Uncertainty-Resolving Questions for Social Robots

Minjung Shin
Interdisciplinary Program in Cognitive Science,
Seoul National University,
Seoul, Korea
mjshin77@snu.ac.kr

Minsu Jang
Electronics and Telecommunications
Research Institute,
Daejeon, Korea
minsu@etri.re.kr

Miyoung Cho
Electronics and Telecommunications
Research Institute,
Daejeon, Korea
mycho@etri.re.kr

Jeh-Kwang Ryu
Department of Physical Education,
Dongguk University,
Seoul, Korea
ryujk@dgu.ac.kr

ABSTRACT

Social robots should deal with uncertainties in unseen environments and situations in an interactive setting. For humans, question-answering is one of the most typical activities for resolving or reducing uncertainty by acquiring additional information, which is also desirable for social robots. In this study, we propose a framework for leveraging the research on learning-by-asking techniques for social robots. This framework is inspired by human inquiries. Information seeking by asking should be considered at the multi-dimensional level, including required knowledge, cognitive processes, and question types. These dimensions offer a framework to embed generated questions into the three-dimensional question space, which is expected to provide a reasonable benchmark for the active learning approach and evaluation methodologies of uncertainty-resolving question generation for social robots.

CCS CONCEPTS

• **Computing methodologies** → **Discourse, dialogue and pragmatics**; **Cognitive science**; **Cognitive robotics**; • **Applied computing** → *Psychology*.

KEYWORDS

Social robot, Uncertainty resolution, Question generation, Inquiry type, Question space

ACM Reference Format:
Minjung Shin, Minsu Jang, Miyoung Cho, and Jeh-Kwang Ryu. 2023. Uncertainty-Resolving Questions for Social Robots. In *Companion of the 2023 ACM/IEEE International Conference on Human-Robot Interaction (HRI '23 Companion), March 13–16, 2023, Stockholm, Sweden.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3568294.3580077

1 INTRODUCTION

Question Answering (QA) is the most common strategy for humans to resolve or reduce uncertainties from the environment or interaction with other people [12]. So learning by asking is also a reasonable approach for social robots which perform tasks in highly interactive situations with people [17, 20]. Figure 1 depicts a specific situation where a social robot resolves ambiguity by asking. It is a common ambiguity occurred in daily conversation but cannot be solved by a probability distribution or searching on the internet. Just asking back is the best solution.

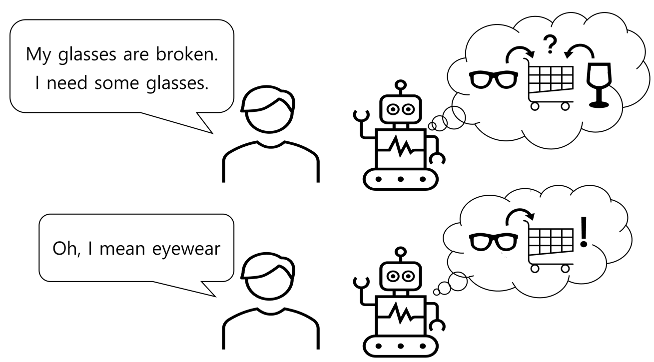

Figure 1: A scenario in that a social robot resolves uncertainty by asking a human.

Asking informative questions seems not easy, even for humans [24]. Over and above, since generating questions has a myriad spectrum of diversity due to a highly context-dependent manner, implementing this feature is challenging for robots and other artificial intelligence (AI) agents. But users indeed have preferred to communicate with robots in 'human-like' languages [10, 22], and asking has been regarded as the most desirable strategy for collaborative control of social robots [6, 11].

Therefore, this study employs a knowledge-based approach to identify acceptable inquiring guidelines from the perspective of how people learn information by asking. Afterward, we propose a framework for classifying questions according to their types and properties to enable social robots to generate helpful inquiries. Additionally, we will discuss the potential for evaluating the validity of the generated questions.

