

Asking the Right Question at the Right Time: Human and Model Uncertainty Guidance to Ask Clarification Questions

Alberto Testoni and Raquel Fernández
Institute for Logic, Language and Computation
University of Amsterdam
{a.testoni, raquel.fernandez}@uva.nl

Abstract

Clarification questions are an essential dialogue tool to signal misunderstanding, ambiguities, and under-specification in language use. While humans are able to resolve uncertainty by asking questions since childhood, modern dialogue systems struggle to generate effective questions. To make progress in this direction, in this work we take a collaborative dialogue task as a testbed and study how model uncertainty relates to human uncertainty—an as yet under-explored problem. We show that model uncertainty does not mirror human clarification-seeking behavior, which suggests that using human clarification questions as supervision for deciding when to ask may not be the most effective way to resolve model uncertainty. To address this issue, we propose an approach to generating clarification questions based on model uncertainty estimation, compare it to several alternatives, and show that it leads to significant improvements in terms of task success. Our findings highlight the importance of equipping dialogue systems with the ability to assess their own uncertainty and exploit in interaction.

1 Introduction

The ability to ask effective and informative questions is a peculiar feature of human intelligence. From an evolutionary perspective, it has been argued that being able to ask questions to reliable informants gives humans a key evolutionary advantage over other species (Tomasello, 2009). Questions are also shown to play a fundamental role in cognitive development and language learning during childhood (Demetras et al., 1986; Chouinard et al., 2007; Leech et al., 2013; Ruggeri et al., 2016, 2017). Both children and adults generally ask questions to resolve a state of uncertainty (Piaget, 1954; Rothe et al., 2018). Deciding how to resolve uncertainties and maintaining mutual understanding are essential steps to achieve *conversational grounding* (Clark and Schaefer, 1987; Clark and Brennan,

1991), the process through which humans negotiate the production and acceptance of their utterances. Clarification questions are among the key strategies speakers use to achieve conversational grounding. However, while humans handle and resolve uncertainties through clarification questions almost effortlessly in everyday conversations, this is particularly challenging for modern dialogue systems based on deep neural networks (Chandu et al., 2021). Even the new generation of conversational agents, such as ChatGPT,¹ seem to struggle to generate clarification questions when dealing with ambiguous instructions (Deng et al., 2023).

In this paper, we investigate uncertainty-driven strategies for asking clarification questions. Clarification requests are usually triggered when ambiguity and/or underspecification make it difficult for the receiver to follow up with enough confidence (Purver, 2004; Benotti and Blackburn, 2017; Pezzelle and Fernández, 2023). It should come as no surprise then that there is a long-standing tradition of work highlighting the importance of incorporating clarification requests in dialogue systems (Paek and Horvitz, 2000; Gabsdil, 2003; Schlangen, 2004; Rieser and Moore, 2005; Rieser and Lemon, 2006, among others). Recent years have witnessed a renewed interest in this ambitious goal, especially in the context of collaborative dialogue tasks (e.g., Nguyen and Daumé III, 2019; Chi et al., 2020; Zhu et al., 2021; Shi et al., 2022; Shen and Lourentzou, 2023; Madureira and Schlangen, 2023b). Yet, several key aspects remain under-explored.

For example, some existing approaches use a dedicated neural network model trained on human clarification behaviour to decide *when* to ask (Shi et al., 2022; Madureira and Schlangen, 2023b). However, it is well known that humans may follow different confirmation-seeking strategies (Schegloff et al., 1977; Schegloff, 1987; Purver et al.,

