

Developing a Conversational Agent With a Globally Distributed Team: An Experience Report

Elayne Ruane
University College Dublin, Ireland
Lero - Irish Software Research Centre

Ross Smith
Microsoft Corporation, Skype
Division, Seattle, USA

Dan Bean
Microsoft Corporation (Ret.), Skype
Division, Seattle, USA

Michael Tjalve
Microsoft Corporation, University of
Washington, Seattle, USA

Anthony Ventresque
University College Dublin, Ireland
Lero - Irish Software Research Centre

ABSTRACT

In this experience report, we discuss the development of a solution that enables conflict-affected youth to discover and access relevant learning content. A team of individuals from a not-for-profit, a large multi-national technology company, and an academic institution, collaborated to develop that solution as a conversational agent named Hakeem. We provide a brief motivation and product description before outlining our design and development process including forming a distributed virtual team, engaging in user-centred design with conflict-affected youth in Lebanon, and using a minimum viable product approach while adapting Scrum for distributed development. We end this report with a reflection on the lessons learned thus far.

CCS CONCEPTS

• **Human-centered computing** → **User centered design; Participatory design; Interaction design theory, concepts and paradigms.**

KEYWORDS

chatbot, global software engineering, software development

ACM Reference Format:

Elayne Ruane, Ross Smith, Dan Bean, Michael Tjalve, and Anthony Ventresque. 2020. Developing a Conversational Agent With a Globally Distributed Team: An Experience Report. In *15th IEEE/ACM International Conference on Global Software Engineering (ICGSE '20)*, October 5–6, 2020, Seoul, Republic of Korea. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3372787.3390430>

1 MOTIVATION

Conflict has forced more than 70 million people around the world to flee their homes. More than half of the global refugee population (52%) are under the age of 18. A key challenge for conflict-affected youth is access to education. Figure 1 shows the gap between education enrolment rates of refugee youth compared to global youth.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ICGSE '20, October 5–6, 2020, Seoul, Republic of Korea

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-7093-6/20/05...\$15.00
<https://doi.org/10.1145/3372787.3390430>

While there is a significant gap at the primary level, it becomes a chasm at the tertiary level. This is creating the real risk of a “Lost Generation”. With limited access to education and skills training, many refugees never go back to formal schooling [6]. Education enables children and youth to thrive, not just survive. Failing to provide education for millions of refugees of school-going age can be hugely damaging, not only for individuals but also for their families and societies, perpetuating cycles of conflict and yet more forced displacement. Education is central to both peace and development – and to helping refugee children fulfil their potential. Increasing global connectivity has provided a foundation for online learning and skills training opportunities that can mitigate the education risk. Unfortunately, online learning opportunities are often scattered across the Web, difficult to discover, and frequently gated behind paywalls. A key issue are language barriers where learning opportunities are not available in a refugee’s native language(s).

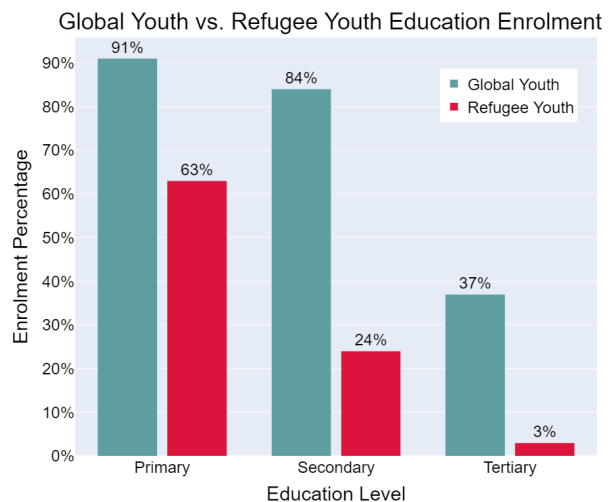


Figure 1: Comparison of refugee youth enrolment rates compared to global youth enrolment rates [6].