HRI '23 Companion, March 13–16, 2023, Stockholm, Sweden

ACM ISBN 978-1-4503-9970-8/23/03…$15.00
https://doi.org/10.1145/3568294.3580077

2 RELATED WORK

Automatic question-generation algorithms have been developed in the text, image, or video domains to reduce the costs of creating QA datasets. The automatized generation systems show highly quantitative performance, but the dual problems that the generated questions include answers to are pointed out as a limitation [25]. Meanwhile, selecting behaviors by asking humans when a robot or AI agent encountered an ambiguous situation has been studied, but they only generated yes/no questions in a predefined framework [3, 28] without consideration of question diversity and adequacy evaluation.

On the other hand, active learning techniques exhibit more flexible learning performance by allowing making queries as presuming humans as oracles. In robotics, active learner leveraging the Learning from Demonstration (LfD) framework [4] showed efficient achievement [6, 15] in motion learning. However, to date, related studies are focused on robot motion. In addition, a variety interactive models were introduced as human-like question generation in game domain [23], or in image comprehension [20]. Recently, a prosocial active learning model that works in an open online community has been presented [17]. It is noteworthy that suggesting helpful question types elicits a cooperative and voluntary reaction from the crowd; however, it is a case-specific optimization focused on image comprehension.

Recent interactive social robots or agents introduced in our daily life mainly provide information to users, such as guide robots or voice assistant agents. However, active asking capabilities of social agents are crucial in some areas. For clinical purposes, interactive robots are being introduced as an assistive role in autistic spectrum disorders [9] and social anxiety [21] that the individuals have difficulty interacting with other people. Similarly, for public service, especially helping the victims of crime, interactive chatbots attract attention as an agent that provides mental support and encourages the disclosure of victims [1, 19]. In educational settings, social robots showed increasing students' learning effects by playing the role of a curious peer [7, 13].

3 DEFINING QUESTION SPACE AND CATEGORIES FOR ANALYZING QA

A helpful inquiry does not aim to generate the question sentence itself but to acquire target information (i.e., good answers) to resolve uncertainty. Therefore, the inquiry should be generated and evaluated considering the comprehensive context of answering questions. The primary stage in building the principles of asking within various contexts is considering the multi-factors involved in the QA process. In this work, we consider three factors that form a three-dimensional space referring to a framework from the educational study [14]. Graesser et al. argued that inquiry should be regarded in a landscape that includes a variety of knowledge presentations and cognitive processes. And they concluded that questions requiring high-level cognition are educationally desirable.

We, however, extend the idea across uncertainties in general social interactions. We first set up the inquiry-generating flow involving the three factors as seen in Figure 2, and we treat each factor as an independent axis. Following that, we classify the sub-elements of each factor from social interaction scenes.

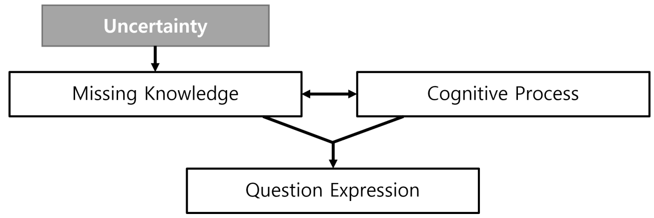

Figure 2: Flow of inquiry-generation in humans under uncertainty

3.1 K-type category: Classification of Knowledge class

The first factor to be considered in QA process is identifying missing information because it stands as a source as well as a target of inquiry. The K-type category is an axis that specifies what kind of knowledge has been missed. Table 1 shows the description of each K-type category that specifies the potential class of missing information in interactive situations. The knowledge gets more complex and subjective from top to bottom.