¹<https://chat.openai.com/>

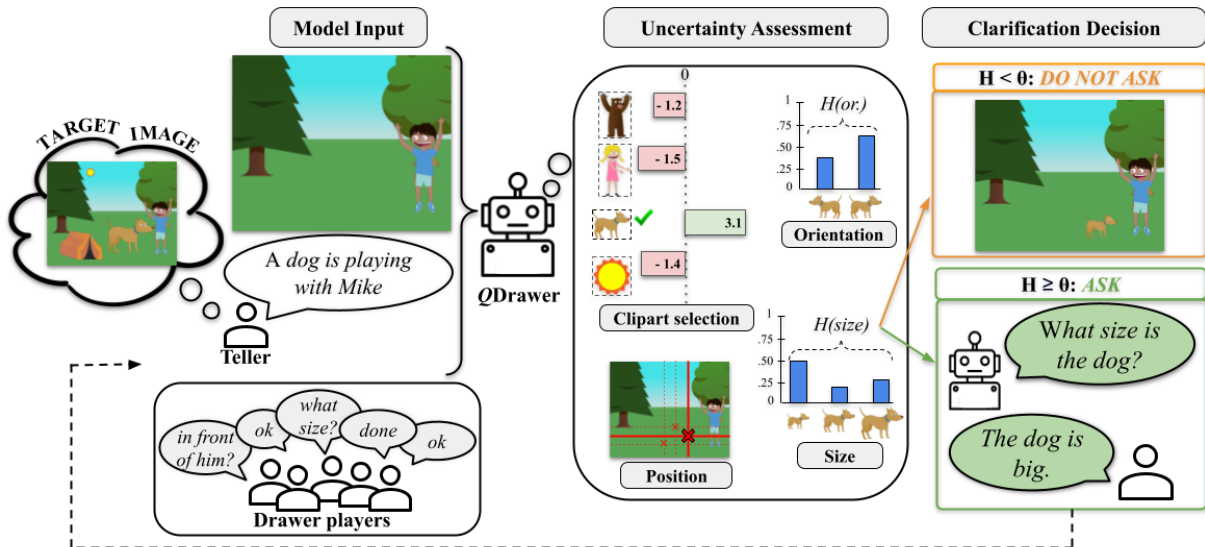


Figure 1: Overview of our experimental setup in CoDraw. After receiving an instruction from the Teller, the Drawer agent selects the clipart(s) to draw, together with their attributes. If the entropy over an attribute exceeds a threshold θ (*size* in the figure), then the Drawer asks a clarification question. The question-answer pair is added to the dialogue history before performing the next drawing action. Different human players may react differently (bottom-left box); the agent decides whether to ask for clarification on the basis of its own uncertainty, independently from the clarification decisions of human players.

2001). This potential human disagreement casts doubt on the effectiveness of using human decisions as training signal for supervised or reinforcement learning. As an alternative approach, model uncertainty has often been employed to decide *when* and/or *what* to ask (Nguyen and Daumé III, 2019; Zhu et al., 2021; Shen and Lourentzou, 2023; Chi et al., 2020). Yet the two types of approaches have not been directly compared, nor is there any study (to the best of our knowledge) on the extent to which model uncertainty correlates with human uncertainty. In our work, we take a closer look at these fundamental aspects.

We consider the CoDraw task and dataset (Kim et al., 2019)—where a Drawer player is asked to recreate a clipart scene following instructions from a Teller player—as a testbed to study uncertainty-driven clarification strategies in collaborative agents. We propose different ways to represent model uncertainty in CoDraw, and then carry out a study on the relation between model uncertainty and human uncertainty, as measured by the decision to ask clarification requests in human-human dialogues. This study reveals that there is a poor relation between model and human uncertainty, which in part appears to stem from our finding that in CoDraw human players do not follow a consistent clarification strategy. This poor relation suggests that using human clarification decisions

as supervision for deciding when to ask may not be the most effective way to resolve model uncertainty. We therefore propose a new Drawer agent model equipped with an uncertainty-based clarification module (illustrated in Figure 1), where we use a template-based approach to generate clarification questions triggered by the uncertainty state of the agent. This approach allows us to modulate the level of uncertainty that leads to posing a question and thus control for the number of questions asked per dialogue, which may be advantageous given the processing cost of questions for humans as well as models (Chiyah-Garcia et al., 2023).

Our results and analyses show that our proposed model is significantly more effective in terms of task success than several strong baselines, including a version based on the human supervision model by Madureira and Schlangen (2023b). Overall, our findings highlight the importance of equipping conversational agents with the ability to assess and act upon their own uncertainty by asking effective clarification questions.

