

How Should Agents Ask Questions For Situated Learning? An Annotated Dialogue Corpus

Felix Gervits¹, Antonio Roque², Gordon Briggs³, Matthias Scheutz², Matthew Marge¹

¹U.S. Army Research Laboratory, Adelphi, MD 20783

²Tufts University, Medford, MA 02155

³U.S. Naval Research Laboratory, Washington, DC 20375

{felix.gervits, matthew.r.marge}.civ@mail.mil,

{antonio.roque, matthias.scheutz}@tufts.edu

gordon.briggs@nrl.navy.mil

Abstract

Intelligent agents that are confronted with novel concepts in situated environments will need to ask their human teammates questions to learn about the physical world. To better understand this problem, we need data about asking questions in situated task-based interactions. To this end, we present the Human-Robot Dialogue Learning (HuRDL) Corpus - a novel dialogue corpus collected in an online interactive virtual environment in which human participants play the role of a robot performing a collaborative tool-organization task. We describe the corpus data and a corresponding annotation scheme to offer insight into the form and content of questions that humans ask to facilitate learning in a situated environment. We provide the corpus as an empirically-grounded resource for improving question generation in situated intelligent agents.

1 Introduction

Situated interaction is an area of interest to the Dialogue Systems community (Bohus, 2019), with recent papers investigating aspects of language interaction in situated environments both empirically and computationally (Gervits et al., 2020; Gupta et al., 2019; Kalpakchi and Boye, 2019; Kleingarn et al., 2019). This topic is critical for the development of technologies that interact with humans in real and virtual environments, including automated vehicles, smart home appliances, robots, and others. Situated agents are typically deployed in open-world environments and possess multiple sensory modalities, so a critical challenge involves enabling such agents to manage uncertainty across modalities and to learn about unfamiliar concepts. By engaging in dialogue with human interlocutors, these challenges can be addressed through effective clarification requests (Chernova and Thomaz, 2014). However, it is not clear what form these clarifications need to take to be most effective. Building

on previous corpus-based methods (Attari et al., 2019; Ginzburg et al., 2019; Gupta et al., 2020; Fuscone et al., 2020; Thomason et al., 2020), we address this question through the development of a human-human corpus of a collaborative pick-and-place task, and a corresponding annotation scheme of question form and content. The underlying assumption is that the kinds of questions people ask in this task provide good empirical support for indicators to guide agent questions in similar domains, and will inspire approaches for automated question generation moving forward.

2 Background

Prior work has investigated misunderstanding and clarification in human dialogue (Schegloff et al., 1977; Clark, 1996; Marge and Rudnicky, 2015; Paek, 2003). In one such analysis, Purver et al. (2003) proposed a scheme of clarification request forms that was applied to a 150,000 word subset of the British National Corpus and shown to cover 99% of the sub-corpus. While this scheme was applied to a dialogue system for automated clarification generation (Purver, 2004, 2006), it has been criticized for being too general and not accounting for certain types of phenomena such as pragmatic uncertainty (Rieser and Moore, 2005). Another classification scheme for clarification requests was introduced by Rodríguez and Schlangen (2004), building on Schlangen (2004)'s categorization of clarification causes. This scheme showed good coverage when applied to a data set of 22 dialogues from the Bielefeld corpus of task-oriented dialogue. While these analyses are useful for furthering our understanding of clarification requests, prior schemes did not consider situated domains with high degrees of uncertainty. Moreover, they mainly focused on the *form and function* rather than the *content* of clarifications. The content in-




Figure 1: First-person graphical user interface (GUI) used by participants to tele-operate the robot in the study. A message box for communicating with the Commander is in the bottom left, and an “inventory” showing the currently-held object is in the bottom right.

formation reflects the particular type of uncertainty that the agent experienced, and tracking that uncertainty helps to inform how agents can use questions (and answers) to manage that uncertainty.

The contributions of the current work are the following: (1) the presentation of a new annotated corpus, which has been made available for research purposes¹; (2) an annotation scheme that extends prior schemes to domains involving situated interaction, and also accounts for clarification requests generated to reduce uncertainty across modalities, such as visual feature clarification; (3) an analysis of the corpus including the distribution of categories from our scheme, along with a discussion of how these results can be used to improve the learning capabilities of situated dialogue agents.

