

Decision-Theoretic Question Generation for Situated Reference Resolution: An Empirical Study and Computational Model

Felix Gervits
DEVCOM Army Research Laboratory
Burlington, MA, USA
felix.gervits.civ@army.mil

Gordon Briggs
U.S. Naval Research Laboratory
Washington, DC, USA
gordon.briggs@nrl.navy.mil

Antonio Roque
Tufts University
Medford, MA, USA
antonio.roque@tufts.edu

Genki A. Kadomatsu*
Tufts University
Medford, MA, USA
genki.kadomatsu@tufts.edu

Dean Thurston
Tufts University
Medford, MA, USA
dean.thurston@tufts.edu

Matthias Scheutz
Tufts University
Medford, MA, USA
matthias.scheutz@tufts.edu

Matthew Marge
DEVCOM Army Research Laboratory
Adelphi, MD, USA
matthew.r.marge.civ@army.mil

ABSTRACT

Dialogue agents that interact with humans in situated environments need to manage referential ambiguity across multiple modalities and ask for help as needed. However, it is not clear what kinds of questions such agents should ask nor how the answers to such questions can be used to resolve ambiguity. To address this, we analyzed dialogue data from an interactive study in which participants controlled a virtual robot tasked with organizing a set of tools while engaging in dialogue with a live, remote experimenter. We discovered a number of novel results, including the distribution of question types used to resolve ambiguity and the influence of dialogue-level factors on the reference resolution process. Based on these empirical findings we: (1) developed a computational model for clarification requests using a decision network with an entropy-based utility assignment method that operates across modalities, (2) evaluated the model, showing that it outperforms a slot-filling baseline in environments of varying ambiguity, and (3) interpreted the results to offer insight into the ways that agents can ask questions to facilitate situated reference resolution.

CCS CONCEPTS

• **Human-centered computing** → Empirical studies in HCI; • **Computing methodologies** → Discourse, dialogue and pragmatics; **Bayesian network models**.

KEYWORDS

human-agent dialogue, situated interaction, reference resolution

*The work was conducted as part of an internship at ARL.

ACM acknowledges that this contribution was authored or co-authored by an employee, contractor, or affiliate of the United States government. As such, the United States government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for government purposes only.

ICMI '21, October 18–22, 2021, Montréal, QC, Canada

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8481-0/21/10...\$15.00

<https://doi.org/10.1145/3462244.3479925>

ACM Reference Format:

Felix Gervits, Gordon Briggs, Antonio Roque, Genki A. Kadomatsu, Dean Thurston, Matthias Scheutz, and Matthew Marge. 2021. Decision-Theoretic Question Generation for Situated Reference Resolution: An Empirical Study and Computational Model. In *Proceedings of the 2021 International Conference on Multimodal Interaction (ICMI '21)*, October 18–22, 2021, Montréal, QC, Canada. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3462244.3479925>

1 INTRODUCTION

Effective human-agent interaction is critical for many application domains ranging from customer service to space missions. In non-laboratory environments these agents frequently encounter novel concepts (and words, objects, actions, etc.) with which they are unfamiliar. This is unavoidable due to difficulties in collecting and labeling training data, but the problem is magnified when agents are placed in novel and unexplored environments. For example, a robot for space exploration may encounter objects and procedures that do not exist on Earth. Moreover, ambiguity often arises in multiple modalities (e.g., language and vision), and so agents will need mechanisms to resolve ambiguity across all of these modalities.

One solution to this problem is to enable agents to engage in dialogue and ask questions of a human interlocutor to reduce ambiguity. This is particularly important for the task of reference resolution, in which an agent needs to identify a target referent that it may know little or nothing about. However, not all questions that agents ask are equally effective or efficient, and we want those that ask *good* questions [3]. To understand what makes good questions, we analyzed data from a corpus in which people controlled a robot to carry out a collaborative tool organization task while asking questions to a commander (an experimenter). The corpus task was designed to contain high amounts of referential ambiguity across multiple modalities with respect to entities in the environment, and thus to elicit as many questions as possible. This builds on the assumption that the kinds of questions people ask in this task are good models for dialogue agents in collaborative task domains.

The contributions of this paper are twofold. First, we conducted several empirical analyses of the dialogue data from the corpus. This included categorizing the distribution of question types used in the task, and exploring the effect of various dialogue-level factors on question generation, including speaker initiative and instruction granularity. Despite being implicated in the clarification process, few if any studies have explored the role that these factors play, especially in the context of a human-robot interaction (HRI) task. Next, from the dialogue data in the corpus, we developed a computational model of situated reference resolution for dialogue agents using a decision network. Decision-theoretic approaches have shown promise in modeling some aspects of conversational grounding under multi-modal uncertainty [19, 37]. However, to our knowledge they have not been employed for the task of generating questions for reference resolution. The advantage of such an approach is that it is effective without pre-training, intrinsically explainable, and robust to uncertainty across multiple modalities. Our model is also generalizable in that it dynamically generates a decision network from perceptual input and parameterizes the network based on the entities and properties in the domain in real time. This enables agents to operate in unseen environments for which they have limited or no knowledge - a property that is rare among current models.

In the next section, we survey work that studies how people ask and respond to questions, and what approaches exist for modeling question generation for situated reference resolution. In Section 3, we describe our empirical study to investigate how robots should ask questions to support disambiguation of novel entities, and present the results. In Section 4 we introduce our decision-theoretic approach informed by the study data, and we evaluate it in Section 5. Finally, contributions and future work are discussed in Section 6.

2 BACKGROUND

In tasks where one dialogue partner is trying to learn from another, and more broadly in any dialogue task where grounding occurs, clarification is an essential process. Clarification primarily emerges in dialogue where common ground contributions are in the process of being established [6] or when there is a misunderstanding or non-understanding [18]. Recovery strategies from these grounding failures have been theorized and studied [36, 40]. As a special case of grounding, reference resolution entails a collaboration where a listener tries to determine the *referent* that a speaker intends [7]. **Our work investigates the decision-making that an agent would make toward optimal clarifications when trying to resolve referents with a human in a collaborative way.**

2.1 Question Selection and Clarification in Situated Human-Agent Dialogue

Humans often use a variety of question forms to facilitate mutual understanding in task settings, so it is natural for situated agents to ask questions of their human interlocutors. One past focus has involved robots. Indeed, robot-generated clarification requests are increasingly being used to facilitate data collection for active learning algorithms [2, 41]. [4] provided guidelines for robot learning, including that the robot must be able to ask questions that reflect

its mental model, such as missing information, task uncertainty, etc. This raises the question of what robot question-asking strategies are most effective for facilitating successful learning. Addressing this question is the focus of the analysis in Section 3.