2 PRODUCT DESCRIPTION

We have developed a chatbot to discover and access relevant learning content for conflict-affected youth who are seeking education,

skills-training, or other job-readiness resources. Our Learning Companion bot, named Hakeem, provides anytime and anywhere access to curated curricula of learning opportunities. Hakeem is a chatbot designed for school-age conflict-affected youth who may have been out of formal education due to displacement caused by global conflict. This free product aims to leverage existing but scattered online educational resources such as MOOCs (Massive Open Online Course) to provide tailored educational resources for these young people so they can begin to rebuild their lives and achieve their dreams in the face of difficult circumstances. While there are many learning opportunities and resources available today from companies, academic institutions, and other sources, it can require considerable time and effort to find, access, and align them with individual needs. By conversing with Hakeem, the user will be able to find a course that matches their interests and experience.

Interaction with Hakeem is through a chat interface (see Figure 2) where users can converse with the chatbot using text utterances or clicking available option and navigation buttons. Currently, Hakeem is hosted on Microsoft Skype and can be accessed using any device that supports Skype, once it has been added as a Skype contact. Hakeem provides support for English or Arabic interaction. Initially, the chatbot responds to the user using both English and Arabic and then asks for a language preference.

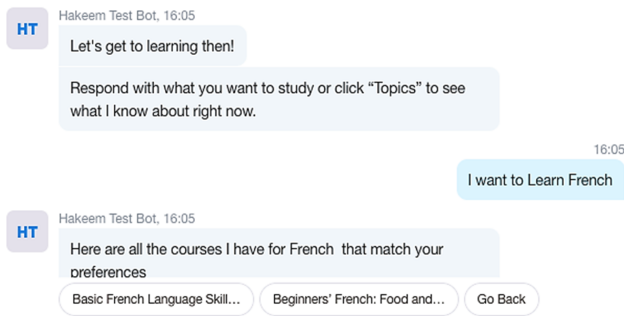


Figure 2: Example Dialogue within the Chat Interface

The course library is curated by humanitarian and subject-matter experts and is aligned with young people’s needs in the context in which they live. The library is updated regularly and the system prioritizes courses that are available for free, in Arabic, with certification, and that can be downloaded for offline access. Content categories currently include Language, Technology, Business, and Design. Humanitarian agencies provide programming that makes the Learning Companion bot discoverable and accessible to conflict-affected youth such as awareness campaigns, computer lab space, and ongoing support including transportation and mentoring.

3 TEAM FORMATION

Global software development is motivated by access to expertise, time pressure in competitive markets, and cost reduction [3, 4]. As this project tackles a complex global and social problem with a technical solution, it was crucial to assemble a diverse and knowledgeable global team. Time pressure was not borne from competition but from humanitarian impetus, i.e. the potential positive

impact on the lives of conflict-affected youth. As a non-commercial project with minimal funding, cost reduction was a key consideration. Consequently, Hakeem is the product of an international, cross-discipline collaboration between a virtual team of individuals from (i) NetHope, a US-based NGO, (ii) Microsoft in Seattle, (iii) the Norwegian Refugee Council (NRC), (iv) the Complex Software Lab at University College Dublin, and (v) Lero - The Irish Software Research Centre.

In September 2017, the No Lost Generation (NLG) Tech Task Force, led by NetHope, brought together 50+ representatives from private and humanitarian sectors to collaborate on challenges faced by conflict-affected youth in MENA (Middle East, North Africa). NetHope used these ideas to formulate four projects, each in collaboration with a private sector company. NRC took the lead on this project with Microsoft and Microsoft brought UCD on board. NetHope acted as facilitator and convener of the overall project, supporting each stakeholder as needed. As the subject matter experts on the problem space and context, NRC led the facilitation of youth engagement in development, solution testing, and content curation. UCD has been responsible for the technical work and development of Hakeem with support from Microsoft through technical guidance, funding, and product donation (e.g. Azure credits).