3 Corpus Collection

The HuRDL corpus task was designed to investigate how agents can effectively generate questions to clarify uncertainty across multiple modalities in a task-based interaction. The task domain was designed to naturally present participants with novel concepts and procedural knowledge that they needed to learn; in doing so, they would need to use a variety of question types.

3.1 Collaborative Tool Organization Task

In the task, the human-controlled robot was placed in a virtual spacecraft (see Figure 1). The task was to organize six tools (among 12 distractors) scattered around the spacecraft – an activity that is

¹The corpus and additional details can be found at <https://github.com/USArmyResearchLab/ARL-HuRDL>.



Figure 2: Interface used by the human confederate playing the Commander. The main view shows the robot (circled in orange) in the environment. The top right corner shows the first-person view seen by the human participant playing the robot. The bottom left corner shows a message box for text dialogues with the participant.

relevant for current and future space robotics applications (Bualat et al., 2015). The tools had to be placed in the proper container (including crates, cabinets, and lockers), some of which were locked and required learning specialized procedures to open. The tools all had fictitious names to ensure that participants were unfamiliar with them, and the tools varied along a number of feature dimensions including color, shape, size, texture, symbol, and pattern. To facilitate the learning process, people could ask questions of a remotely-located Commander (who was played by a human study confederate; see Figure 2) in a live text-based dialogue.

To explore the effects of several dialogue-level factors, we manipulated *Speaker Initiative* and *Instruction Granularity*, as these have been shown to be relevant for human-agent dialogue (Baraglia et al., 2016; Gervits et al., 2016; Marge et al., 2020). The Commander took the initiative and gave scripted instructions for half the participants (*Commander Initiative* or *CI*) and only responded to questions in the other half (*Robot Initiative* or *RI*). Half the trials for each participant involved high-level granularity (“The sonic optimizer goes in the secondary cabinet, shelf A”) and half involved low-level granularity (“Move to locker Z” → “Pick up the sonic optimizer from the top shelf” → “Move to the secondary cabinet” → “Place the sonic optimizer on shelf A”). In all conditions, the confederate responded to clarification requests with a set policy which generally only provided minimal information.

3.2 Interactive Study Platform

The study was run on Amazon Mechanical Turk (MTurk), which was used for recruitment, questionnaires, and linking to the study environment. To support the proposed study, we developed an infrastructure that enabled interactivity between participants and the experimenter. The environment was developed in Unity 3D and built in WebGL (a browser-based graphics library)². Photon Unity Networking (PUN) was used to support communication between participants and the experimenter and also for synchronizing objects between both views. We used a Willow Garage PR2 robot model and allowed participants to tele-operate the robot directly using the keyboard.

Twenty-two participants recruited from MTurk performed the task. Eleven participants were female, and the average age was 36.8 ± 7.14 . All participants were native English speakers from US zip codes. Participants volunteered by clicking a link on the MTurk page. They then read detailed instructions and performed a tutorial to ensure that they understood the controls and instructions. The tutorial involved a simplified version of the main task with a live experimenter and four simple objects to place (different from the task stimuli). Following successful completion of the tutorial, the task was then performed, which generally took 30–45 minutes. Participants were paid \$10 for completing the study with a possible additional \$2 performance bonus. Video data was recorded of the robot movement and action in the environment, and a transcript of the dialogue was logged. The following measures were taken: *task performance* based on the percentage of the six task-relevant objects placed correctly, *task duration* based on how long it took to complete the task, *questions / total utterances* which indicates the proportion of questions in the dialogue, and *proportion of question types* based on the scheme described in Section 4.

4 HuRDL Corpus Overview

The HuRDL corpus contains twenty-two dialogues with a total duration of 13 hours. It contains a total of 1122 participant utterances, 760 of which are questions. Each dialogue has a mean of 51 participant utterances, 34 of which are questions. The mean score on the task is $77.3\% \pm 24\%$ and

²This WebGL setup was ideal for MTurk since it rendered directly in the browser and could be linked to from the study page.