2 Related Work

Clarification requests have long been studied as a key repair mechanism in human dialogue (Gabsdil, 2003; Purver, 2004; Schlangen, 2004; Rieser and Lemon, 2006; Schlöder and Fernández, 2014;

Benotti and Blackburn, 2017, 2021). Datasets of human conversations where spontaneous clarification questions are annotated have been an important resource for data-driven research on clarification strategies. For example, Rodríguez and Schlangen (2004) proposed an annotation scheme based on Clark’s four levels of communication (Clark, 1996) to identify the form and function of clarification requests. These and similar annotation schemes have been applied to small datasets (Benotti, 2009; Gervits et al., 2021; Benotti and Blackburn, 2021).

At a larger scale, the TEACH dataset (Padmakumar et al., 2022) has been annotated by Gella et al. (2022) with dialogue acts, including acts related to clarification, and the Minecraft Dialogue Corpus (Narayan-Chen et al., 2019) and its extensions (Shi et al., 2022; Kiseleva et al., 2022; Mohanty et al., 2022) have been annotated with different types of clarification requests. However, in the latter case the dataset was augmented with further clarification questions *a-posteriori*, introduced asynchronously by other annotators who were asked to decide whether clarifications were needed. Aliannejadi et al. (2021) and Gao et al. (2022) adopt similar techniques for including clarification questions in dialogue data.

Recently, Madureira and Schlangen (2023b,a) released CoDraw-iCR (v1 and v2), a fine-grained annotation of spontaneous clarification requests in the CoDraw dataset (Kim et al., 2019). This constitutes a worthwhile resource for analysing clarification behaviour in collaborative tasks. Madureira and Schlangen (2023b) exploit these annotations to train a neural model for deciding *when* to clarify. In contrast, we shall use them to study the relation between human and model uncertainty in Section 5 and then, in Section 6, propose a model for the CoDraw Drawer agent that asks clarification questions on the basis of its own uncertainty.

Model uncertainty represents a valuable signal for Natural Language Generation models to deal with the different sources of ambiguity in language use (Baan et al., 2023). Earlier on, when the speech recogniser used to be a common source of errors in spoken dialogue systems, ASR confidence scores were exploited to decide on error handling strategies such as confirmation and clarification (Skantze, 2008; Stoyanchev et al., 2014, among others). In more recent systems, where the focus is on the next action to take in a collaborative task, a common approach consists of using model uncertainty (as measured by the entropy over the model pre-

diction on a set of actions) to signal ambiguity and trigger the generation of clarification questions (Nguyen and Daumé III, 2019; Zhu et al., 2021; Shen and Lourentzou, 2023). Without explicitly estimating uncertainty, Testoni and Bernardi (2021) propose a beam search re-ranking algorithm guided by the model intermediate predictions about the target in a referential task. Similar approaches use the probability difference between the top two predicted actions to estimate uncertainty (Chi et al., 2020). Recently, Naszadi et al. (2023) investigated the alignment between predictive uncertainty and ambiguous instructions in visually-grounded communication tasks. They show that well-calibrated prediction probabilities benefit the detection of ambiguous instructions. Differently from our work, the authors do not investigate the effectiveness of model uncertainty to generate follow-up clarification questions.

Instead of relying on model uncertainty, other proposed methods involve training external neural networks or sub-modules with the objective of deciding *when* to ask for clarification (Madureira and Schlangen, 2023b; Shi et al., 2022). Khalid et al. (2020) propose to incorporate cognitive modeling in a Reinforcement Learning framework to generate context-sensitive clarification strategies in a referential game. While a few approaches generate clarification questions from scratch in an end-to-end fashion (Khalid et al., 2020; Zhu et al., 2021), heuristics are often applied at the generation stage (Shen and Lourentzou, 2023; Zhang and Zhu, 2021; Sekulić et al., 2021). In this paper, we propose an approach to generating clarification questions based on model uncertainty estimation and compare it to alternative methods for deciding when to clarify by training on human clarification decisions.