There were a total of 11 individuals on the initial team; 2 from NRC, 1 from NetHope, 3 from Microsoft, and 5 from UCD/Lero¹. Some team members had pre-existing social ties with ongoing collaborative relationships but prior to formation had never worked together on a software development project. There were transient members who joined the team for short stints (3-6 months), in both technical and non-technical roles. As such, team cohesion was an important focus at the beginning of the project to create a fully performing team capable of repeatedly onboarding new members. The team grew to 16 members at its biggest and is currently made up of 9 individuals (see Figure 3).

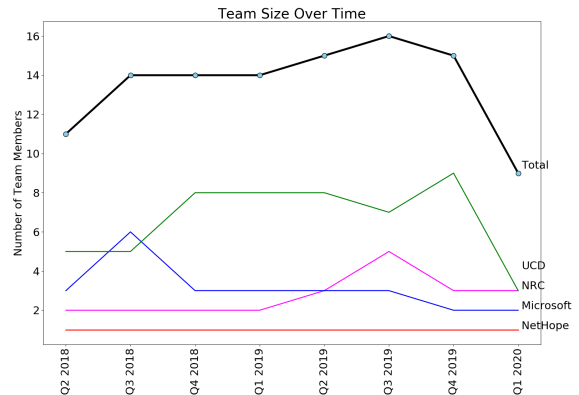


Figure 3: Team size over the product lifecycle.

Major strengths of this team are diversity of expertise, intrinsic motivation, and a strong shared belief in the project as evidenced by the voluntary commitment of significant time and effort. However, well-known disadvantages of distributed virtual teams are

¹The Complex Software Lab in UCD is a member of Lero.

also present, including dispersed resources and working across multiple time zones. Unlike more traditional software development projects that may leverage distributed teams within a single company, this team represents several organizations, each with their own organizational structure, job titles and responsibilities, processes, priorities, and resources. The working language of the project is English but members of the team represent a variety of cultures and languages.

4 USER-CENTERED DESIGN APPROACH

While technology like Artificial Intelligence (AI) can be a powerful tool, it should never be the starting point for doing social good. It is critical to first understand the target user group, the problems they face, and the context in which any solution will be applied. Context includes factors like local infrastructure (e.g. connectivity, power), access to technology (e.g. mobile phones, tablets), and cultural barriers (e.g. gender discrimination, languages barriers). The team began by creating a lean canvas² for the project to establish a shared understanding of the problem, target audience, and a possible solution. A conversational virtual agent was identified as a feasible solution. To ensure it would actually provide the intended value, the team practiced user-centered design. User-centered design focuses on the experience of the end-user and incorporates user feedback throughout an iterative design process. The team included conflict-affected youth as co-creators, guiding design decisions such as chatbot persona, feature prioritization, and the course library.

4.1 Target User Group

A fundamental aspect of user-centered design is forming a good understanding of the user. The target user group were individuals displaced by the war in Syria, with whom the NRC work in both Lebanon and Jordan. This conflict-affected youth population is 50% male and 50% female with Arabic as their primary language and English as their secondary language. Most have completed primary education but have not been able to pursue secondary or tertiary education. For many, there are few academic courses that fit into their real world schedule which can include full-time employment. Also, if the conflict-affected youth has children, there can be a shortage of time to take any courses and the finances to afford them. They often have access to smartphones, but not tablets or desktop computers, necessitating access to shared computing resources. Even with the availability of computing devices, connectivity to the Web is often a challenge creating the need to study offline.

4.2 Co-design

The team experts in humanitarian response (NetHope, NRC) worked closely with a group of 15 conflict-affected youth in Lebanon to determine what would provide them with the most effective and engaging lifelong learning experience. The desire to establish a co-design effort with conflict-affected youth led to fully embracing user-centered design. A key question was whether building a conversational user experience would provide additional value over a website. After several months of working directly with conflict-affected youth in Lebanon, the unequivocal answer was yes. The

²A one page business planning tool.

youth were seeking guidance and felt an interactive, conversational experience with a virtual Learning Companion could help them the most. Similarly, the name Hakeem was chosen because they liked the idea of the agent representing a wiser older brother. A direct result of co-designing with the target user group was the decision to implement the virtual Learning Companion as an AI-supported chatbot that provides real learning value and offers an engaging interaction model that encourages lifelong learning.