The impact of dialogue factors on question generation has not been adequately explored to our knowledge. For example, it is known that factors such as *speaker initiative* are related to role [21] and interactivity [8], which has implications for how questions are asked. Initiative has been explored in dialogue and discourse modeling [43], and in its effects on dialogue management in various domains [5, 15], as well as in HRI studies [1, 34]. Nevertheless, we are unaware of any empirical studies that manipulated initiative to investigate question generation. Another factor that may influence clarifications is *instruction granularity*. Instructions can be produced at different levels of granularity, ranging from discrete low-level commands to high-level plans. Though [29] found that people are more likely to use low-level commands when instructing a robot versus a human, robots still need to be equipped with the capability to respond to high-level intentions [46]. Few if any studies have examined the role that speaker initiative and instruction granularity play in the process of clarification generation, especially in novel environments characterized by multi-modal ambiguity.

2.2 Situated Reference Resolution

The task of asking good questions relates to the body of work on situated reference resolution. Early work established requirements and information sources for multi-modal reference resolution [10, 11, 16, 26, 42]. Graph-based approaches that represent referents and their semantic properties against an environment have shown promise but are generally restricted to controlled environments such as tabletops or game boards [23, 28, 49].

Probabilistic models are another popular approach [12, 14, 45]. While effective at identifying referents from a set of entities, they may lead to dispreferred questions, compared to approaches that leverage features or knowledge about the environment [47]. The same drawback holds for POMDP-based approaches [27, 32, 44]. While they are generally quite effective for dialogue state tracking, they are less effective at generating clarification questions for collaborative reference resolution, especially in novel environments. This is because POMDPs typically require a predefined state-action space, but the relevant properties of a given environment cannot always be known in advance. As a result, these approaches tend to struggle with open-world interaction. Neural approaches have recently been developed for situated, multi-modal dialogue [20, 22, 24, 33, 39]. While promising, these approaches typically require large, labeled training sets and are generally limited to the specific domains in which they were trained.

The present work most closely relates to the body of work on clarification for situated, collaborative reference resolution. Some work has explored the use of decision-theoretic approaches for this purpose, mainly for selecting questions that maximize expected utility [2, 25, 38]. Information-theoretic approaches have also been developed which rely on selecting queries that reduce entropy over a set of groundings in an instruction [9, 17]. However, we are not aware of work that fuses the two approaches. Our model uses decision networks supplemented with an entropy-reduction method




Figure 1: Participant view (first-person) and GUI, including a message box to communicate with the experimenter, and an inventory.

over object properties to enable an agent to initiate referential clarifications when there are multiple possible strategies. This promises to be (1) more generalizable and scalable than the algorithms described above because of its focus on decision-making over detailed object representations, (2) flexible enough to be used in open-world scenarios, and (3) natural, since the model is explicitly grounded in empirical studies of human-human dialogue.

3 HURDL CORPUS ANALYSIS

The HuRDL (Human-Robot Dialogue Learning) Corpus [13] was used as the basis for empirical investigation of question generation. The corpus task was designed to investigate how robots should ask questions to learn novel entities when facing multi-modal ambiguity. A question-type analysis was performed on the dialogue data to examine the frequency of various question types, including *WH-questions* (WHQs; who, what, where, when, why, how) and *Yes/No-Questions* (YNQ). These results served as the empirical basis for the computational model described in Section 4. Additional analyses examined the effect of dialogue-level factors on question generation (see Section 3.5).

3.1 Task Overview

In the corpus task, human participants played the role of a robot performing a situated reference resolution task. The study was run online using an interactive setup in which participants tele-operated a virtual PR2 robot. The goal of the task was to organize a number of previously-unseen tools scattered around the spacecraft. A complete study run consisted of six *trials*, each one corresponding to a single tool being located, then placed in the correct container. Participants used their keyboard to control the robot and pick up / place objects. They had a first-person view of the environment, along with an on-screen graphical user interface (GUI) that included a text box to communicate with a remote “commander” (an experimenter) through low-latency text chat (see Figure 1). The commander gave scripted instructions in each trial (to maximize experimental control) and responded to participant questions in a standardized manner.

3.2 Environment Properties

The environment consisted of an open room with various containers, interfaces, objects, and landmarks. The environment also contained six *tool types*, with three *instances* of each type. The tools were given fictitious names such that the names could not be used solely to identify the objects. The six tools were called: *optimizer*, *calibrator*, *module*, *synthesizer*, *capacitor*, and *emitter*. The tools had an additional prefix added to this label, resulting in object names like *megaband module* and *temporal emitter*. Only six instances were task-relevant, so the remaining twelve were distractors.

In order to elicit a variety of question types, the tools were designed to be heterogeneous with respect to a number of features. Six features were assigned to the tool instances, including *color* (red, yellow, blue, etc.), *shape* (short, tall, narrow, wide), *size* (small, medium, large), *texture* (wood, coarse, metal, smooth), *symbol* (x, +, -, etc. on the object), and *pattern* (striped, spotted, no pattern). Each tool type varied along three feature dimensions, which resulted in a total of 20 combinations of three-feature sets ($\binom{6}{3}$). The feature set of each of the six tool types was randomly selected from these combinations. For example, *modules* varied by pattern, shape, and symbol whereas *synthesizers* varied by color, size, and shape.