Table 1: K-type category

#	Categories	Description
K1	Identity	Information about who or what a person or thing is
K2	Class	Inclusion relationships of categories
K3	Attributes	Properties, feature of the object
K4	Quantities	Quantitative specifications
K5	Spatial layout	Spatial relations among entities
K6	Temporal relation	Temporal information or sequences
K7	Contents	Additional detailed information
K8	Procedure	Order or method of specific process
K9	Causality	Causal chains of events or states
K10	Intention	Motivation, aim or plan of other agent
K11	Internal state	Mental states such as preference, mood of other agent

3.2 C-type category: Classification of Cognitive Process

After specifying the knowledge class of missing information, the cognitive process required for searching for the answer should be considered. The cognitive factor, referred to as the C-type category, is the strategy for drawing the missing information from the answerer. For determining the C-type category, We simplified the original system of Bloom's taxonomy (i.e., recognition, recall, comprehension, applying, analyzing, synthesizing, and evaluation) [5] according to the complexity. Knowledge is the easiest level, including recognizing and recalling facts or events without comprehension. Comprehension is the intermediate level of understanding the meaning of the knowledge. Operation requires manipulation, such as applying, analyzing, or synthesizing knowledge based on comprehension. Evaluation, the highest cognitive level, denotes judgment and valuation about comprehension or operation of the knowledge. Table 2 shows description of each C-type category.

Table 2: C-type category

#	Categories	Description
C1	Knowledge	To recognize or recall facts or events
C2	Comprehension	To understand the meaning of something
C3	Operation	To apply, analyze or synthesize something
C4	Evaluation	To make a judgment about something

3.3 Q-type category: Classification of Question Expression

The following issue is how to express the inquiry. The adequate question sentence should come from a combination of accurate targeting of missing information and the proper answering strategy. Q-type categories classify the form of the question sentence according to pragmatic approach. We suggest fourteen question types as Q-type categories, as described in Table 3. From top to bottom, the questions become more profound and subjective, like the K-type and C-type categories.

Table 3: Q-type category

#	Categories	Description
Q1	Verification	Asking whether true or not
Q2	Case specification	Asking to specify the case
Q3	Concept completion	Asking to fulfill insufficient information
Q4	Feature specification	Asking to describe properties
Q5	Quantification	Asking quantitative specification
Q6	Definition	Asking to state the nature, scope, or meaning
Q7	Comparison	Asking dis- or similarity between groups
Q8	Interpretation	Asking to explain the meaning or details
Q9	Cause elucidation	Asking what the antecedent of causality is
Q10	Result account	Asking what the outcome of causality is
Q11	Intention disclosure	Asking about motivation or goal orientation
Q12	Method explication	Asking the procedure, sequence or tools
Q13	Expectation	Asking a belief or case in the future
Q14	Judging	Asking an opinion or evaluation of something