5 GATHERING REQUIREMENTS

Conversational interface and dialogue design involves anticipating user needs and often requires considerable effort and communication with users to arrive at stable requirements. Requirements were elicited through user stories, an established agile methodology where textual stories are used to represent requirements in a way all stakeholders can comprehend and which developers can use to guide development of specific features [2, 5]. Both technical and non-technical members of the team participated in drafting user stories which were generated from information gathered through surveys of conflict-affected youth by the NRC. Each story represents a user perspective for a specific use-case of the chatbot. Below is a condensed excerpt of “Mariam’s Learning Companion Story” which accompanied example dialogues with the chatbot.

Mariam had a knack for baking. To help her family, she often made desserts to sell in her Mom’s store. Mariam dreamed of baking desserts for a fine restaurant or even starting a bakery someday. Like many teens, she prized her smartphone. As a refugee, she felt fortunate to have one. Mariam started to read the brochure on the Learning Companion Bot left by the humanitarian worker. Written in Arabic, it was easy for her to read. The first thing she noticed was the chatbot also used Arabic. Mariam learned she could send messages to it and receive information back about learning opportunities. She wondered if it could help her learn how to market her baking skills. Reading the instructions, Mariam realized she already had Skype on her smartphone and would just need to add the Learning Companion Bot as a new Skype contact. That seemed easy.

6 DEVELOPMENT METHODOLOGY

Hakeem has been developed using an agile methodology which includes close relationships with end users, an iterative development process, and the early involvement of multiple stakeholders. The team was committed to designing Hakeem in a responsible manner. In addition to inclusion of a diverse set of stakeholders, some other characteristics of responsible innovation are:

- *Transparency* – ensuring the entire team understands what is being built and why (e.g. what are the benefits or risks of features like User Preferences) and sharing with the youth how each feature works.
- *Anticipating Consequences* – anticipating impact of design decisions is honed by asking questions (e.g. about data collection) and examining the value trade-offs.
- *Responsive design* – using an iterative process that allowed for debate, especially around conflicting values, and for changes based on feedback from youth and humanitarian staff.

6.1 Minimum Viable Product

Minimum viable product (MVP) is a development approach where a product is developed with just enough features to get feedback from users. The feedback then provides the basis for enhancing the product and gathering more feedback to form an iterative development cycle. The creation of a Learning Companion MVP and gathering feedback from conflict-affected youth was a first priority. When envisioning the first MVP, various design elements were considered including:

- Search for available learning resources using text or speech
- Select learning opportunities based on user needs using a set of predefined questions
- Offer curated set of learning resources from a larger portfolio
- Prompt or remind user to learn when it notices inactivity
- Notify user when new learning resources are available
- Allow user to submit requests for additional learning resources to be reviewed by curators and added in the future
- Suggest learning resources based on past engagement

After much discussion, and based on requirements gathered from user stories, the team decided to include the following design elements in the scope for the MVP:

- Discovery of learning content and courses
- Access to learning content and courses outside of the bot
- User keyword search for course content
- Interactive dialogue to narrow search results
- User preferences e.g. preferred language

The feedback from users on the MVP was used by the team to inform functional requirements and focus the vision of the project.

7 SCRUM FOR DISTRIBUTED DEVELOPMENT

Due to initial requirements volatility, Scrum was a clear choice as the development framework. Formal roles were not initially established, however, the team naturally self-organized according to their skill set to take on roles such as Scrum Master (from Microsoft) and Product Owner (from NetHope).

7.1 Communication

The team adapted Scrum to deal with limited co-location and the use of online communication. As daily stand-ups were not possible, frequent virtual team-wide updates were made a priority. Microsoft Teams was used for ad-hoc discussion between members of the development team. Microsoft Skype was used for team-wide, pre-scheduled meetings which were organised by the Project Manager from NetHope to ensure participation from all time-zones. The development team operated in two or three week sprints that were planned and reviewed by the entire team via Skype meetings. Team members used the twice-monthly calls to contribute feature ideas and discuss requirements to help guide the development process.

