Should I Interfere? AI-Assistants' Interaction with Knowledge Workers: A Case Study in the Oil & Gas Industry

Juliana Jansen Ferreira
IBM Research
Rio de Janeiro, RJ, Brazil
jjansen@br.ibm.com

Ana Fucs
IBM Research
Rio de Janeiro, RJ, Brazil
anafucs@br.ibm.com

Vinícius Segura
IBM Research
Rio de Janeiro, RJ, Brazil
vboas@br.ibm.com

ABSTRACT

Artificial Intelligence (AI) assistants have been a hot topic for a few years. Popular solutions – such as Google Assistant, Microsoft's Cortana, Apple's Siri, and Amazon Alexa – are becoming resourceful AI-assistants for general users. Apart from some mishaps, those assistants have a successful history in supporting people's everyday tasks. The same cannot be said in industry-specific scenarios, in which AI-assistants are still a bet. Companies combining AI with human expertise and experience can be stand out in their industry. This is particularly important for industries that rely their strategic decision-making processes on knowledge workers actions. More than another system, AI-assistants are new players in the human-computer interaction. But how and when should an AI-assistant interfere in a knowledge worker task? In this paper, we present findings from a case study using the Wizard of Oz approach in an oil and gas company. Our findings begin to answer that question for what kind of interference knowledge workers in that domain would accept from an AI-assistant.

CHI'19 Extended Abstracts, May 4-9, 2019, Glasgow, Scotland, UK.

ACM ISBN 978-1-4503-5971-9/19/05. DOI: https://doi.org/10.1145/3290607.3299052

CCS CONCEPTS

• **Human-centered computing** → User studies; Field studies
• **Human-centered computing** → Field studies

KEYWORDS: AI-a; domain expert; knowledge worker; wizard of oz; user study

ACM Reference format:
Juliana Jansen Ferreira, Ana Fucs, Vinícius Segura. 2019. Should I Interfere? AI-a Interactions with Knowledge Workers: A Case Study in the Oil and Gas Industry. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts (CHI'19 Extended Abstracts), May 4–9, 2019, Glasgow, Scotland, UK.* ACM, New York, NY, USA. 8 pages. https://doi.org/10.1145/3290607.3299052

1 INTRODUCTION

The undeniable and growing significance of Artificial Intelligence (AI) technology presents a new set of concerns and challenges for the human-computer interface community. Those concerns and challenges are related to direct effects of AI technologies on people's lives. Human cognitive processes and computers can now be coupled very tightly, and the resulting partnership will present new ways for the human brain to think and computers to process data. AI-as (AI-a) are now real social actors [10] as they are active players in the interaction with people. They can learn and change their behavior while they interact with users. It presents new experiences and challenges not only for users but also for technology developers [3]. As the developers cannot foresee all the possible scenarios, the UX design becomes even more important in the communication between people and AI-as.

The investigation of AI-as (especially conversational agents or chatbots) interaction with users is presenting interesting results related to general tasks, like add an appointment to a calendar, call another person, shop online, and so on [4][8][12] using mainly voice commands. That kind of AI-as aims to help a lot of different people [1][14], even people with disabilities [5]. But when we consider a specific user – a knowledge worker, an expert in a given domain or industry, with a lot of expertise on a set of tasks and daily dealing with tacit knowledge – the AI-a interaction planning and design becomes even more complex.

Knowledge workers are people that work based on what they know, their previous experience, and personal tacit knowledge. Their main tasks are, for example, acquiring, searching, analyzing, and creating knowledge, as well as organizing, planning, and making decisions. Environmental and context aspects impact on the cognition and decision making of knowledge workers [9]. The tacit component of knowledge, which is highly embedded in individuals, is a key source of competitive advantage [2][14]. The partnership between AI-as and those knowledge workers can combine what each one does best: AI dealing with a lot of data, finding patterns and relations; and the knowledge worker teaching the assistant what they know about their industry domain.

