

On Chatbots Exhibiting Goal-Directed Autonomy in Dynamic Environments

Biplav Srivastava
IBM Research

Abstract

Conversation interfaces (CIs), or chatbots, are a popular form of intelligent agents that engage humans in task-oriented or informal conversation. In this position paper and demonstration, we argue that chatbots working in dynamic environments, like with sensor data, can not only serve as a promising platform to research issues at the intersection of learning, reasoning, representation and execution for goal-directed autonomy; but also handle non-trivial business applications. We explore the underlying issues in the context of *Water Advisor*, a preliminary multi-modal conversation system that can access and explain water quality data.

Introduction

Chatbots (McTear, Callejas, and Griol 2016), which can engage people in natural dialog conversation, have gained popularity recently drawn by numerous platforms to create them quickly for any domain (Accenture 2016). Most common types of such agents deal with a single user at a time and conduct informal conversation, answer the user's questions or provide recommendations in a given domain. They need to handle uncertainties related to human behavior and natural language, while conducting dialogs to achieve system goals. Chatbots have been deployed in customer care in many industries where they are expected to save over \$8 billion per annum by 2022 (Juniper 2017).

However, the data sources used by common chatbots are static databases like product catalogs or user manuals. Therefore, for their problem of dialog management, i.e., creating dialog responses to user's utterances, effective approaches include learning policies over predictable nature of data (Young et al. 2013) or reasoning on its abstract representations (Inouye 2004).

The application scenarios become more compelling when the chatbot works in a dynamic environment, e.g., with sensor data, and interacts with groups of people, who come and go, rather than only an individual at a time. In such situations, the agent has to execute actions to monitor the environment, model different users engaged in conversation over time and track their intents, learn patterns and represent them, reason about best course of action given goals and

system state, and execute conversation or other multi-modal actions.

We now explore the underlying issues of goal-directed autonomy in dynamic environment in the context of *Water Advisor* (WA) (Ellis et al. 2018), a prototypical multi-modal conversation system that can access and explain water quality data to a variety of stake-holders. We identify opportunities for learning, reasoning, representation and execution in WA and motivate more such applications.

Decision-Support for Water Usage With a Multi-Modal Conversation Interface

The global situation of water quality around the world is alarming in both developing and developed countries (UNEP 2016) because water demand continues to rise while existing sources for fresh water are getting polluted. A key strategy for tackling water pollution is engaging people. A person makes many daily decisions touching on water usage activities like for profession (e.g., fishing, irrigation, shipping), recreation (e.g. boating), wild life conservation (e.g., dolphins) or just regular living (e.g., drinking, bathing, washing). Accessible tools for public are particularly useful to handle public health challenges such as the Flint water crisis (Pieper, Tang, and Edwards 2017).

A decision in this space needs to consider the activity (purpose) of the water use; relevant water quality parameters and their applicable regulatory standards for safety; available measurement technology, process, skills and costs; and actual data. There are further complication factors: there may be overlapping regulations due to geography and administrative scope; one may have to account for alternative ways to measure a particular water quality parameter that evolves over time; and water data can have issues like missing values or at different levels of granularity. The very few tools available today target water experts such as WaterLive mobile app for Australia¹, Bath app for UK², and GangaWatch for India (Sandha, Srivastava, and Randhawa 2017) and assume technical understanding of sciences.

¹<http://www.water.nsw.gov.au/realtime-data>

²<https://environment.data.gov.uk/bwq/profiles/>




Figure 1: A screenshot of *Water Advisor*. See video of it in action at <https://youtu.be/z4x44sxC3zA>.

Water Advisor

Water Advisor (WA) is intended to be a data-driven assistant that can guide people without requiring any special water expertise. One can trigger it via a conversation to get an overview of water condition at a location, explore it by filtering and zooming on a map, and seek details on demand (Figure 1) by exploring relevant regulations, data or other locations. The current prototype uses water quality data available from Flint, MI³ but future extensions will use open water data from US Geological Survey⁴ (USGS) that is refreshed for thousands of places in US per day. However, the number of water quality parameters, for which data is available, varies widely between locations and over time, making generation of useful advice challenging. For regulations, WA relies on information provided by multiple agencies at national (US, India) and state levels (Michigan, New York), which has been consolidated for reuse⁵.

Technical Issues

In a water advising application, one or more users may need to interact with the chatbot if handling a complex decision like water contamination. The tool has to detect the user's information goals and meet them at lowest cognitive cost. The system uses a natural language classifier (NLC) to understand user utterance, and its error rate varies with input. The system has to decide whether to ask clarifying questions if it has low confidence and there are many ways to respond. The user may have preferences about how they specify an input (like location) and the kind of response they want (visual v/s textual). We discuss a range of issues below for exposition

³http://www.michigan.gov/flintwater/0,6092,7-345-76292_76294_76297—,00.html

⁴<https://waterdata.usgs.gov/nwis/current/?type=quality>

⁵<https://github.com/biplav-s/water-info>

but note that the current WA prototype handles only a subset of following integration issues.

Learning plays an important role in understanding user's utterance, finding reliable water data samples in the database based on region and duration of interest, discovering issues in water quality and improving overall performance over time. In the prototype, for utterances, we use trained user models from commercial systems and for water quality, a simple regression method.

Representation is needed to map water's usage purpose to quality parameters and model safe limits of pollution parameters with different mathematic properties (e.g., polarity). It also helps map water purpose to regulations and further, aggregate and reconcile the latter when a region falls under overlapping jurisdiction of regulations. We represent this as geographically-scoped attribute-value pair in JSON format and make it publicly available for others to use and extend⁶.