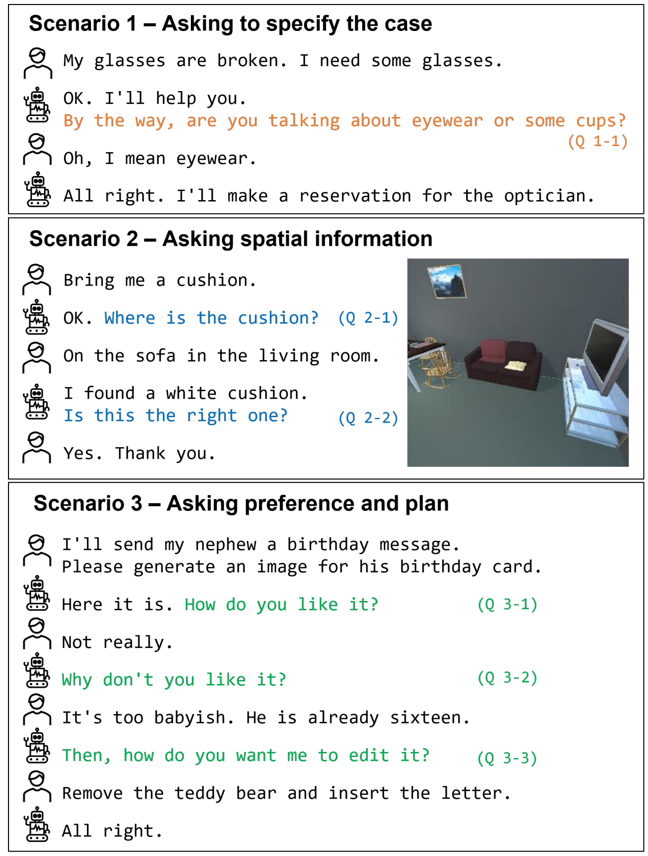

Figure 3: Three scenarios for human-robot interaction situations. (Scenario 1) A case-specifying question to address ambiguity arising from the discourse (Scenario 2) Asking spatial layout and verifying the task of the space navigation robot[8] (Scenario 3), Asking about preference and plan of the user while performing the image generation task.

4 GOOD QUESTION FOR UNCERTAINTY RESOLUTION

Then, which combination of given categories would be a good question? This issue can be addressed by referring to the human question-generation process. When humans encounter an uncertain situation, an information gap arises [18, 26]. Then, emotional drive (i.e., curiosity) to fill the gap occurs naturally [18, 26]. Asking questions to address the uncertainty is the behavioral response to such emotional motivation. Regarding such cognitive processes, a good question should precisely target the information gap and clearly convey the questioner's intention [16].

4.1 Scenarios for Generating questions

We outline potential scenarios of social interactions between people and social robots. Examples of human-like inquiry in the three different circumstances are shown in Figure 3.

Scenario 1 shows ordinary ambiguity in daily conversation. The uncertainty comes from double words, so the agent doesn't know what the glasses are (K1 Identity). Since the user already has a mental image of the glasses he/she said, he/she can reply by simply recalling (C1 Knowledge) it. Then, with a straightforward query (Q2 Case specification) can elicit the user's response. This scenario can be extended to typical situations in searching for information on the internet. According to previous research [2, 27], users make unclear or ambiguous queries very often. Hence, asking clarifying question strategy is critical to narrow the scope toward the user's target information.

Scenario 2 presents a case where a spatial navigating robot interacts with a user. In the first question (i.e., Q 2-1), the robot does not know the target object's location (K5 Spatial layout) and thus ascertain the spatial information by recalling (C1 Knowledge) the user's memory. Robots can address this issue by asking for fulfilling insufficient information (Q3 Concept completion) about where the target object is. The second question (i.e., Q 2-2) is verifying (Q1 Verification) the user's intention (K10 Intention). The verification can be obtained by recalling (C1 Knowledge) his/her predefined mental image.

Scenario 3 is an instance of a generative model that asks about the user's opinion and improves the produced item together. In the first question (i.e., Q 3-1), the uncertainty belongs to the user's preference (K11 Internal state), which needs to appraise the outcome according to his/her desire(C4 Evaluation). So the model requests feedback (Q14 of Judging) for the generation result. The second question (i.e., Q 3-2) arises from not knowing the reason (K9 Causality) for dissatisfaction. It requires the user's analysis (C3 Operation) of his/her opinion. The answer can be achieved by asking to elucidate the cause (Q9 Cause elucidation) of the unsatisfied state. The last question (i.e., Q 3-3) requests a method (K8 Procedure) for improvement by synthesizing (C3 Operation) the user's idea. In this case, the model can ask how to improve (Q12 Method explication) directly.

4.2 Embedding Questions to Question Space

Each axis defined independently forms a three-dimensional space, and the generated question can be allocated in this space. Figure 4 illustrates the embedding of the questions presented in Section 4.1 and Figure 3.

