Gustavo Marques Mostao, 'ICaro Ramires Costa De Souza, Leonardo Barreto Campos, and Carlos Eduardo Cugnasca. 2018. AgronomoBot: a smart answering Chatbot applied to agricultural sensor networks. In *Conference on Precision Agriculture*, Vol. 24. 1–13.
[18] Kweku Opoku-Agyemang, Bhaumik Shah, and Tapan S. Parikh. 2017. Scaling Up Peer Education with Farmers in India. In *Information and Communication Technologies and Development (ICTD '17)*. ACM, 15:1–15:10.
[19] Filip Radlinski and Nick Craswell. 2017. A theoretical framework for conversational search. In *CHIIR*. 117–126.
[20] Rodrygo LT Santos, Craig Macdonald, Iadh Ounis, et al. 2015. Search result diversification. *FTIR* 9, 1 (2015), 1–90.
[21] Matthew J Smith. 2018. Getting value from artificial intelligence in agriculture. *Animal Production Science* (2018).
[22] Sarah Van Dalsem. 2011. *An iPhone in a haystack: the uses and gratifications behind farmers using Twitter.* Master's thesis. University of Nebraska.
[23] Anton van der Vegt, Guido Zuccon, Bevan Koopman, and Peter Bruza. 2018. A task completion framework to support single-interaction IR research. *Journal of Documentation* 74, 2 (2018), 289–308.

[24] Jim Virgona, Geoff Daniel, et al. 2011. Evidence-based Agriculture-Can We Get There? *Agricultural Science* 23, 1 (2011), 19.

[25] Liu Yang, Minghui Qiu, Chen Qu, Jiafeng Guo, Yongfeng Zhang, W Bruce Croft, Jun Huang, and Haiqing Chen. 2018. Response ranking with deep matching networks and external knowledge in information-seeking conversation systems. In *SIGIR*. ACM, 245–254.