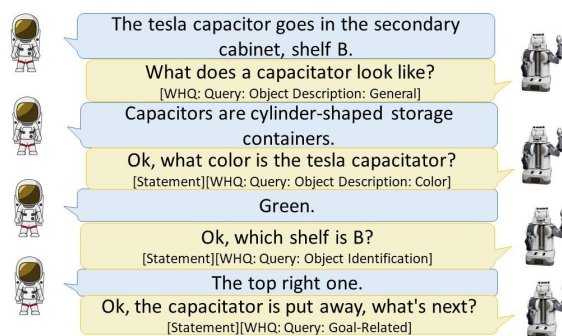


Figure 3: Example dialogue showing a single exchange along with annotations of the participant utterances.

the average duration is 35.2 ± 7 min. An example dialogue (with annotation) is shown in Figure 3.

To analyze question types, two annotators labeled the twenty-two dialogues as described below. The annotators began by using one of the dialogues to develop an annotation scheme for the form and clarification type categories described below; this scheme was then refined by performing a consensus annotation on a second dialogue. Inter-rater reliability was calculated by having both raters annotate the same four dialogues. Overall, there was 82.2% raw agreement between annotators, with a Cohen’s $\kappa = .79$. For non-statement utterances there was 82.9% raw agreement with Cohen’s $\kappa = .81$.

4.1 Annotations: Utterance Forms

First, utterances were labeled with their *form*, which was one of several categories: *yes/no-questions (YNQs)* are questions that elicit a yes or no response; *alternative questions (AQs)* are questions that present a list of options; *wh- questions (WHQs)* ask who, what, where, when, why, which, or how; and Statements are non-questions.

The annotators found that of this corpus’ utterances, 15.3% are YNQs, 2.8% are AQs, 49.6% are WHQs, and 35% are Statements. These add up to more than 100% because some complex utterances contained multiple question forms, or contained both a statement and a question.

4.2 Annotations: Clarification Types

Second, utterances were labeled with their *clarification type* using the types shown in Table 1. By adding question content, this approach expands on previous utterance taxonomies, such as the one presented by Rodríguez and Schlangen (2004).

The annotators noted that the clarification type annotations were guided by (but not universally de-

Table 1: Distribution of Clarification Type Annotations Per Total Utterances.

Type	%	Example
CONFIRMATION QUESTIONS		
<i>Confirm Object by Location:</i>	5.9	
Spatial	3.8	<i>The one on the left?</i>
Proximity	0.09	<i>The nearby one?</i>
Landmark	1.6	<i>The one next to the wall?</i>
Deictic Action	0.36	<i>The one I'm holding?</i>
Other	0.09	<i>On the second half?</i>
<i>Confirm Object by Feature:</i>	8.9	
Size	0.62	<i>You mean the tall one?</i>
Shape	0.62	<i>The narrow one?</i>
Color	3.7	<i>The green one?</i>
Pattern	0.36	<i>The striped one?</i>
Symbol	0.45	<i>The one with a circle?</i>
Hybrid	1.78	<i>The green cylinder?</i>
Comparison	0.62	<i>The one that looks like a snake?</i>
Other	0.80	<i>The one with numbers on it?</i>
<i>Confirm Action:</i>	4.0	
General	0.27	<i>Did I do it right?</i>
Task-Related	3.6	<i>Does the block go in the locker?</i>
Other	0.18	<i>Does this light up?</i>
QUERIES		
<i>Object Description:</i>	10.9	
General	4.8	<i>What does it look like?</i>
Size	0.18	<i>What size is it?</i>
Shape	0.18	<i>What shape is it?</i>
Color	5.3	<i>What color is it?</i>
Pattern	0.27	<i>What pattern is it?</i>
Symbol	0.18	<i>What symbol is on it?</i>
<i>Location-Related</i>	12.3	<i>Where is that one?</i>
<i>Object Identification</i>	6.1	<i>Which one is that?</i>
<i>Object Naming</i>	0.98	<i>What is the small one called?</i>
<i>Goal-Related</i>	9.1	<i>What's next?</i>
<i>Request Teaching:</i>	11.2	
General	3.3	<i>How do I open lockers?</i>
Target	7.9	<i>What's the code for crate 3?</i>

terminated by) the form annotations: YNQs and AQs tended to be confirmation questions, and WHQs tended to be queries. However, the corpus contains several interesting exceptions to this. For example, consider the corpus utterance: “shelf D I assume is the bottom right one.” Although the utterance form is a statement, the utterance is an indirect speech act (Searle, 1975) functioning as a question that is seeking to confirm an object according to its location, and doing so with a spatial reference.