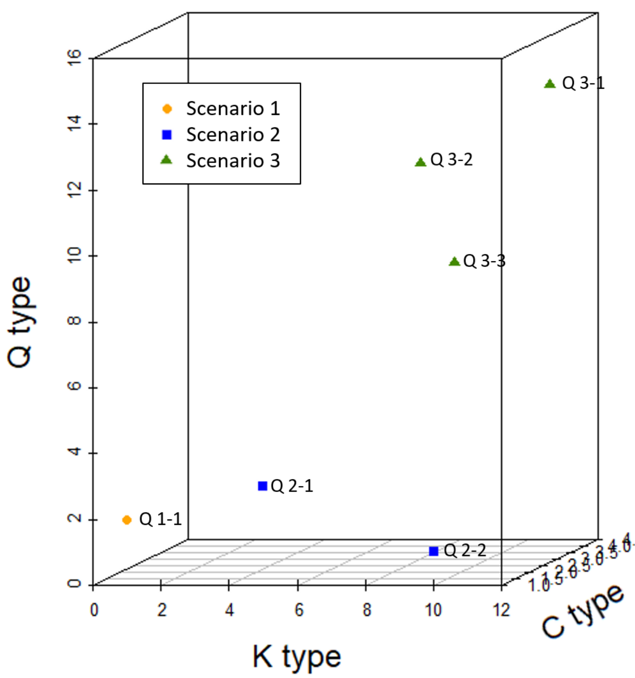

Figure 4: Embedding the generated questions presented Section 4.1 to three dimensional question space

Since sub-elements in each axis are arranged in order from shallow to deep or simple to complex, the embedded location reflects the complexity of the QA process. Specifically, questions that call for simple or straightforward information through an easy process are allocated to the lower left side of the space(i.e., Q 1-1, Q 2-1, Q 2-2). On the other hand, questions asking about complex, ambiguous, and subjective uncertainty are placed to the upper right corner(i.e., Q 3-1, Q 3-2, Q 3-3). As a result, within the question space, questions from similar contexts likely form clusters. In other words, this framework can serve as a basis for determining which position would be more appropriate to generate a question depending on the context of the situation.

5 CONCLUSION AND FUTURE WORK

Since asking questions requires a highly context-dependent manner and takes on myriad forms in interactive situations, adequate question generation for robots and other AI agents was considered intractable. However, for social robot agents to be employed in the real world, active information seeking to mitigate uncertainty is crucial, and asking can be a great strategy. Hence this study argues the significance of a framework to organize the generated questions based on the perspective of human information acquisition. The framework regards the properties and categories of the questions in various factors to reflect on the diverse contexts of interactions.

We, therefore, propose a three-dimensional question space. Each axis represents knowledge class of information gap, the cognitive process needed to answer a question, and the various questioning styles devided by pragmatics. We also classified elements of each dimension into appropriate categories. Assigning generated questions to their categories allows placing them in the corresponding location in the question space.

Since the given position roughly represents how complex the situation is, it can be used as the foundation for determining whether the inquiry matches the context. Social robots should also aim to interact across a wide range of contexts and questions, as typical human-human interactions naturally pass through all the ranges. The framework we suggest could work together with a dataset of context-compatible inquiries driven by embracing various social interaction scenarios. It can be a valuable guideline to improve the social robot agent's ability to come up with insightful questions.

Subsequent research is expected to provide a reasonable benchmark for active learning and evaluation techniques of uncertainty-resolving question generation conducted with actual humans. Furthermore, structuring uncertainty-resolving questions addressed in this study can serve as a cornerstone for creating an extended framework that includes various questions for other social goals, such as intimacy.

ACKNOWLEDGMENTS

We are deeply grateful to the reviewers for providing kind and helpful comments. And this work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No. 2022-0-00951, Development of Uncertainty-Aware Agents Learning by Asking Questions)

REFERENCES

[1] Yuna Ahn, Yilin Zhang, Yujin Park, and Joonhwan Lee. 2020. A Chatbot Solution to Chat App Problems: Envisioning a Chatbot Counseling System for Teenage Victims of Online Sexual Exploitation. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI EA '20)*. Association for Computing Machinery, New York, NY, USA, 1–7. https://doi.org/10.1145/3334480.3383070

[2] Mohammad Aliannejadi, Hamed Zamani, Fabio Crestani, and W. Bruce Croft. 2019. Asking Clarifying Questions in Open-Domain Information-Seeking Conversations. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval* (Paris, France) *(SIGIR'19)*.