3 Task & Dataset

CoDraw (Kim et al., 2019) is a goal-driven instruction-following collaborative task between two players. The Teller player has access to an abstract target scene containing a variable number of cliparts (on average, target scenes contain 6 cliparts), leveraging the Abstract Scene dataset (Zitnick and Parikh, 2013; Zitnick et al., 2013). The Teller provides written instructions in English to the Drawer in order to reconstruct the scene in a turn-based fashion. The Drawer has access to a gallery of 58 cliparts that they can place on

a canvas. The Drawer can remove, resize (small, medium, large), or flip (facing right or left) each object in the scene. After receiving an instruction, the Drawer has to reply to the Teller before proceeding to the following turn. Typically, the Drawer either replies with confirmation feedback or asks follow-up questions. The CoDraw dataset consists of around 10k dialogues, with an average length of 7.7 turns per dialogue. Additional details can be found in Appendix D.

In CoDraw, task success is measured by comparing the target image with the one drawn by the Drawer at the end of the conversation. To this aim, Kim et al. (2019) designed a similarity score metric, ranging from 0 to 5, tailored to the CoDraw task. This score takes into account different cliparts (by checking whether cliparts in the target scene appear in the reconstructed one) and their attributes (size, position, flipping, plus some additional attributes for cliparts representing people, i.e., facial expression and body pose) by assigning different weights to each component. The scenes drawn by human players reach a similarity score of 4.17 (sd=0.64). In our work, we apply some changes to the similarity score metric to fix some flaws in its design and better capture the properties under analysis in the experimental setup, as described in Appendix A. In the following, we report the results using the new version of the metric.

4 Computing Uncertainty in CoDraw

4.1 The Silent Drawer Model

Kim et al. (2019) proposed a neural Drawer model for the CoDraw task. At each round of the dialogue, this model is conditioned on the Teller’s latest message, which is encoded into a vector using a bidirectional LSTM module. The Drawer also receives as input a vector representation of the current state of the canvas that is being drawn (the cliparts added so far and their attributes). These input representations are passed to a dense feed-forward neural network. The resulting vector represents the Drawer’s follow-up action, which determines the cliparts to be added to the canvas, together with their attributes. Thus, Kim et al.’s Drawer model is *silent*, i.e., it can only draw objects on the canvas but cannot contribute to the dialogue. The model is trained in a Supervised Learning fashion on human data. At inference time, the model assigns an unbounded score to each of the 58 available cliparts: each clipart with a score above 0 is drawn on the

canvas, together with its predicted attributes (size, position, orientation, etc.). The model is tested by providing turn-by-turn Teller’s instructions from the CoDraw test set, reaching an average similarity score of 3.31 (sd=0.67).

4.2 Computing Model Uncertainty

We are interested in representing the uncertainty of the model regarding its next action. After receiving an instruction, the Drawer model has to make decisions about which clipart(s) to draw and their corresponding attributes. For each dialogue turn and selected cliparts, we estimate model uncertainty regarding these components as follows:

- For **clipart selection**, we take the score assigned to each clipart as a proxy of model uncertainty. The lower the value, the more uncertain the model is about selecting the given clipart.
- For the **size** (*small, medium, large*) and **orientation** (*right, left*) attributes, we compute the entropy of the probability distribution over the possible labels per attribute. The higher the entropy, the higher the uncertainty on that attribute.
- For the **position** attribute, the Drawer outputs two values per clipart, corresponding to the x - and y -axis coordinates of the canvas. Measuring uncertainty from continuous values is inherently more challenging than discrete variables. We take inspiration from uncertainty estimation using Ensembles of different models (Lakshminarayanan et al., 2017) and consider the variance among the prediction of five Drawer models trained with different random initialization seeds as a representation of model uncertainty.

5 Model vs. Human Uncertainty

In this section, we investigate the relation between model and human uncertainty. In the absence of clues such as eye gaze or reaction times (which could be indicative of uncertainty, but are not available in CoDraw), we take clarification questions as a signal of human uncertainty. We exploit the annotation of clarification questions released by Madureira and Schlangen (2023b,a), which consists of turns by the Drawer player annotated as clarification requests plus the attributes mentioned in such requests (size, position, orientation, etc.).² Around 40% of the dialogues in CoDraw contain

²In particular, we take advantage of the annotation in CoDraw-iCR (v2) by Madureira and Schlangen (2023a).

at least one clarification question, with an average number of 2.2 clarification requests per dialogue in this subset.