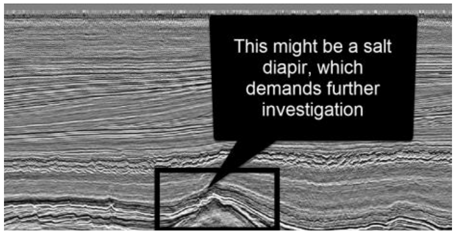

Figure 1: Seismic image example (Netherland – Central Graben – inline 474)

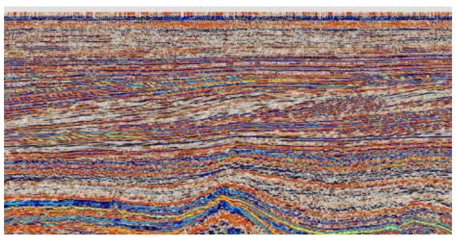

Figure 2: Same seismic line as Figure 1 with different visualization characteristics.

For designing an AI-a interaction with knowledge workers, we need to investigate when, how, and for what purpose AI-as should interfere with knowledge workers activities [13] For that, we designed and executed an experiment using the Wizard of Oz technique [6][7] to identify what kind of AI-a's interference are considered relevant by a group of knowledge workers. Our case study was executed in an oil and gas industry, in the context of a knowledge intensive process called seismic interpretation. Our initial findings helped us guide the design of our AI-a interaction with our knowledge workers, the interpreters.

2 OIL AND GAS CONTEXT OF CASE STUDY

The main practice we are investigating is the seismic interpretation, which is a central process in the oil and gas exploration industry. This practice main goal is to support other decision-making processes by reducing uncertainty. To achieve that goal, different people work alone and engage in multiple informal interactions and formal collaboration sessions, embedding biases, decisions, and reputation. Seismic interpretation is the process of inferring the geology of a region at some depth from the processed seismic survey[2] . **Figure 1** and **Figure 2** show examples of seismic data lines (or slices), which is a portion of a seismic survey.

The main knowledge workers in seismic interpretation practices are the interpreters, which play strategic roles in the decision-making processes. They combine formal knowledge from books, scientific papers, reports from previous projects, and their tacit knowledge to define where and when a company should invest to discover and produce oil and gas. The main process in this studies context is the seismic interpretation process.

The oil and gas industry is very sensitive about their data. Therefore, all images used to illustrate the seismic interpretation process in this paper were taken from a free and public data set of the Netherlands Offshore F3 Block in Open Seismic Repository[3].

3 WIZARD OF OZ METHOD

We used the Wizard of Oz (Woz) approach [6][7] to evaluate interference content of our AI-Assistant "CoRgI" [11] for the oil and gas industry. One researcher played the role of the wizard to interact as our AI-a with users. We selected the Woz technique because it can provide valuable information on which to base our future designs. It presents characteristics we need for this research like allow to gather actual human responses about the non-existent interaction and test design of feedback through output technologies [7].

---

[2] https://www.britannica.com/science/seismic-survey
[3] https://opendtect.org/osr/

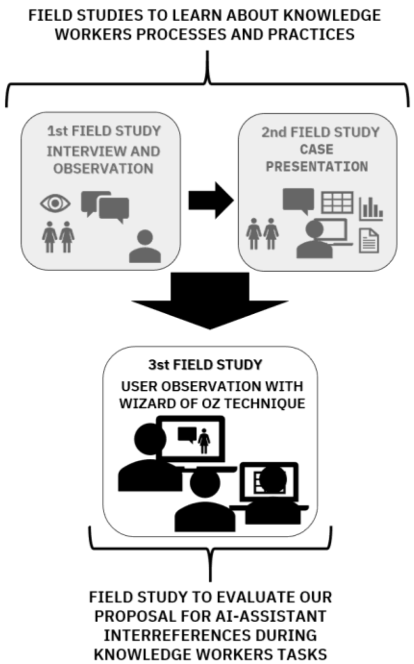

Figure 3: Field Studies results as input for evaluation with Woz.

We had two previous field studies (Figure 3) where we explored and investigated the knowledge workers' practices, workplace, data sources, and other inputs to understand the nature of the seismic interpretation process to propose an AI-a to empower interpreters on the interpretation process. In the first field study we were able to get acquainted with workflows, semantic terms and categories, tools, people interaction, and other practices that could provide us with input and insights on how an AI technology could help improve knowledge-intensive activities. In the second field study, participants presented known cases, where the knowledge workers presented the main events of those cases. They talked about the source of data used, key decision moments, pain points, common practices, other knowledge workers involved, and so on. We have five cases mapped, which gave us enough knowledge to propose our AI-a's moments of interference with the knowledge workers' tasks.