In lieu of daily stand-ups, project visibility around task allocation and project milestones was achieved using virtual tools. Trello, a list-making application, was used to keep track of task allocation, was visible to the entire team, and was actively monitored and updated by the Project Manager. Google Drive was used to maintain detailed documentation of the design and development processes including

an in-depth technical specifications document, a collaborative effort between all team-members, led by the Project Manager.

7.2 Distributed Development

Early in the project, a Hackathon was held in Microsoft where members of the organization came together to brainstorm feature ideas and develop basic prototypes. The results of the Hackathon were documented in detail, discussed by the entire project team, and passed over to the development team. The core development team was based in UCD but remote development was carried out by Complex Software Lab members in Canada. Intermittent code reviews were conducted in Seattle by members from Microsoft.

Distributed development is difficult due to barriers to coordination [1]. It is well known that informal conversations are strong contributors to a successful software development project. These interactions were not available due to geographic distance. While the formal communication tools employed ensured the project progressed appropriately, the development team required more intensive communication. To mitigate coordination challenges, a local coordinator led the core development team in UCD and coordinated remote development efforts. Microsoft Teams was used for frequent reporting from remote development team members and GitHub was used for version control. Cross-site relationships were developed and maintained through video conference calls without in-person site visits. Pair programming and other collaborative development techniques that require high collaboration maturity [4] were not used in this project due to the distributed nature and relatively short-term work relationships between developers.

One of the most time-consuming parts of designing a conversational agent is specifying and implementing the dialog flow. The entire team developed some initial dialog scripts based on user stories. Then, the development team implemented an interaction mode ("Follow Me Mode") in the chatbot to crowdsource the development of the dialog flow. This mode allows any team member to teach the bot how to walk through certain branches of the dialog tree by typing in the interaction between user and bot. The contributor specifies whether an utterance comes from the user or from the bot and the bot then learns the path and implements it in the dialog flow so that other users in the future can follow that path. This facilitated workload distribution, increased the speed of the development process in the face of geographic distance, and provided ideas from a broader set of people than the core development team.

8 CHALLENGES AND RISKS

There were many challenges that had to be addressed and monitored by the team to mitigate any risk to the progress of the project. We can categorize coordination risk factors for multi-site projects into three categories: technical, organizational, and people [1]. The technical challenges of this project were considerable. State of the art Conversational AI is far from human language understanding capabilities and in this case the agent had to converse in two linguistically distant languages. Additionally, gathering educational resources into a course library took considerable effort.

Organizational risk factors include time zone and geographic separation, divergent incentive structures, divergent management practices, and different schedule pressures. The geographic distance

between team members had the potential to slow progress due to delays in communication across multiple time zones. The distributed nature of the efforts meant virtual communication was required to update other team members on the status of specific activities occurring in each location. Synergies and collaboration often benefit from face-to-face meetings but in this case, virtual meetings were used as a substitute. As a multidisciplinary and multi-stakeholder team, it was important to learn how to work together effectively as a group of individuals from multiple organizations with various work-styles, goals, and resources.

People risk factors included language skills, differences in national and corporate culture, and collaboration immaturity. The working language of the project was English but not all team members shared English as their native language. Additionally, the team did not possess Arabic language skills and so could not validate the accuracy of the Arabic dialogue. The team did not share a common organisational culture and had not worked together before. Having an experienced Project Manager from NetHope and sincere buy-in from team members mitigated these risks.

User risk factors focused on the vulnerability of conflict-affected youth. A constant concern of the project team was unknowingly creating a solution that posed a risk for any further harm. Because there was a significant cultural distance between the development team and the end users, this meant localization efforts required guidance and mediation by remote team members working directly with conflict-affected youth and other country/cultural experts.

The innovative nature of the project meant expectations of the chatbot's capabilities evolved rapidly. The initial short-term goal was to help conflict-affected youth discover and access educational opportunities and resources. However, with more awareness of Hakeem's potential, it is already envisioned that it can offer educational assessment and certification processes that lead to employment. While these capabilities may be possible, they represent the challenge of managing rapidly rising expectations and scope creep. This is reflected in the emergent nature of the system requirements. This requirements volatility could have slowed down or even brought the project to a halt if not managed correctly.