Reasoning is crucial to keep conversation focused based on system usability goals and user needs. One can model cognitive costs to user based on alternative system response choices and seek to optimize short-term and long-term behavior. Reasoning can further help to short-list regulations based on water activity and region of interest, generate advice and track explanations. We currently use rules on geographical scope and missing values to determine system response.

Execution is autonomous as the agent can choose to act by (a) asking clarifying questions about water usage goals or locations, (b) asking user's preference about advice, (c) seeking most reliable water data for region and time interval of interest from available external data sources, and corre-

⁶<https://github.com/biplav-s/water-info>

sponding subset of compatible regulations (d) invoking reasoning to generate an advice for water usage using filtered water data and regulations, (e) visualizing and explaining its output using water regulations, and (f) using one or more suitable modalities available at any turn of user interaction, i.e., chat, maps and document views. The current prototype uses a simplistic strategy for execution based on error rates, system confidence and usability rules.

Human Usability Factors have to be modeled and supported during WA's operation. In the current prototype, the user-interface controller module automatically keeps the different modalities synchronized so that the user is looking at consistent information across them. The system has to be aware of missing data or assumptions it is making, and needs to take them into account while communicating output advice in generated natural language. One avenue for future exploration is to measure and track complexity of interaction (Liao, Srivastava, and Kapanipathi 2017) and use sensed signals to pro-actively improve user experience. Another is to combine close-ended and open-ended questioning strategies for efficient interaction (Zhang, Liao, and Srivastava 2018).

Ethical Issues can emerge whenever a piece of technology is used among people at large. In the context of conversations, a recent paper surveys ethical issues (Henderson et al. 2018) like biases, adversarial examples, privacy violations, safety challenges and reproducibility concerns. A water-use chatbot can conceivably create bias among users of different activity subgroups (e.g., preferring recreation over drinking), compromise on privacy of users who submit queries about an activity or a region, and create public safety concerns (e.g., when users find scarcity of good quality water). We have not considered them in the prototype, however.

Discussion and Conclusion

In this paper, we used decision-support in water as a use-case to demonstrate that chatbots can serve as a promising platform to integrate AI sub-disciplines for goal-directed autonomy. Apart from learning, reasoning, representation and execution, chatbots also need to work with human usability factors and ethical issues. An interesting aspect of these applications is that the chatbot may be helping a group of people take collective decision making, like conducting an interview, and data changes over time. Beyond water and customer support, complex applications are emerging in sciences (astronomy(Kephart et al. 2018)), business (career counseling⁷, hospitality⁸) and societal domains (health⁹).

References

- Accenture. 2016. Chatbots in customer service. In *At: https://accntu.re/2z9s5fH*.
- Ellis, J.; Srivastava, B.; Bellamy, R.; and Aaron, A. 2018. Water advisor - a data-driven, multi-modal, contextual assistant to help with water usage decisions. In *Proc. 32nd AAAI Conference on Artificial Intelligence (AAAI-18), New Orleans, Louisiana, USA*.
- Henderson, P.; Sinha, K.; Angelard-Gontier, N.; Ke, N. R.; Fried, G.; Lowe, R.; and Pineau, J. 2018. Ethical challenges in data-driven dialogue systems. In *Proc. of AAAI/ACM Conference on AI Ethics and Society (AIES-18), New Orleans, Louisiana, USA*.
- Inouye, R. B. 2004. Minimizing the length of non-mixed initiative dialogs. In Leonor van der Beek, Dmitriy Genzel, D. M., ed., *ACL 2004: Student Research Workshop*, 7–12. Barcelona, Spain: Association for Computational Linguistics.
- Juniper. 2017. Chatbots: Retail, ecommerce, banking & healthcare 2017-2022. In *At: http://bit.ly/2sJHejY*.
- Kephart, J.; Dibia, V.; Ellis, J.; Srivastava, B.; Talamadupula, K.; and Dholakia, M. 2018. Cognitive assistant for visualizing and analyzing exoplanets. In *Proc. 32nd AAAI Conference on Artificial Intelligence (AAAI-18), New Orleans, Louisiana, USA*.
- Liao, Q.; Srivastava, B.; and Kapanipathi, P. 2017. A Measure for Dialog Complexity and its Application in Streamlining Service Operations. *ArXiv e-prints*.
- McTear, M.; Callejas, Z.; and Griol, D. 2016. Conversational interfaces: Past and present. In *The Conversational Interface*. Springer, DOI: https://doi.org/10.1007/978-3-319-32967-3_4.
- Pieper, K. J.; Tang, M.; and Edwards, M. A. 2017. Flint water crisis caused by interrupted corrosion control: Investigating "ground zero" home. *Environmental Science & Technology* 51(4):2007–2014.
- Sandha, S. S.; Srivastava, B.; and Randhawa, S. 2017. The gangawatch mobile app to enable usage of water data in every day decisions integrating historical and real-time sensing data. *CoRR* abs/1701.08212.
- (UNEP), U. N. E. P. 2016. A snapshot of the worlds water quality: Towards a global assessment. In *Nairobi, Kenya*. Online at: https://uneplive.unep.org/media/docs/assessments/unep_wwqa_report_web.pdf.
- Young, S.; Gašić, M.; Thomson, B.; and Williams, J. D. 2013. Pomdp-based statistical spoken dialog systems: A review. *Proceedings of the IEEE* 101(5):1160–1179.
- Zhang, Y.; Liao, V.; and Srivastava, B. 2018. Towards an optimal dialog strategy for information retrieval using both open-ended and close-ended questions. In *Proc. Intelligent User Interfaces (IUI 2018, Tokyo, Japan, March)*.

⁷<https://www.ibm.com/talent-management/career-coach>

⁸<https://www.bebot.io/hotels>

⁹<https://www.healthtap.com/>