We decided to evaluate eleven interferences that CoRgI could do while users were interacting with a seismic interpretation software. We had the collaboration of seven participants. The moments and content of the interferences were also based on what we learned in the previous field studies. We learned the following facts about interpreters:

- Interpreters start by doing individual interpretations before they share their findings with others,
- Interpreters are very sensitive about sharing their on-going work with others,
- Interpreters collaborate with more experienced colleagues in peer reviews,
- Previous interpretations in a given context, from the same interpreter or others, may provide more knowledge about that context,
- Interpreters look for formal references such as published papers and books to substantiate their interpretations,
- They may not know about every person's previous experience, and
- They may not know about every person's current context of seismic investigation.

Based on what we learned about interpreters' practices, we defined a list of possible interference (Table 1) related to a few common tasks interpreters perform during the interpretation process. The tasks we selected were a) draw a polygon in a seismic data, b) run visual recognition algorithms using a drawn polygon as reference, c) associate a paper to a seismic data, d) open a seismic data to interpret, and e) associate people to a project.

There were two researchers in the room with the users: one was guiding the study and the other was observing the user's interaction to send the interferences in the moments where we intended in the CoRgI's interaction design (the wizard). The researchers were well known by users, since this study is part of a larger research project, and they did not notice that the Woz technique had been applied. At the end of the sessions, we explained to each user the study goals and why and how we used the Woz technique.

Table 1: CoRgI's interferences for Woz study.

	Interference Content
1	*"Other people have experience in this basin as Basin Modeler. Do you want to list them?"*
2	*"John Doe worked on 4 projects in this basin. Do you want to see his public history of investigations?"*
3	*"There are 12 investigations in this cube: 7 public and 5 private. Do you want to know more about them?"*
4	*"There is a Mini Basin in this seismic cube starting around Inline 9500. I suggest you analyze it using XPTO algorithms."*
5	*"There are 8 investigations with similar selected area: 3 public and 5 private. Do you want to know more about them?"*
6	*"It seems you are looking for X structures. There are 5 investigations in this cube where X structures were annotated. Do you want to see them?"*
7	*"It seems that you are looking for Y. There are 9 investigations in this cube that show Y evidence. Do you want to see them?"*
8	*"It seems you are looking for W. There are 3 papers in other projects related to it. Do you want to see them?"*
9	*"This paper is also associated to Project A. Project A is a private project, but you can reach out to its project owner Ana Fucs for more details."*
10	*"This paper is also associated to Project A. Do you want to explore Project A information?"*
11	*"The green region appears to show X structure. There is another result showing X in a different investigation. Do you want to analyze it?"*

The data from our user studies were collected by capturing users' screen while interacting with the interpretation software. A video (with audio) from each user's study was recorded, with user's formal consent. In this way, we also had the user's verbalization about the interaction, interferences and any feedback they could provide regarding the CoRgI's design.

4 FINDINGS

Our goal was to assess how interpreters, knowledge workers from the oil and gas industry, would react to the interference of an AI-a in a defined set of their work practices. We executed the study with five interpreters. The users provided feedback to researcher in three ways: 1) verbally talking about the interference in a given task, 2) accepting the CoRgI's suggestion about the task; and 3) rating CoRgI's interference (Figure 4).

As we mentioned before, we designed the interferences considering the context of a few common tasks interpreters perform during the interpretation process. After analyzing the study data, we organized the interferences in three classes: a) Linked by **Data:** Data related to the interaction context as link to AI-a interference, b) Linked by **Knowledge:** Knowledge created in the interaction context as link to AI-a interference and, 3) Linked by **People:** People with experience in the same interaction context as link to AI-a interference. (see Table 2).

This was an initial test to assess users' feedback regarding the relevance of the designed interferences. Therefore, we classified the participants' reactions into three categories: "relevant", "relevant, but", and "irrelevant". All interferences in the "data" and "people" classes were defined as "relevant" by all users (P4 was not able to respond about interference #9 – linked by data class - and did not gave researchers any more feedback). The linked by knowledge class was the class with different feedbacks (**Figure 5**).