5.1 Setup

We explore whether we can use model uncertainty (as defined in the previous section) to predict when human players ask clarification questions. To this aim, we train a logistic regression model with human decisions about whether to ask a clarification question in the follow-up turn as *dependent variable* and each model uncertainty type (clipart selection, size, orientation, position) as *independent variables*. We use the Drawer model described in Section 4.1 to compute model uncertainty at each dialogue turn. Using the methods described in Section 4.2, we compute uncertainty for each clipart with a score above 0 and each of its attributes. If this includes more than one clipart at a given turn, we take the highest uncertainty value per variable (clipart selection, size, orientation, position).³ We use the CoDraw test set (1002 dialogues, with an average of 7.7 turns per dialogue) to extract model uncertainty. We further split the CoDraw test set to train and evaluate the logistic regression model (70% and 30%, respectively). We study the effect of different initialization seeds and training epochs and compare the model performance to a random baseline.

5.2 Results

The results are displayed in Figure 2. Let us consider the solid and dotted green lines in the plot. The logistic regression model trained to predict *when* human players ask clarification questions based on the uncertainty of the Drawer model yields an Average Precision of 0.188 (epoch 15; average across 5 different initialization seeds). This result is above the random baseline (0.188 vs. 0.142, $p < 0.05$), indicating that model uncertainty has some degree of predictive power. Nevertheless, performance is poor. Among the independent variables analyzed, only the uncertainty about clipart selection has statistically significant predictive power ($p < 0.001$). Additional details about the coefficients can be found in Appendix B.

Figure 2 provides additional interesting insights: we can see that the Drawer model learns to perform the task over the training epochs, as evidenced by

³Other approaches, such as taking the average of the uncertainty values, do not lead to any significant difference in the final results.




Figure 2: Average Precision (AP, green) of a logistic regression model in predicting when human players asked clarification questions based on the uncertainty extracted from a Drawer trained for a variable number of epochs (x -axis). The plot also shows the similarity score (SS, blue) achieved by the Drawer, and its uncertainty on clipart selection (normalized and reversed clipart scores), and on the size and orientation attributes (as measured by the entropy H). For each epoch, we report the average and variance of 5 Drawer models trained with different seeds. Note that the y -axis has two scales: we report SS on the left, while all the other metrics refer to the right side.

the upward trend of the similarity score, which measures task success. At the same time, the uncertainty of the model decreases (the plot shows uncertainty on clipart selection, size, and orientation. Same pattern is observed for position). Thus, the Drawer model is reasonably well calibrated: higher task performance goes hand in hand with lower uncertainty. Yet, this trend does not have an impact on the power of model uncertainty to predict human clarification behaviour, as the Average Precision of the logistic regression model remains constantly low over training epochs. We verified that other methods to capture model uncertainty, such as extracting epistemic and aleatoric uncertainty (Kendall and Gal, 2017) using an Ensemble of models, lead to the same results and do not have significant predictive power. Additional details can be found in Appendix B.

5.3 Analysis

There could be many reasons why model uncertainty and human uncertainty are poorly related. Human players do not only have immensely more knowledge and resources than the relatively simple Drawer model; different speakers may also follow different conversation strategies and such variation

may be difficult to predict. It is well known that there is variability in human clarification-seeking and repair strategies (Schegloff et al., 1977; Schegloff, 1987; Purver et al., 2001). To check whether such variation is present in the CoDraw dataset, we carry out an exploratory analysis. We use a keyword-matching approach to cluster together instructions by Teller players with similar form and content that lack some key information, and check how frequently they are followed by clarification questions. We only consider Teller’s instructions uttered in the first dialogue turn to exclude the possible confounding effect of the previous dialogue history. Full details on the procedure and the results can be found in Appendix B. Here we highlight the main takeaways. For example, we observe that there are a total of 2274 Teller first instructions that mention at least one clipart but do not specify its size (e.g., ‘a tree to the left a few inches from the edge’). Only 23% of follow-up turns by the Drawer player include a clarification question, and among these, only 27% are about the size of the clipart. Similarly, there are 240 Teller first instructions that mention at least one clipart but do not specify its location (e.g., ‘little girl running towards her friend that has a pie in his hand’). 49% of follow-up turns include a clarification question, and among these, 48% are about the location of the clipart.