In addition to the tools, the environment also contained a number of *containers* (lockers, cabinets, and crates) and *interfaces*, which were used to open the containers. The lockers, cabinets, and crates were locked and required learning a procedure to open (e.g., entering a passcode, arranging keycards, and pressing buttons). Various landmarks were placed throughout, including a wall with a transparent window, text on the walls (e.g., X99), electrical components on the walls, windows, and lights. These landmarks were put in place to elicit spatial referring expressions and other questions unique to situated environments.

3.3 Experimental Conditions

The study used speaker initiative as a between-subjects factor, and instruction granularity as a within-subjects factor. The effect of these factors on question generation and performance is of particular importance since many human-robot interactions involve mixed-initiative dialogue [30] and varying degrees of instruction granularity [29]. Our analysis was designed to investigate the effect that these factors have on the frequency and form of robot-initiated questions.

Regardless of condition, all participants saw the same starting configuration of objects. The main difference was in the experimenter script and policy. In the Commander-Initiative (CI) condition, the commander provided each instruction (e.g., “The sonic optimizer goes in the secondary cabinet, shelf A”), whereas in the Robot Initiative (RI) condition, the participant had to establish which tools and locations were relevant.

Instruction granularity was manipulated by first preparing the high-level instructions and then converting these into low-level instructions by breaking them down into their constituent commands. For example, the high-level instruction, “The tesla capacitor goes in the secondary cabinet, shelf B” would become the following four low-level instructions: “Move to the primary cabinet” → “Pick up the tesla capacitor from shelf A” → “Move to the secondary

cabinet" → "Place the tesla capacitor on shelf B". In the CI condition, exactly three out of the six trials in each run were low-level instructions, and the other half were high-level instructions (the order was counterbalanced across participants).

3.4 Predictions

The following measures were used in the analysis: *task performance* - based on the percentage of the six task-relevant objects placed correctly, *task duration* - based on how long in seconds it took to complete the task, *questions / total utterances* - which indicates the frequency of questions in the dialogue, and *proportion of question types* - which indicates the relative frequency of each question type.

For our analysis, we had certain *a priori* predictions. First, we expected that questions would be very common due to the need to clarify ambiguity. However, WHQs would be more common since the answers to them provide more information than a YNQ. We expected people to ask approximately three questions on average to resolve a referential ambiguity, corresponding with the three properties by which the entities varied. We also expected people to ask about object descriptions using salient properties like color and shape, though we did not know which properties would be salient. While object features are more salient if a single feature can be used to distinguish them (e.g., "feature singletons" from [48]), we are not aware of studies finding greater salience of certain features in objects possessing heterogeneous features.

We predicted that trials involving high-level instructions would take longer compared to low-level trials due to the more complex, and uncertain, instructions. This is because low-level instructions are very explicit and typically elicit only a single type of uncertainty. On the other hand, high-level instructions may have several unknown terms, and people need to decide on a method to clarify these terms. Regarding questions, the study was designed to elicit a broad range of question types. While the aggregate results will inform the kinds of questions people generally ask, we expected to see more questions overall with the high-level instructions due to the increased complexity. Finally, we predicted that speaker initiative would affect performance and efficiency in that people in the CI condition would perform better, finish the task faster, and ask fewer questions. This is because the task goals are explicitly laid out and structured when compared to the RI condition, in which participants do not have explicit goals.

3.5 Empirical Results

The final dataset included dialogues from twenty-two participants, 10 in the CI condition and 12 in the RI condition. 11 of the participants were female and the rest were male. The average age was 37 ± 7 years and all participants were native English speakers. The mean number of questions to resolve an ambiguity was 1.72 ± 0.40 , suggesting that people were very efficient at asking disambiguating questions.

3.5.1 Distribution of Questions. We analyzed the transcribed dialogue data for the distribution of question types, which led to several interesting findings about how people ask questions to resolve referential ambiguity. As expected, we found that most utterances were questions, with 50% in the form of WHQs, 15% in the form of YNQs and 3% as AQs (*Alternative Questions*, which present

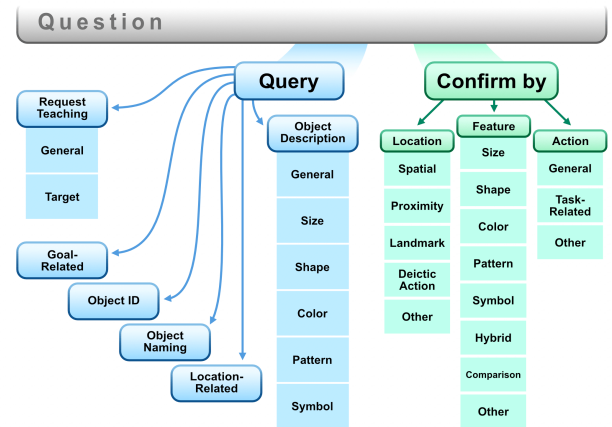


Figure 2: Question types used in the task. Queries generally corresponded to WHQs, whereas Confirm By's generally corresponded with YNQs.

a list of options); the rest were statements. This fit our prediction that WHQs would be the preferred question type. Though YNQs were less frequently used, feature-based confirmations were the most common subtype. In terms of feature preference, color was the most commonly asked-about feature, even though the objects were heterogeneous and only half of them varied by color. Among the location-based confirmations, spatial types were the most common followed by landmark types. Finally, action confirmations were mostly task-related (e.g., "Does this go in the locker?"). Figure 2 shows a taxonomy of question types observed in the data. These results will be used in Section 4 for selecting questions and setting the utilities of the model.

3.5.2 Instruction Granularity. Within the CI condition¹, instruction granularity affected the trial duration and questions per utterance. Specifically, we found a difference in high- vs low-level instructions on trial duration and questions per utterance as indicated by paired-sample t-tests. High-level trials had a longer duration ($M = 1063.3 \pm 274.2$) than low-level ($M = 748.4 \pm 228.3$) trials, $t(9) = 3.04$, $p < .05$. As we hypothesized, the less structured nature of high-level trials led to longer duration, suggesting that these instructions contained more uncertainty. In addition, high-level trials had a higher proportion of questions to utterances ($M = 0.74 \pm 0.08$) compared to low-level trials ($M = 0.58 \pm 0.20$), $t(9) = 2.26$, $p < .05$ even though the average number of questions and utterances did not differ between Granularity conditions ($ps > .05$). This was an expected result given that low-level trials often had acknowledgments after each low-level command, and these were counted as *Statements* in our scheme. There was no effect on performance or the average number of questions per trial ($ps > .05$).