An example of a **"relevant"** feedback was provided by participant 3 (P3) related to interference #4: **P3**: This is related to other investigations other people made in the same region? / **Researcher**: Yes / **P3**: It is useful… if someone already worked on this area, I don't need to re-do some of the work.

Related to the **"relevant, but"** feedback, we had the following feedback provided by participant 4 (P4): **P4**: Are those investigations in the same project that I am in? / **Researcher:** Yes…do you think they are relevant? / **P4:** Yes and no… imagine that I am trying to map salt structures and I have already notice some… this is already done… it is asking if I want to see the previous ones… If I did a little while ago, I do not want to see it again now…

For the **"irrelevant"** classification, participants 3 and 4 (P3 and P4) provided strong feedbacks.

They both indicated a negative reaction related to the interference, but what they meant was that the context of the AI-a interference content was incorrect: ***P3:*** *You are wrong, this is not even chaotic… /* ***P4:*** *Chaotic structures? No, not here… /*

Table 2: Classes of interference

Class of interference	Interferences (see Table 1)
Linked by **Data**	3, 9, 10
Linked by **Knowledge**	4, 5, 6, 7, 8, 11
Linked by **People**	1, 2

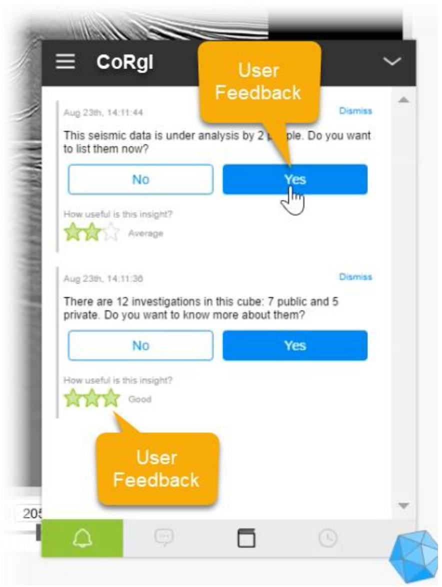

Figure 4: CoRgI's interface

We had positive reaction from users related to contextualized interferences from CoRgI. The context in which user are interacting has a strong linking power, regardless the link used (data, knowledge, or people). Some users' feedback quotes are about the importance of interaction context: **P4 about interference #4:** Ok, this is useful if I am studying mini basins... we must guarantee that it (CoRgI) knows what I am looking for. If not, this (interference) is something loose... / **P5 about interference #11:** Yes, it does make sense... you do this kind of analysis and finds other similar analysis...same project, same seismic...

Regarding individual (private) and public knowledge, the context in which tasks are executed was also a subject for discussion: **P1 about interference #3**: … sometimes people look at seismic for reasons that may not be important…

5 DISCUSSION ABOUT FINDINGS

The "Linked by knowledge" interference class presented the most diverse results from user's feedback in this study, but we need to investigate further. We are aware that it was an initial study with few participants and we need more data and user studies to say more about our designed interferences for those users.

We collected interesting feedbacks related to design options for our AI-a that support knowledge workers that collaborate with each other. For example, interpreters may want to start from someone else's interpretation. One interpreter got to point C by passing through points A and B. Other interpreters may need to get to point E, but they can start from the work the other interpreter did to get to point C and speed up their work.

The interferences that got a "relevant, but…" feedback have a clear necessity to further investigation. The participants gave us examples of situations where the interference could not work, but the task context needs to be considered for those cases. We need to get back and discuss with those same participants about those cases and get new users feedback to check if the conditional relevance of interferences is pointed by more users.

In this study, we noticed a difficulty that participants had from distinguishing the relevance of the interference from the content of the interference. An AI-a has a learning curve; the user needs to teach the assistant about the context they are working. The concept of an AI-a is new in the oil and gas industry, therefore, this confusion about the relevance of the interference in a given point in the task and its content may be associated to this novel aspect of the technology in the industry practice.

6 CONCLUSION & WHAT'S NEXT

This case study goal was to investigate when, how, and why (for what purpose) AI-as should interfere with knowledge workers activities in the oil and gas industry. We observed that the whole AI system scenario is very new to our users and we believe that it may had a considerable impact in our assessment of an AI-a interferences. Therefore, we need to take this discovery as a premise for future studies.