Thus we observe that when analyzing utterances that do *not* mention a specific necessary attribute, the human Drawer players may or may not ask about the missing information. Humans may be more or less cautious, take more or fewer risks, or guess/hope that more information will be provided later on in the dialogue (Purver, 2004; Benotti and Blackburn, 2017, 2021). This high human disagreement on deciding whether to clarify or not suggests that learning when to ask clarification questions from human clarification behaviour may not be optimal. Indeed the model proposed by Madureira and Schlangen (2023b) achieves only 0.35 Average Precision. Instead, in the next section we investigate whether model uncertainty represents an effective guidance to decide *when* to ask clarification questions by studying their impact on task success.

6 Asking Clarification Questions

We implement a *Questioning Drawer* model (*QDrawer*) by adapting the *Silent Drawer* described in Section 4.1. We focus on the size attribute, because it applies to all cliparts and it is often un-




Figure 3: Example of clarification exchanges (in green) with questions generated by *QDrawer* and answers by the Teller. The images on the right depict the *QDrawer*’s canvas after each round of conversation.

derspecified in the Teller’s instructions (see Section 5). Our goal is to test whether asking clarification questions triggered by size uncertainty results in (1) higher accuracy regarding the size attribute, and (2) higher overall similarity score. To isolate the effect of asking questions and avoid confounds related to language generation, we use a template-based approach to generate clarification questions and their answers.

6.1 The Uncertainty-based *QDrawer*

We adapt the *Silent Drawer* as follows: (a) the model receives as input the immediately preceding turn exchange between the Drawer and Teller, instead of just the Teller’s utterance; (b) we add a term in the cross-entropy loss to include not only the cliparts that are added to the canvas but also the ones that are updated by changing some attributes;⁴ (c) we include a clarification generation module that works as follows: At any turn of the

⁴In this way, the model learns to assign a high score to all cliparts mentioned in the instruction and not just to the ones that do not appear yet on the canvas—an essential step to successfully profit from clarification exchanges. See Appendix E for additional details.

conversation, if the entropy over the size attribute labels (*small*, *medium*, or *large*) exceeds a given threshold,⁵ *QDrawer* generates a clarification question targeting a maximum number of 2 cliparts (i.e., the ones with the highest entropy).

We employ the following template, which corresponds to the most common formulation of size questions in CoDraw according to [Madureira and Schlangen \(2023a\)](#): “*what size is the ___?*” or “*what size are the ___ and the ___?*”, where the slots are filled with the name(s) of the clipart(s). We exploit the ground-truth annotation in the CoDraw dataset to provide an answer to the question on behalf of the Teller using these templates: “*The ___ is ___*” or “*The ___ is ___ and the ___ is ___*”, where the slots are filled with the name(s) of the clipart(s) under discussion and their size (small, medium, large). Before performing a new drawing action, *QDrawer* is conditioned on the current state of the canvas, as well as the clarification question asked and the corresponding answer. Figure 3 shows an example of clarification questions asked by *QDrawer* and the answers received, as well as the updated canvas after each round of the conversation. Another example is reported in Appendix C.

We experimented with replacing the Bi-LSTM language encoder of the Drawer agent (see Section 4.1) with more complex models (both fine-tuned on CoDraw and trained from scratch), such as BERT and RoBERTa ([Devlin et al., 2019](#); [Liu et al., 2019](#)) and their corresponding distilled versions ([Sanh et al., 2019](#)). We found that these BERT-based *QDrawer* variants give rise to the same patterns on the role of clarification discussed in the results below, but lead to a deterioration of model performance as measured by the similarity score (max. 3.07).⁶ We thus decided to keep the original model architecture.⁷

6.2 Baselines

We compare *QDrawer* to the *Silent Drawer* model, as well as to different configurations of *QDrawer* that differ with respect to the information used to

⁵We experiment with several thresholds, as it will become clear in subsequent sections.

⁶This result is in line with a strand of work highlighting that LSTM-based models outperform BERT-like ones for small datasets and/or specific domains ([Zeyer et al., 2019](#); [Ezen-Can, 2020](#); [Le et al., 2021](#); [Wei et al., 2022](#)).