3.5.3 Speaker Initiative. We ran a one-way ANOVA on dialogues from the 22 participants to examine the effect of speaker initiative. Contrary to our predictions, there was no effect of initiative on the number of questions, task performance, or task duration ($ps >$

¹This analysis was run only on the CI condition because instruction granularity was not controlled in the RI condition.

.05). We did however find significant differences in question types between the two Initiative conditions, specifically involving the proportion of *Confirm Actions* / total questions [$F(1,20) = 4.69, p < .05$] and *Goal-Related* / total questions [$F(1,20) = 4.49, p < .05$]. In both cases, the RI group displayed higher proportions of these question types, with *Confirm Actions* making up 9.8% of all questions and *Goal-Related* making up 17.8% of all questions (compared to 1.7% and 8.5% in the CI condition, respectively). As predicted, the lack of guidance in the RI condition led to the increased prevalence of these question types.

4 DECISION-THEORETIC MODEL

In this section we introduce a computational model of dialogue-based situated reference resolution that is informed by the empirical findings described in the previous section. The model uses a decision network [35] to represent the agent’s knowledge about relevant properties of the entities in the world (see Section 4.1). This knowledge can be deterministic (e.g., $knows(x)$) or probabilistic (e.g., $knows(x)_{p=.8}$). Using this approach, the agent selects and asks a single question about some property that maximizes its expected utility over all possible questions. It then interprets the response and continues asking questions in this way until it has resolved the referent. The list of question types was based on a subset of the ones in Figure 2, and the properties and initial utility table also came from the task environment.

While the structure of the decision network implementation reflects the particular properties of our task domain (e.g., shape, size, color, etc.), these are general object properties that can be used to disambiguate entities in a variety of domains. The approach, however, does not depend on these particular properties, and works for any arbitrary number and type of properties. The question types are also domain-independent, and are based on the properties. Moreover, we can construct the decision network in real time so that it can be used in a variety of domains without prior training (see Section 4.2).

4.1 Decision Network Specification

Decision networks are probabilistic graphical models related to Bayesian networks that encode conditional relationships between variables in a directed acyclic graph. They are effective at integrating multiple sources of uncertainty (e.g., vision, language, knowledge) and supporting probabilistic decision-making. Decision networks contain three types of nodes: *chance nodes* which are random variables, *decision nodes* which reflect actions that an agent can perform, and *utility nodes* which represent the utility function over those actions. The goal of the network is to find the action with the maximum expected utility (MEU) using Equation 1. To compute the expected utility of each action a (i.e., question), given evidence e , and all possible values of the chance nodes, s , the network multiplies the posterior probabilities, $P(s|e)$, by the utility weights for that action, $U(a, s)$. The posterior probabilities start out uniformly distributed, and are updated as the agent perceives and learns about additional object properties. The utility weights are set using an entropy-driven approach (see 4.1.3 below).

$$MEU(e) = \operatorname{argmax}_a \sum_s P(s|e)U(a, s) \quad (1)$$

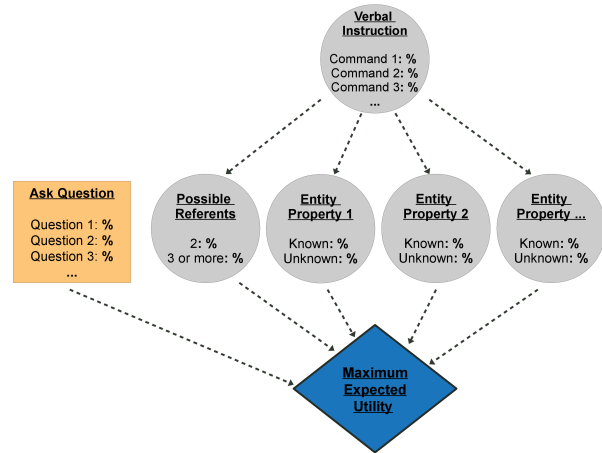


Figure 3: Decision Network to generate questions for situated reference resolution. The grey circles are chance nodes (random variables), the orange square is a decision node, and the blue diamond is a utility node. See text for more details.

The general structure of the decision network is depicted in Figure 3. The network contains the following node types:

4.1.1 Chance Nodes. These are random variables that represent (1) the possible verbal instructions, (2) the number of possible referents that correspond to the input command, and (3) the agent’s knowledge about the various properties that entities can possess in the environment. A separate node is used for each entity property, and represents a probability function of the agent’s knowledge of that property given the input commands. The prior probabilities for the *verbal instructions* chance node are uniformly distributed over the vocabulary of instructions in the task domain, and are updated when an instruction is received by the system. The prior probabilities for the *referents* and *properties* chance nodes are also uniform, but these values are updated based on the agent’s knowledge acquired during the task.

4.1.2 Decision Node. A decision node is a list of questions that the agent can ask. We selected the most frequently used twelve question types from the corpus, though the node can support an arbitrary number of questions. Nine WHQs were used, including all of the Object-Description queries, and three YNQs were selected, from those used at >1% frequency in the data (*Confirm:Color*, *Confirm:Spatial*, *Confirm:Landmark*).

4.1.3 Utility Node. A utility node is a conditional probability table that represents the utility of asking a question given the agent’s knowledge of the various properties and the number of ambiguous referents that correspond to the input. We explore two ways of setting these utilities, though others may also be used. The first is a *data-driven* method that involves using the corpus data to set the utilities. Specifically, we used the relative frequency (normalized) of each question type. The procedure sets the data-derived utility value for all *unknown* properties that correspond with that question, whereas the *known* properties are set to 0. For example, the utility of *Query:Shape* when knowledge of *Shape* is unknown would be

the data-driven value (e.g., 20), and otherwise it would be 0. This approach creates a natural preference ranking over the questions.