As shown in Table 1, clarification types can either be confirmation questions or queries. Confirmation questions can be one of three main classes, either confirming an object based on its location or a feature, or confirming an action. Each of these has additional sub-classes further specifying the confirmation. Queries can be one of several different classes, some of which are non-confirmation questions related to reference resolution, and some of which are requests for task-related instruction (either about opening lockers, or asking the next step in the task). Several of these also have additional sub-classes.

5 Results and Discussion

In our analysis of the corpus, we discovered several key findings about how people ask questions under uncertainty.

5.1 Question Types

In terms of utterance forms, a one-way MANOVA showed significant differences between the mean proportion of the three main utterance forms (collapsed across all conditions) by total participant utterances, $F(2,63) = 33.98, p < .001$. Post-hoc tests using the Bonferroni correction revealed that YNQs were the least frequent, followed by Statements, and then WHQs; ps for all comparisons $< .05$. Given that half of all utterances were WHQs, this finding suggests that WHQs are key questions used by people to reduce uncertainty in this domain.

In terms of clarification types, queries were by far the most common, accounting for 73% of participant utterances. Interestingly, we found a strong negative correlation between *Location-Related* queries / total questions and task performance in the CI condition, $r(8) = -.720, p < .05$. That is, the more location-related questions people asked the worse they performed on average. This correlation could reflect ineffective questions. For example, the experimenter did not know where an object was located, so questions such as “Where is X?” were generally ineffective.

Compared to the corpus analysis from Rodríguez and Schlangen (2004) in which 52% of clarification requests were related to referential ambiguity, in our corpus this was about 75%. In their corpus, 45% of response utterances were YN answers (suggesting a similar proportion of YNQs), whereas in ours, only 15% of the questions were YNQs. This was likely a result of the novel objects in our task, which led to more queries. Compared to the corpus analysis in Purver et al. (2003), our results indicate a large proportion of *non-reprise clarifications*, i.e., explicit questions that do not echo or repeat the instruction. We also found fewer disfluencies in our analysis due to it being written communication. The few observed ones were mostly typos and fragments. Finally, Cakmak and Thomaz (2012) found that 82% of all questions in their learning task were feature-based, whereas we observed about 20% of questions in this category. This can be attributed to differences in task domain and participant familiarity with the environment.

5.2 Design and Research Implications

This corpus analysis provides evidence that human dialogue strategies to manage uncertainty can be used to inform the development of real-time, online learning algorithms for agents in situated interaction. Our results directly inform the development of such algorithms in several ways. First, they outline the distribution of question forms and types that people used to manage uncertainty. This distribution (see Table 1) serves as a guideline about which kinds of questions to use, especially when encountering specific kinds of uncertainty. For example, if multiple objects have the same color, an agent can generate a color query to disambiguate. Second, the results capture the surface form of the questions, i.e., how they were realized. This enables the corpus to serve as a training set for data-driven dialogue systems that can learn to generate questions based on input instructions and their own uncertainty representations.

While the human data may serve as a good guideline for agent clarification, it is important to acknowledge the limitations of applying the results too directly. For example, people tended to ask about features that were (1) salient and (2) interpretable. Since salience for agents is likely different than for humans, the content of their queries should adjust accordingly. Moreover, though general object descriptions and object identifications were used by humans in the task, they should perhaps be limited by agents since they may not have the perceptual capabilities to interpret the response to general questions such as “What does it look like?” Instead, feature-based queries that the agent can interpret may be more effective. Prior work in active learning has highlighted the benefit of feature queries (e.g., Bullard et al. (2018)), however the present work is complementary to such approaches in that it serves as an empirical basis for the form and content of questions that robots should ask once the learning algorithms have determined the missing information.