Association for Computing Machinery, New York, NY, USA, 475–484. https://doi.org/10.1145/3331184.3331265

[3] Peter Anderson, Qi Wu, Damien Teney, Jake Bruce, Mark Johnson, Niko Sünderhauf, Ian Reid, Stephen Gould, and Anton van den Hengel. 2018. Vision-and-Language Navigation: Interpreting Visually-Grounded Navigation Instructions in Real Environments. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*. IEEE, Salt Lake City, UT, USA, 3674–3683. https://doi.org/10.1109/CVPR.2018.00387

[4] Brenna D. Argall, Sonia Chernova, Manuela Veloso, and Brett Browning. 2009. A survey of robot learning from demonstration. *Robotics and Autonomous Systems* 57, 5 (2009), 469–483. https://doi.org/10.1016/j.robot.2008.10.024

[5] B.S. Bloom and D.R. Krathwohl. 1956. *Taxonomy of Educational Objectives: The Classification of Educational Goals*. Number V. 1 in Taxonomy of Educational Objectives: The Classification of Educational Goals. Longmans, Green, New York, NY.

[6] Maya Cakmak and Andrea L. Thomaz. 2012. Designing robot learners that ask good questions. In *2012 7th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, Boston, MA, USA, 17–24. https://doi.org/10.1145/2157689.2157693

[7] Jessy Ceha, Nalin Chhibber, Joslin Goh, Corina McDonald, Pierre-Yves Oudeyer, Dana Kulić, and Edith Law. 2019. Expression of Curiosity in Social Robots: Design, Perception, and Effects on Behaviour. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) *(CHI ’19)*. Association for Computing Machinery, New York, NY, USA, 1–12. https://doi.org/10.1145/3290605.3300636

[8] Matt Deitke, Eli VanderBilt, Alvaro Herrasti, Luca Weihs, Jordi Salvador, Kiana Ehsani, Winson Han, Eric Kolve, Ali Farhadi, Aniruddha Kembhavi, and Roozbeh Mottaghi. 2022. ProcTHOR: Large-Scale Embodied AI Using Procedural Generation. In *Advances in Neural Information Processing Systems*. 12 pages. Outstanding Paper Award.

[9] Joshua J. Diehl, Lauren M. Schmitt, Michael Villano, and Charles R. Crowell. 2012. The clinical use of robots for individuals with Autism Spectrum Disorders: A critical review. *Research in Autism Spectrum Disorders* 6, 1 (2012), 249–262. https://doi.org/10.1016/j.rasd.2011.05.006

[10] Malcolm Doering, Phoebe Liu, Dylan F. Glas, Takayuki Kanda, Dana Kulić, and Hiroshi Ishiguro. 2019. Curiosity Did Not Kill the Robot: A Curiosity-Based Learning System for a Shopkeeper Robot. *J. Hum.-Robot Interact.* 8, 3, Article 15 (jul 2019), 24 pages. https://doi.org/10.1145/3326462

[11] Terrence Fong, Charles Thorpe, and Charles Baur. 2003. Robot, asker of questions. *Robotics and Autonomous Systems* 42, 3 (2003), 235–243. https://doi.org/10.1016/S0921-8890(02)00378-0

[12] Russell Golman and George Loewenstein. 2018. The Desire for Knowledge and Wisdom. In *The New Science of Curiosity*, Goren Gordon (Ed.). Nova Science Publishers, Inc., Hauppauge, NY, Chapter 2, 37–42.

[13] Goren Gordon, Cynthia Breazeal, and Susan Engel. 2015. Can Children Catch Curiosity from a Social Robot?. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction* (Portland, Oregon, USA) *(HRI ’15)*. Association for Computing Machinery, New York, NY, USA, 91–98. https://doi.org/10.1145/2696454.2696469

[14] Art Graesser, Yasuhiro Ozuru, and Jeremiah Sullins. 2009. What is a good question? In *Bringing reading research to life*, Margaret G. McKeown and Linda Kucan (Eds.). Guilford Press, New York, NY, Chapter 7, 170–193.