The second, and more general, method for setting utilities is an *entropy-driven* approach based on computing the Shannon entropy for each property. For WH-questions about a property ϕ (e.g., “What color is it?”), with values ϕ_1, \dots, ϕ_n (e.g., red, green, blue, etc.), this utility is:

$$U_{WH}(\phi) = - \sum_{i=1}^n P(\phi_i) \cdot \log P(\phi_i) \quad (2)$$

For YN-questions about a property ϕ (e.g., “Is it red?”), the system will learn either that the object does or does not have a particular property value. Therefore, we computed the expected entropy for YNQs from individual entropies associated with each possible value of the property:

$$U_{YN}(\phi) = \sum_{i=1}^n P(\phi_i) \cdot (-P(\phi_i) \cdot \log P(\phi_i) - P(\neg\phi_i) \cdot \log P(\neg\phi_i)) \quad (3)$$

4.2 Dynamic Model Generation

To support generalizability, we developed a procedure to dynamically construct the decision network from a given world state in real-time. First, the world state is processed based on the agent’s perception. This includes entities, their properties and values, and a mapping of properties to questions. Next, we compute the *minimum disambiguating set* of properties of the world to determine the number of questions and structure of the decision network (see Algorithm 1). The entropy of all entity properties is then computed, and a utility table is constructed based on the entropy values in Equations 2 and 3. Finally, the decision network is instantiated with the above parameters.

Algorithm 1 Algorithm to compute the minimal disambiguating set of properties.

```

1: procedure COMPUTEMINSET(entities)
2:    $C \leftarrow list$ 
3:    $H \leftarrow list$ 
4:    $propertyList \leftarrow list$ 
5:    $entityList \leftarrow combinations(entities)$ 
6:
7:   for all  $e1, e2$  in  $entityList$  do
8:     for all  $p$  in  $e1.properties$  do
9:       if  $e1.p \neq e2.p$  then
10:         $propertyList.append(p)$ 
11:       end if
12:     end for
13:      $C.append(CNF(propertyList))$ 
14:   end for
15:
16:    $H \leftarrow SAT.solve(C)$ 
17:   return  $H$ 
18: end procedure

```

To compute the minimum disambiguating set of properties², we formulate the problem as an instance of the well known Hitting Set (Set Cover) problem and use a SAT solver to find a solution (or approximation in larger domains). The process is detailed in the algorithm in Algorithm 1. C is a collection of subsets (from a domain) that represents the properties by which all entities differ from one another. C is computed in lines 7-12 by generating all pairs of entities and comparing each of their properties. The resulting list of disambiguating properties is converted to a Boolean formula in conjunctive normal form (CNF) and added to C (line 13). Then we compute the smallest subset H by running a SAT solver on C and returning the minimum disambiguating set. This process enables an agent to disambiguate every entity in the environment based on their properties.

5 MODEL EVALUATION

To evaluate our proposed approach, we created a framework for simulating a situated dialogue task. In particular, we simulated the process of reference resolution that would arise when a situated agent would be asked to carry out a command involving an unknown entity, such as in the corpus task. **The principle aim of the simulation framework was to evaluate the efficiency of our proposed model, namely the number of questions required to correctly identify and resolve a reference to an unknown entity.**³

The task world consisted of a set of entities, each with various properties. Each entity was unique, such that no two objects possessed the same set of property values. Task world configurations can be manually specified, such as in the case of the spacecraft domain, or randomly generated. Random generation of task worlds allows us to evaluate how well our proposed system would generalize to a wider array of possible situations. In particular, to evaluate the ability of our question-asking system to ask well-targeted questions, we created two types of randomly generated worlds: (1) *low property variance* where entity property values varied only across a small ($N=3$) set of properties; and (2) *high property variance* where entity property values varied across a large ($N=7$) set of properties.

5.1 Simulated Dialogue Agent

We implemented a simulated dialogue agent that was tasked with resolving the identity of an unknown entity through dialogue. The agent could ask WHQs and YNQs regarding particular object properties to an oracle agent, which would respond based on the ground truth from the task world configuration. The simulated dialogue agent would then reduce the set of possible candidate objects based on this information. The evaluation trial would end once the dialogue agent successfully reduced the set of candidate objects to one. The number of questions required during each trial was recorded.

Two models were implemented and compared in the evaluation: (1) the model from Section 4 (with two utility assignments) and (2) a slot-filling baseline. The baseline asks questions corresponding to unknown properties of the entities. Once a property is learned

²For example, if two objects have the same shape and size, but differ in color, then the minimum set is *color*.

³Note that this is different from a typical reference resolution evaluation (e.g., [23]) in that the focus here is on dialogue efficiency, not resolution accuracy.

Table 1: Evaluation results for reference resolution efficiency of unknown entities (mean questions per instruction). Lower values indicate greater efficiency. The best results in each environment are bolded.

Environment	Human (Corpus)		Baseline		Model (Data)		Model (Entropy)	
	M	SD	M	SD	M	SD	M	SD
Spacecraft Domain	1.72	0.4	2.75	0.33	1.75	0	1.75	0
Random (low variance)	-	-	5.09	0.71	1.91	0.14	1.80	0.09
Random (high variance)	-	-	2.66	0.37	1.96	0.19	1.85	0.13

(e.g., *color:red*), it no longer asks questions about that property. To choose a question, the baseline randomly selects from the same set of questions available to the model. This stochasticity makes the baseline approach effective in the random worlds since the relevant properties are not known in advance.

We were interested in comparing how well the model performs relative to this baseline and to the humans in the corpus. The human comparison informs the feasibility of this approach as a cognitive model of human question generation under uncertainty. We predicted:

(H1): The model will outperform the baseline in the spacecraft domain.

This is because the model’s utilities are set to prioritize salient object properties that humans queried to succeed in the task. We also expected the model (but not the baseline) to achieve similar performance to humans in the corpus.

(H2): The model will outperform the baseline in the random domains, both for high and low property variance.

However, this advantage would only hold when the decision network utilities were entropy-driven rather than data-driven, since the data-driven utilities reflect a static configuration of only one environment.

(H3): Systems will perform better in domains with high property variance.