Moving forward, the HuRDL corpus has utility as a test bed for further exploration into clarification requests in domains with high uncertainty across modalities. Future work will explore dialogue strategies (i.e., patterns of question types) used by different participants and compare their effectiveness. Additional corpus analysis can investigate other factors that influence question generation including the effects of time pressure, workload,

and object properties. Moreover, the video analysis can reveal the visual input that people had access to and the influence that this had on question generation.

It is important that future work apply our annotation scheme to different task domains to establish generalizability. While our task emphasized uncertainty of novel entities to elicit questions, other tasks may not. As a result, we do not expect such a high frequency of questions in other kinds of tasks, nor do we expect the same distribution of question types. The scheme, however, should capture the scope of questions used in a broad range of situated learning tasks since the categories represent general properties by which objects can be identified and distinguished from one another.

6 Conclusion

To investigate the problem of how agents can most effectively ask questions in a situated interaction, we analyzed dialogue data from the HuRDL corpus that we collected. The task involved uncertainty across multiple modalities and led to a variety of clarification questions to manage this uncertainty. We categorized these questions in a novel scheme and used it to annotate the corpus. Analysis of question types showed that people used a high frequency of WH-questions, and that these were targeted at learning object features and locations, object task-relevance, goals, and procedural knowledge. These patterns were influenced by dialogue-level factors such as speaker initiative and instruction granularity. Given these results, we presented guidelines to inform automated approaches to effective question generation, which will help make situated agents more resilient in uncertain environments. Future work will develop algorithms for clarification based on the question types and dialogue strategies identified in this corpus.

Acknowledgments

This research was sponsored by the Basic Research Office of the U.S. Department of Defense with a Laboratory University Collaboration Initiative Fellowship awarded to MM. The authors would like to thank Genki Kadomatsu and Dean Thurston for their contributions to the study platform.