9 LESSONS LEARNED

It took several MVPs and two years to reach Product-Market Fit³ due to the time required for user-centered and responsible design. Iteration grounded in evidence enabled the team to develop a robust solution and improve it by learning from end-user feedback. The iterative process also created a space for knowledge transfer between end-users, humanitarian workers, and technologists.

The project has been successful because the team had a strong vision, was highly organised, engaged in open communication, and had strong buy-in from each member who felt passionate about helping conflict-affected youth. The team has learned that practicing true user-centered design is particularly important where the target user group may be vulnerable and have limited resources. For a technology-enabled solution like Hakeem to be relevant to the end-user and sustainable, it needs to be co-designed by a diverse set of stakeholders, including the end-user and facilitators of solution

³A strong value proposition delivered to the people who want it - developing the right product for the right market.

deployment, and integrated into existing delivery options such as the NRC education programs.

AI technology is evolving at a rapid pace. What was not possible when this project started in 2018 (e.g. developing chatbot without specialized technical knowledge) is now possible thanks to new no-code tools such as Microsoft's Power Virtual Agents which lower the bar for adoption. This has implications for handover and maintenance of the system by non-technical staff at partner NGOs. When developing a solution that addresses both urgent and ongoing needs (immediate skill-acquisition vs. lifelong learning), it is important to make it available as soon as possible while making sure that it is built to last e.g. it can be maintained and updated as needed by humanitarian staff. The team used the phrase "Act with urgency, plan for sustainability" to describe this philosophy.

10 CURRENT PRODUCT STATUS AND FUTURE WORK

The current version of Hakeem is stable and development is ongoing. A public release is expected in 2020 and plans after that are likely to follow three avenues. First, deploy the chatbot to other channels (e.g. Facebook Messenger) in order to meet youth where they are. This is considered by project stakeholders and users as important – while also feasible from a development perspective. Second, increase chatbot capabilities by providing automatic assessment and support (e.g., recommending extra content if a particular concept is hard to understand). This will involve expanding the research and engineering team and adding more content partners. Third, make the chatbot accessible to other charities and user groups e.g., refugees in other parts of the world or for retraining displaced workers. A critical requirement has been for a team like NRC to self-manage the chatbot and evolve it based on changing needs. To support this, Hakeem is customizable and the course library can be managed by education experts using a web-interface. It is our hope that Hakeem has a positive and lasting impact on the lives of conflict-affected youth and that others will be inspired to tackle educational barriers and help avoid any future Lost Generations.

11 ACKNOWLEDGMENTS

This work was supported, in part, by Science Foundation Ireland grant 13/RC/2094. This project would not have been possible without dedicated leadership from Leila Toplic of NetHope and crucial work on the ground by Barbara and Stephen of NRC.

REFERENCES

- [1] Matthew Bass, James D Herbsleb, and Christian Lescher. 2009. A coordination risk analysis method for multi-site projects: Experience report. In *2009 Fourth IEEE International Conference on Global Software Engineering*. IEEE, 31–40.
- [2] Mike Cohn. 2004. *User stories applied: For agile software development*. Addison-Wesley Professional.
- [3] James D Herbsleb. 2007. Global software engineering: The future of socio-technical coordination. In *Future of Software Engineering (FOSE'07)*. IEEE, 188–198.
- [4] Stefan Lasser and Michael Heiss. 2005. Collaboration maturity and the offshoring cost barrier: the tradeoff between flexibility in team composition and cross-site communication effort in geographically distributed development projects. In *IPCC 2005. Proceedings*. IEEE, 718–728.
- [5] Garm Lucassen, Fabiano Dalpiaz, Jan Martijn EM van der Werf, and Sjaak Brinkkemper. 2016. The use and effectiveness of user stories in practice. In *International working conference on requirements engineering: Foundation for software quality*. Springer, 205–222.
- [6] UNHCR. 2019. Stepping Up: Refugee Education in Crisis.