To understand the effect of property variance, we performed within-model comparisons in the random environments. Since higher variance leads to more opportunities to ask good questions, there is a higher chance that a question will target a disambiguating property.

5.2 Evaluation Results

Three systems were evaluated: baseline, entropy-driven model, and data-driven model. 20 trials (i.e., instructions) were provided to each system, each corresponding to an entity, e.g., “Pick up the temporal emitter”⁴. We ran 100 iterations of these trials in each comparison, resulting in 2000 entities to resolve for each system. For the runs in the spacecraft domain, the environment layout from the corpus was used. For the random domain runs, a new world state (and decision network) was randomly generated with each iteration according to the property variance constraints. Independent-samples t-tests were used to evaluate *H1* and *H2*, and paired-samples t-tests were used to evaluate *H3*. The dependent variable in all analyses was *mean questions per instruction* (averaged across all 100 iterations), which reflects question generation efficiency. This was used as the

⁴These trials correspond to the Commander Initiative, low-level Instruction Granularity condition in the corpus.

primary measure because resolution accuracy was always 100%, as the oracle agent had ground truth knowledge, and the correct referent could always be identified with enough questions.

Our results supported *H1* in that the model showed greater efficiency in question generation than the baseline system in the spacecraft domain (see Table 1). Specifically, the model with data-driven utilities asked fewer questions than the baseline to resolve all entities, $t(198) = 30.52, p < .001$. As predicted, there was no difference in performance of the model compared to the human data ($t(108) = 0.70, p = .48$), however, the humans outperformed the baseline system, $t(108) = 9.25, p < .001$.

We found partial support for *H2*. As predicted, the model with entropy-driven utilities outperformed the baseline in domains with both low ($t(198) = 46.03, p < .001$) and high ($t(198) = 20.40, p < .001$) property variance. However, contrary to our prediction, we found that this advantage held even when the model used data-driven utilities across both property variance conditions ($ps < .001$). This result was likely driven by the minimal disambiguating set calculation, as it served to reduce the space of questions sufficient to maintain efficiency.

Finally, we found partial support for *H3* in that the Baseline system performed significantly better with high property variance, but the Model performed significantly worse ($ps < .05$). The results for the Baseline system are intuitive in that domains with a higher degree of property variance are simpler for reference resolution since questions about most properties will be informative. However, though the Model was significantly better in the high property variance condition, the means are much closer than for the Baseline system (low: 1.91 vs high: 1.96). The difference is likely explained by the use of Algorithm 1, which reduced the number of properties, and hence mitigated the drawbacks of low property variance.

6 DISCUSSION

6.1 Contributions

6.1.1 Scientific Contributions. Our empirical results extend prior work on clarification in human-human and human-agent dialogue in the presence of multi-modal uncertainty. First, they highlight the distribution of question types that people use to resolve referential ambiguity, suggesting that WHQs, and particularly feature queries, are very common. By contrast, the results from [3] on a human learning task (using different annotation categories) found that 81% of questions were YNQs, 13% were WHQs, and 4% were AQs. This reflects the different domains and familiarity with the task. Feature-based questions were extremely common in their corpus (82% of all questions), whereas we observed about 20% of questions this

category. Like in their analysis, many of the questions we observed were grounded in properties of the entities and environment. This extends prior work on visual salience [31, 48], informing which properties are preferred in feature-heterogeneous entities. Second, the empirical results shed light on the impact of dialogue-level factors such as speaker initiative and instruction granularity on question generation. Interestingly, none of these factors impacted task performance. In terms of task duration, however, trials with high-level instructions were found to take longer on average to resolve than low-level instructions. This was likely due to the fact that high-level instructions had more uncertainty packed into them (i.e., identifying the object and the location), thus requiring more processing to resolve. As a result, high-level trials also elicited a greater proportion of questions than low-level trials, which likely reflected the extra uncertainty. In terms of speaker initiative, we found several differences in question types between the CI and RI conditions. In particular, the RI group used more *Confirm Actions* and *Goal-Related* queries than the CI group. These question types reflect the additional “burden” of taking initiative, and highlight required capabilities of robots in this domain.

6.1.2 Technical Contributions. The computational model of situated reference resolution employs an empirically-inspired decision-theoretic approach to resolving referential ambiguity through dialogue. The approach is well suited to handling uncertainty across multiple modalities, including vision, speech, knowledge, and others. In an evaluation with a simulated dialogue agent, the model outperformed a baseline system in terms of question efficiency, and showed no difference from human-level performance. The similarity to the human data suggests that the model may be capturing some of the salience-based mechanisms which humans utilize to prioritize their questions. However, one important difference is that the model used much fewer YNQs than humans, due to the reduced entropy of such questions compared to WHQs. While more work is needed, this model serves as a step toward a computational cognitive model of clarification under uncertainty.

The advantage of the model over the baseline system stems from several areas. First, the utilities enable it to prioritize more informative questions. By effectively ranking the highest entropy properties in the environment, the decision network encodes the most relevant properties to ask about. These properties include those unique to situated interaction, such as spatial orientation and landmark proximity. Second, the model is more efficient because it constrains its space of questions based on the minimum disambiguating set of properties. This, combined with the entropy-driven utilities, enables the model to scale to larger domains and to ask the most informative questions each time. Finally, dynamic construction of the network enables generalizability to open-world domains with no prior training necessary. This is a key advantage over existing approaches that require large amounts of training data and are limited to closed worlds.

Overall, this model can integrate well with a variety of systems, especially those that ground perceptual input into symbolic representations, such as [2] and [47]. These systems can benefit from the model’s focus on property-related questions, with the symbol groundings serving as input to the model. The model can also be used to enhance POMDP-based dialogue systems (e.g., [27]), which

excel at dialogue state tracking, but struggle with scalability and open-world domains.

6.2 Limitations and Future Work

One limitation of the analysis and model is the possibility that the results will not generalize to physical domains or other kinds of tasks. Though the corpus task involved a virtual robot, the interaction was situated in an environment in order to elicit uncertainty in locations, spatial reference, and procedural knowledge - all properties highly relevant for physical robots as well. We do not see any reason why the analysis and the model would not extend to physical robots, though of course such robots will need to contend with additional factors such as noisier sensor readings and potential task failure. Though the model was designed to be generalizable, future studies with physical agents in other task domains will be needed to replicate our findings and explore these additional factors.