References

- Nazia Attari, Martin Heckmann, and David Schlangen. 2019. [From explainability to explanation: Using a dialogue setting to elicit annotations with justifications](#). In *Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue*, pages 331–335, Stockholm, Sweden. Association for Computational Linguistics.
- Jimmy Baraglia, Maya Cakmak, Yukie Nagai, Rajesh Rao, and Minoru Asada. 2016. [Initiative in robot assistance during collaborative task execution](#). In *Proceedings of the 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 67–74. IEEE.
- Dan Bohus. 2019. [Situated interaction. Keynote presentation](#). Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue.
- Maria Bualat, Jonathan Barlow, Terrence Fong, Chris Provencher, and Trey Smith. 2015. [Astrobee: Developing a free-flying robot for the international space station](#). In *Proceedings of the AIAA SPACE 2015 Conference and Exposition*, page 4643.
- Kalesha Bullard, Andrea L Thomaz, and Sonia Chernova. 2018. [Towards intelligent arbitration of diverse active learning queries](#). In *Proceedings of the 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6049–6056. IEEE.
- 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)*, pages 17–24. IEEE.
- Sonia Chernova and Andrea L Thomaz. 2014. [Robot learning from human teachers](#). *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 8(3):1–121.
- Herbert H Clark. 1996. *Using Language*. Cambridge University Press.
- Simone Fuscone, Benoit Favre, and Laurent Prevot. 2020. [Filtering conversations through dialogue acts labels for improving corpus-based convergence studies](#). In *Proceedings of the 21st Annual SIGdial Meeting on Discourse and Dialogue*, pages 203–208, 1st virtual meeting. Association for Computational Linguistics.
- Felix Gervits, Kathleen Eberhard, and Matthias Scheutz. 2016. [Team communication as a collaborative process](#). *Frontiers in Robotics and AI*, 3:62.
- Felix Gervits, Ravenna Thielstrom, Antonio Roque, and Matthias Scheutz. 2020. [It’s about time: Turn-entry timing for situated human-robot dialogue](#). In *Proceedings of the 21st Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 86–96, 1st virtual meeting. Association for Computational Linguistics.
- Jonathan Ginzburg, Zulipiye Yusupujiang, Chuyuan Li, Kexin Ren, and Pawel Lupkowski. 2019. [Characterizing the response space of questions: a corpus study for English and Polish](#). In *Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue*, pages 320–330, Stockholm, Sweden. Association for Computational Linguistics.
- Itika Gupta, Barbara Di Eugenio, Brian Ziebart, Aiswarya Baiju, Bing Liu, Ben Gerber, Lisa Sharp, Nadia Nabulsi, and Mary Smart. 2020. [Human-human health coaching via text messages: Corpus, annotation, and analysis](#). In *Proceedings of the 21st Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 246–256, 1st virtual meeting. Association for Computational Linguistics.
- Prakhar Gupta, Shikib Mehri, Tiancheng Zhao, Amy Pavel, Maxine Eskenazi, and Jeffrey Bigham. 2019. [Investigating evaluation of open-domain dialogue systems with human generated multiple references](#). In *Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue*, pages 379–391, Stockholm, Sweden. Association for Computational Linguistics.
- Dmytro Kalpakchi and Johan Boye. 2019. [SpaceRefNet: a neural approach to spatial reference resolution in a real city environment](#). In *Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue*, pages 422–431, Stockholm, Sweden. Association for Computational Linguistics.
- Diana Kleingarn, Nima Nabizadeh, Martin Heckmann, and Dorothea Kolossa. 2019. [Speaker-adapted neural-network-based fusion for multimodal reference resolution](#). In *Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue*, pages 210–214, Stockholm, Sweden. Association for Computational Linguistics.
- Matthew Marge, Felix Gervits, Gordon Briggs, Matthias Scheutz, and Antonio Roque. 2020. [Let’s do that first! A comparative analysis of instruction-giving in human-human and human-robot situated dialogue](#). In *Proceedings of the 24th Workshop on the Semantics and Pragmatics of Dialogue (SemDial)*.
- Matthew Marge and Alexander Rudnicky. 2015. [Miscommunication recovery in physically situated dialogue](#). In *Proceedings of the 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 22–31.
- Tim Paek. 2003. [Toward a taxonomy of communication errors](#). In *ISCA Tutorial and Research Workshop on Error Handling in Spoken Dialogue Systems*.
- Matthew Purver. 2004. [Clarie: The clarification engine](#). In *Proceedings of the 8th Workshop on the Semantics and Pragmatics of Dialogue (Catalog)*, pages 77–84. Citeseer.

- Matthew Purver. 2006. [Clarie: Handling clarification requests in a dialogue system](#). *Research on Language and Computation*, 4(2-3):259–288.
- Matthew Purver, Jonathan Ginzburg, and Patrick Healey. 2003. [On the means for clarification in dialogue](#). In *Current and New Directions in Discourse and Dialogue*, pages 235–255. Springer.
- Verena Rieser and Johanna D Moore. 2005. [Implications for generating clarification requests in task-oriented dialogues](#). In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL’05)*, pages 239–246.
- Kepa Joseba Rodríguez and David Schlangen. 2004. [Form, intonation and function of clarification requests in german task-oriented spoken dialogues](#). In *Proceedings of Catalog (the 8th workshop on the semantics and pragmatics of dialogue; SemDial04)*.
- Emanuel A. Schegloff, Gail Jefferson, and Harvey Sacks. 1977. [The preference for self-correction in the organization of repair in conversation](#). *Language*, 53(2):361–382.
- David Schlangen. 2004. [Causes and strategies for requesting clarification in dialogue](#). In *Proceedings of the 5th Annual SIGdial Workshop on Discourse and Dialogue at HLT-NAACL 2004*, pages 136–143.
- John R Searle. 1975. [Indirect speech acts](#). *Syntax & Semantics*, 3: *Speech Act*, pages 59–82.
- Jesse Thomason, Michael Murray, Maya Cakmak, and Luke Zettlemoyer. 2020. [Vision-and-dialog navigation](#). In *Proceedings of the Conference on Robot Learning*, volume 100 of *Proceedings of Machine Learning Research*, pages 394–406. PMLR.