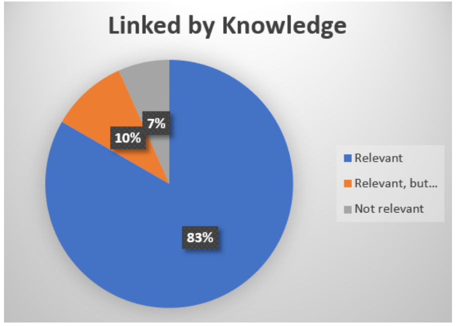

Figure 5: Pie Chart of users' feedback to "liked by knowledge" classes interferences

The context of knowledge workers' practices is very anchored in the investigated context. Since we used real, but general, scenarios for this study, some users did not have the tacit knowledge or even the motivation to engage with the AI-a. We believe that a combined user study considering an interesting scenario for the user can provide us with more insightful feedback for future designs.

REFERENCES

[1] Heloisa Candello, Claudio Pinhanez, and Flavio Figueiredo. 2017. Typefaces and the Perception of Humanness in Natural Language Chatbots. In Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems - CHI '17, 3476–3487. https://doi.org/10.1145/3025453.3025919

[2] Russell W Coff. 1997. Human assets and management dilemmas: Coping with hazards on the road to resource-based theory. Academy of management review 22, 2: 374–402.

[3] Charles Hill, Rachel Bellamy, Thomas Erickson, and Margaret Burnett. 2016. Trials and tribulations of developers of intelligent systems: A field study. In Visual Languages and Human-Centric Computing (VL/HCC), 2016 IEEE Symposium on, 162–170.

[4] Mohit Jain, Ramachandra Kota, Pratyush Kumar, and Shwetak N. Patel. 2018. Convey: Exploring the Use of a Context View for Chatbots. In Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems - CHI '18, 1–6. https://doi.org/10.1145/3173574.3174042

[5] Vidhi Jain and Prakhar Agarwal. 2017. Symptomatic Diagnosis and Prognosis of Psychiatric Disorders through Personal Gadgets. In Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems - CHI EA '17, 118–123. https://doi.org/10.1145/3027063.3048417

[6] J. F. Kelley. 1984. An iterative design methodology for user-friendly natural language office information applications. ACM Transactions on Information Systems 2, 1: 26–41. https://doi.org/10.1145/357417.357420

[7] Jonathan Lazar, Jinjuan Heidi Feng, and Harry Hochheiser. 2017. Research methods in human-computer interaction. Morgan Kaufmann. pp.47

[8] Ewa Luger and Abigail Sellen. 2016. "Like Having a Really Bad PA": The Gulf between User Expectation and Experience of Conversational Agents. In Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems - CHI '16, 5286–5297. https://doi.org/10.1145/2858036.2858288

[9] Johannes Moskaliuk, Carolin P. Burmeister, Florian Landkammer, Bettina Renner, and Ulrike Cress. 2017. Environmental effects on cognition and decision making of knowledge workers. Journal of Environmental Psychology 49: 43–54. https://doi.org/10.1016/j.jenvp.2016.12.001

[10] Clifford Nass, Jonathan Steuer, and Ellen R. Tauber. 1994. Computers are social actors. In Proceedings of the SIGCHI conference on Human factors in computing systems celebrating interdependence - CHI '94, 72–78. https://doi.org/10.1145/191666.191703

[11] Vinícius Segura, Juliana Jansen Ferreira, Ana Fucs, Marcio Ferreira Moreno, Rogério de Paula, and Renato Cerqueira. 2018. CoRgI: Cognitive Reasoning Interface. In International Conference on Human-Computer Interaction, 398–409.

[12] Gheorghe Tecuci, Dorin Marcu, Mihai Boicu, and David A Schum. 2016. Knowledge Engineering: Building Cognitive Assistants for Evidence-based Reasoning. Cambridge University Press.

[13] Alexandra Vtyurina and Adam Fourney. 2018. Exploring the Role of Conversational Cues in Guided Task Support with Virtual Assistants. In Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems - CHI '18, 1–7. https://doi.org/10.1145/3173574.3173782

[14] Anbang Xu, Zhe Liu, Yufan Guo, Vibha Sinha, and Rama Akkiraju. 2017. A New Chatbot for Customer Service on Social Media. In Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems - CHI '17, 3506–3510. https://doi.org/10.1145/3025453.3025496