Directions for future work are two-fold. First, we intend to further analyze the empirical dataset to better understand the questions that people asked and how these influenced task performance and efficiency. This may yield new findings of factors that influence question generation, leading to refinements of the model. The second direction for future work is to improve the computational model. We will investigate stochastic domains in which the agent has multiple sources of uncertainty and limited knowledge of the entities and their properties. The decision network is naturally designed to handle uncertainty so it would be informative to evaluate it in these more realistic domains. Finally, we will integrate this model within the language pipeline of a robot architecture and test performance in real, physical environments where world knowledge comes from noisy sensor data.

7 CONCLUSION

We presented a decision-theoretic approach to situated reference resolution for dialogue agents that was grounded in empirical findings from a corpus of clarification dialogues. The analysis revealed novel findings about the effect that dialogue-level factors have on question generation under multi-modal uncertainty. The results also informed the distribution of question types that humans use when resolving novel entities, and led to a computational model for automated question generation. The model uses a decision network and an entropy-based utility assignment method to generate efficient questions for reducing referential ambiguity. Moreover, it can be generated in real-time, making it generalizable to open-world domains. This approach can be integrated into a variety of dialogue systems to perform efficient reference disambiguation in open worlds, and is particularly well suited to multi-modal domains. Overall, these results offer insight into the ways that agents can ask questions to facilitate learning, and serve as a step toward improving automated approaches for situated reference resolution.

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, and an Air Force Office of Scientific Research grant to MS under award FA9550-18-1-0465.