[15] Soheil Habibian, Ananth Jonnavittula, and Dylan P. Losey. 2022. Here’s What I’ve Learned: Asking Questions That Reveal Reward Learning. *J. Hum.-Robot Interact.* 11, 4, Article 40 (sep 2022), 28 pages. https://doi.org/10.1145/3526107

[16] Robert D. Hawkins, Andreas Stuhlmüller, Judith Degen, and Noah D. Goodman. 2015. Why do you ask? Good questions provoke informative answers. *Cognitive Science* (2015).

[17] Ranjay Krishna, Donsuk Lee, Li Fei-Fei, and Michael S. Bernstein. 2022. Socially situated artificial intelligence enables learning from human interaction. *Proceedings of the National Academy of Sciences* 119, 39 (2022), e2115730119. https://doi.org/10.1073/pnas.2115730119

[18] George Loewenstein. 1994. The psychology of curiosity: A review and reinterpretation. *Psychological bulletin* 116, 1 (1994), 75. https://doi.org/10.1037/0033-2909.116.1.75

[19] Rashid Minhas, Camilla Elphick, and Julia Shaw. 2022. Protecting victim and witness statement: examining the effectiveness of a chatbot that uses artificial intelligence and a cognitive interview. *AI & SOCIETY* 37, 1 (2022), 265–281. https://doi.org/10.1007/s00146-021-01165-5

[20] Ishan Misra, Ross Girshick, Rob Fergus, Martial Hebert, Abhinav Gupta, and Laurens Van Der Maaten. 2018. Learning by asking questions. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*. IEEE, Salt Lake City, UT, USA, 11–20. https://doi.org/10.1109/CVPR.2018.00009

[21] Samira Rasouli, Garima Gupta, Elizabeth Nilsen, and Kerstin Dautenhahn. 2022. Potential applications of social robots in robot-assisted interventions for social anxiety. *International Journal of Social Robotics* 14 (2022), 1–32. https://doi.org/10.1007/s12369-021-00851-0

[22] Celine Ray, Francesco Mondada, and Roland Siegwart. 2008. What do people expect from robots?. In *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, Nice, France, 3816–3821. https://doi.org/10.1109/IROS.2008.4650714

[23] Anselm Rothe, Brenden M. Lake, and Todd M. Gureckis. 2017. Question Asking as Program Generation. In *Advances in Neural Information Processing Systems*, I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.), Vol. 30. Curran Associates Inc., Long Beach, CA, USA.

[24] Anselm Rothe, Brenden M Lake, and Todd M Gureckis. 2018. Do people ask good questions? *Computational Brain & Behavior* 1, 1 (2018), 69–89. https://doi.org/10.1007/s42113-018-0005-5

[25] Thomas Scialom and Jacopo Staiano. 2020. Ask to Learn: A Study on Curiosity-driven Question Generation. In *Proceedings of the 28th International Conference on Computational Linguistics*. International Committee on Computational Linguistics, Barcelona, Spain (Online), 2224–2235. https://doi.org/10.18653/v1/2020.coling-main.202

[26] Juliette Vazard and Catherine Audrin. 2022. The noetic feeling of confusion. *Philosophical Psychology* 35, 5 (2022), 757–770. https://doi.org/10.1080/09515089.2021.2016675

[27] Hamed Zamani, Susan Dumais, Nick Craswell, Paul Bennett, and Gord Lueck. 2020. Generating Clarifying Questions for Information Retrieval. In *Proceedings of The Web Conference 2020* (Taipei, Taiwan) *(WWW ’20)*. Association for Computing Machinery, New York, NY, USA, 418–428. https://doi.org/10.1145/3366423.3380126

[28] Yi Zhu, Yue Weng, Fengda Zhu, Xiaodan Liang, Qixiang Ye, Yutong Lu, and Jianbin Jiao. 2021. Self-Motivated Communication Agent for Real-World Vision-Dialog Navigation. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*. IEEE, Montreal, QC, Canada, 1574–1583. https://doi.org/10.1109/ICCV48922.2021.00162