REFERENCES

- [1] Jimmy Baraglia, Maya Cakmak, Yuki Nagai, Rajesh Rao, and Minoru Asada. 2016. Initiative in robot assistance during collaborative task execution. In *2016 11th ACM/IEEE international conference on human-robot interaction (HRI)*. IEEE, 67–74.
- [2] Kalesha Bullard, Andrea L Thomaz, and Sonia Chernova. 2018. Towards intelligent arbitration of diverse active learning queries. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 6049–6056.
- [3] 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, 17–24.
- [4] Sonia Chernova and Andrea L Thomaz. 2014. Robot learning from human teachers. *Synthesis Lectures on Artificial Intelligence and Machine Learning* 8, 3 (2014), 1–121.
- [5] Jennifer Chu-Carroll. 1999. Form-based reasoning for mixed-initiative dialogue management in information-query systems. In *Sixth European Conference on Speech Communication and Technology*.
- [6] Herbert H Clark. 1996. *Using language*. Cambridge university press.
- [7] Herbert H Clark and Deanna Wilkes-Gibbs. 1986. Referring as a collaborative process. *Cognition* 22, 1 (1986), 1–39.
- [8] Mark G Core, Johanna D Moore, and Claus Zinn. 2003. The role of initiative in tutorial dialogue. In *10th Conference of the European Chapter of the Association for Computational Linguistics*.
- [9] Robin Deits, Stefanie Tellex, Pratiksha Thaker, Dimitar Simeonov, Thomas Kollar, and Nicholas Roy. 2013. Clarifying commands with information-theoretic human-robot dialog. *Journal of Human-Robot Interaction* 2, 2 (2013), 58–79.
- [10] David DeVault, Natalia Karieava, Anubha Kothari, Iris Oved, and Matthew Stone. 2005. An Information-State Approach to Collaborative Reference. In *Proceedings of the ACL Interactive Poster and Demonstration Sessions*. Association for Computational Linguistics, Ann Arbor, Michigan, 1–4. <https://www.aclweb.org/anthology/P05-3001>
- [11] Benjamin Fransen, Vlad Morariu, Eric Martinson, Samuel Blisard, Matthew Marge, Scott Thomas, Alan Schultz, and Dennis Perzanowski. 2007. Using vision, acoustics, and natural language for disambiguation. In *Proceedings of the ACM/IEEE international conference on Human-robot interaction*. 73–80.
- [12] Kotaro Funakoshi, Mikio Nakano, Takenobu Tokunaga, and Ryu Iida. 2012. A unified probabilistic approach to referring expressions. In *Proceedings of the 13th annual meeting of the special interest group on discourse and dialogue*. 237–246.
- [13] Felix Gervits, Antonio Roque, Gordon Briggs, Matthias Scheutz, and Matthew Marge. 2021. How Should Agents Ask Questions For Situated Learning? An Annotated Dialogue Corpus. In *Proceedings of the The 22nd Annual Meeting of the Special Interest Group on Discourse and Dialogue*.
- [14] Peter Gorniak and Deb Roy. 2005. Probabilistic grounding of situated speech using plan recognition and reference resolution. In *Proceedings of the 7th international conference on Multimodal interfaces*. 138–143.
- [15] Arthur C Graesser, Andrew Olney, Brian C Haynes, and Patrick Chipman. 2005. AutoTutor: A Cognitive System That Simulates a Tutor Through Mixed-Initiative Dialogue. *Cognitive systems: Human cognitive models in systems design* 177 (2005).
- [16] Stephanie Gross, Brigitte Krenn, and Matthias Scheutz. 2017. The reliability of non-verbal cues for situated reference resolution and their interplay with language: implications for human robot interaction. In *Proceedings of the 19th ACM International Conference on Multimodal Interaction*. 189–196.
- [17] Sachithra Hemachandra, Matthew R Walter, and Seth J Teller. 2014. Information Theoretic Question Asking to Improve Spatial Semantic Representations.. In *AAAI Fall Symposium*.
- [18] Graeme Hirst, Susan McRoy, Peter Heeman, Philip Edmonds, and Diane Horton. 1994. Repairing conversational misunderstandings and non-understandings. *Speech communication* 15, 3-4 (1994), 213–229.
- [19] Eric Horvitz and Tim Paek. 1999. A computational architecture for conversation. In *UM99 User Modeling*. Springer, 201–210.
- [20] Xin Huang, Chor Seng Tan, Yan Bin Ng, Wei Shi, Kheng Hui Yeo, Ridong Jiang, and Jung-jae Kim. 2021. Joint Generation and Bi-Encoder for Situated Interactive MultiModal Conversations. In *AAAI 2021 DSTC9 Workshop*.
- [21] Joris Hulstijn. 2003. Roles in dialogue. In *Proceedings of the 7th Workshop on Formal Semantics and Pragmatics of Dialogue (Diabrick)*. Citeseer.
- [22] 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*. 422–431.
- [23] Casey Kennington and David Schlangen. 2017. A simple generative model of incremental reference resolution for situated dialogue. *Computer Speech & Language* 41 (2017), 43–67.
- [24] 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*. 210–214.
- [25] Kazunori Komatani, Tsugumi Otsuka, Satoshi Sato, and Mikio Nakano. 2017. Question selection based on expected utility to acquire information through dialogue. In *Dialogues With Social Robots*. Springer, 79–91.
- [26] Geert-Jan M Kruijff, Hendrik Zender, Patric Jensfelt, and Henrik I Christensen. 2007. Situated dialogue and spatial organization: What, where... and why? *International Journal of Advanced Robotic Systems* 4, 1 (2007), 16.
- [27] Pierre Lison. 2013. Model-based bayesian reinforcement learning for dialogue management. *arXiv preprint arXiv:1304.1819* (2013).
- [28] Changsong Liu, Rui Fang, Lanbo She, and Joyce Chai. 2013. Modeling collaborative referring for situated referential grounding. In *Proceedings of the SIGDIAL 2013 Conference*. 78–86.
- [29] 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)*.
- [30] Nikolaos Mavridis. 2015. A review of verbal and non-verbal human-robot interactive communication. *Robotics and Autonomous Systems* 63 (2015), 22–35.
- [31] Teruhisa Misu. 2015. Visual saliency and crowdsourcing-based priors for an in-car situated dialog system. In *Proceedings of the 2015 ACM on International Conference on Multimodal Interaction*. 75–82.
- [32] Teruhisa Misu, Komei Sugiura, Kiyonori Ohtake, Chiori Hori, Hideki Kashioka, Hisashi Kawai, and Satoshi Nakamura. 2010. Modeling spoken decision making dialogue and optimization of its dialogue strategy. In *Proceedings of the SIGDIAL 2010 Conference*. 221–224.
- [33] Seungwhan Moon, Satwik Kottur, Paul A Crook, Ankita De, Shivani Poddar, Theodore Levin, David Whitney, Daniel Difranco, Ahmad Beirami, Eunjoon Cho, et al. 2020. Situated and interactive multimodal conversations. *arXiv preprint arXiv:2006.01460* (2020).
- [34] Thibaut Munzer, Yoan Mollard, and Manuel Lopes. 2017. Impact of robot initiative on human-robot collaboration. In *Proceedings of the Companion of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*. 217–218.
- [35] Daniel L Owen. 1978. The use of influence diagrams in structuring complex decision problems. *Readings on the principles and applications of decision analysis* 2 (1978), 763–771.
- [36] Tim Paek. 2003. Toward a taxonomy of communication errors. In *ISCA Tutorial and Research Workshop on Error Handling in Spoken Dialogue Systems*.
- [37] Tim Paek and Eric Horvitz. 2000. Conversation as Action under Uncertainty. In *Proceedings of the Sixteenth Conference on Uncertainty in Artificial Intelligence (Stanford, California) (UAI'00)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 455–464.
- [38] P. Prodanov and A. Drygajlo. 2005. Decision Networks for Repair Strategies in Speech-Based Interaction with Mobile Tour-Guide Robots. In *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*. 3041–3046.
- [39] David Schlangen, Sina Zarrieß, and Casey Kennington. 2015. Resolving references to objects in photographs using the words-as-classifiers model. *arXiv preprint arXiv:1510.02125* (2015).
- [40] Gabriel Skantze. 2005. Exploring human error recovery strategies: Implications for spoken dialogue systems. *Speech Communication* 45, 3 (2005), 325–341.
- [41] Jesse Thomason, Aishwarya Padmakumar, Jivko Sinapov, Nick Walker, Yuqian Jiang, Harel Yedidion, Justin Hart, Peter Stone, and Raymond Mooney. 2020. Jointly improving parsing and perception for natural language commands through human-robot dialog. *Journal of Artificial Intelligence Research* 67 (2020), 327–374.
- [42] David Traum, David DeVault, Jina Lee, Zhiyang Wang, and Stacy Marsella. 2012. Incremental dialogue understanding and feedback for multiparty, multimodal conversation. In *International Conference on Intelligent Virtual Agents*. Springer, 275–288.
- [43] Marilyn Walker and Steve Whittaker. 1990. Mixed Initiative in Dialogue: An Investigation into Discourse Segmentation. In *28th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Pittsburgh, Pennsylvania, USA, 70–78.
- [44] Jason D Williams and Steve Young. 2007. Partially observable Markov decision processes for spoken dialog systems. *Computer Speech & Language* 21, 2 (2007), 393–422.
- [45] Tom Williams, Saurav Acharya, Stephanie Schreitter, and Matthias Scheutz. 2016. Situated open world reference resolution for human-robot dialogue. In *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 311–318.
- [46] Tom Williams, Gordon Briggs, Bradley Oosterveld, and Matthias Scheutz. 2015. Going Beyond Literal Command-Based Instructions: Extending Robotic Natural Language Interaction Capabilities.. In *AAAI*. 1387–1393.
- [47] Tom Williams, Fereshta Yazdani, Prasanth Suresh, Matthias Scheutz, and Michael Beetz. 2019. Dempster-shafer theoretic resolution of referential ambiguity. *Autonomous Robots* 43, 2 (2019), 389–414.
- [48] Steven Yantis. 2005. How visual salience wins the battle for awareness. *Nature neuroscience* 8, 8 (2005), 975–977.
- [49] Hendrik Zender, Geert-Jan M Kruijff, and Ivana Kruijff-Korbayová. 2009. Situated Resolution and Generation of Spatial Referring Expressions for Robotic Assistants.. In *IJCAI*. 1604–1609.