⁷Again, a logistic regression model trained with the uncertain extracted from the Drawer variants based on BERT and RoBERTa poorly correlates with human behaviour on asking clarification questions, reaching an average precision (0.18 AP) in line or slightly below the one discussed in Section 5.

	When To Ask	Size Acc.	SS	CQs Size Acc. boost	CQs SS boost
SD	–	76.9	3.31	–	–
<i>QD</i>	Human	77.3	3.31 [◦]	+6.2%	+0.04 [•]
	Decider	80.1	3.34	+7.1%	+0.06
<i>QD</i>	$\theta = 0.3$	87.3	3.40	+11.0%	+0.09
	$\theta = 0.7$	86.0	3.39	+11.0%	+0.09
	$\theta = 1.1$	82.4	3.36	+9.8%	+0.09

Table 1: Silent (*SD*) and Questioning Drawer (*QD*) results using different approaches to generate clarification questions. Top part: baselines. *CQs*: clarification questions. *Human*: dialogue turns when human players ask CQs. *SS*: similarity score on the whole test set. *Size acc.*: accuracy (%) of the size attribute of all drawn cliparts on the whole test set. *Boost* columns: SS and Size Acc. boost brought by CQs on the *subset* of dialogues containing at least one CQ. All the differences are statistically significant except [◦] (non-significant difference with *SD*) and [•] (adding CQs does not significantly increase SS on the subset).

decide *when* to ask a question. In particular, we compare our proposed uncertainty-based approach to the following alternatives:

- The supervised learning approach by [Madureira and Schlangen \(2023b\)](#), where an external neural model trained on human clarification behaviour is used to decide when to ask (*Decider*). In this case, size uncertainty is still used to select which clipart(s) to target in the question, and the clarification question is generated and answered in the same way as in the uncertainty-based *QDrawer*.
- Asking questions at the same position (in terms of dialogue turn) where *human* players ask questions about size in CoDraw.
- Asking questions at *random* dialogue turns (discussed in Section 6.5).

6.3 Task Success Results

Table 1 summarizes the results of our experiments on the CoDraw test set. We compare different models on their accuracy in correctly selecting the size of the cliparts drawn on the canvas and the overall similarity score between the drawn and the target scene at the end of the dialogue. The *Silent Drawer* (*SD*) reaches a size accuracy of 76.9% and a similarity score of 3.31. The version of *QDrawer* that generates size questions in the same position as human players only marginally improves performance on size accuracy, while the similarity score does not show any statistically significant increase (paired t-test, $p > 0.05$). Using an external neural

Decider to choose *when* to ask clarification questions improves the performance over *SD* (up to 80.1% accuracy and 3.34 similarity score). Finally, the uncertainty-driven *QDrawer* outperforms the baselines to a large extent: up to 87.3% size accuracy and 3.40 similarity score.⁸ Regarding the comparison with the *Decider* model, it is worth noting that this model is trained to ask questions at each position where human players asked questions, not only when humans asked size questions. While this method is somewhat coarse-grained (because the *Decider* may ask size-CQ when they are not needed), it ensures full recall (the *Decider* will ask a size question whenever humans asked a size question). Hence, the *Decider* baseline is still meaningful with respect to the size attribute.

The *QDrawer* model shows a consistent improvement as the entropy threshold is lowered. This aspect is investigated in more detail in Section 6.5.

6.4 Clarification Questions’ Contribution

To further analyze how clarification exchanges contribute to task success, for each model configuration, we consider the subset of dialogues *CQ* where at least one clarification question is asked. We then remove all clarification questions from these dialogues, resulting in *!CQ*. We compute the difference in size accuracy and similarity score between *CQ* and *!CQ* to check the specific contribution of clarification questions (accuracy and SS boost). As we can see on Table 1, the boost brought about by clarification questions is much higher for the uncertainty-driven *QDrawer* than for the baseline models. In particular, for the subset of dialogues in *CQ*, clarification exchanges increase size accuracy by +11% in absolute terms and the similarity score by +0.09. Asking questions in the same position as human players does not significantly increase the similarity score when comparing *CQ* and *!CQ* (paired t-test, $p > 0.05$). To make sure that this statically non-significant result does not stem from our generation approach, we replace machine-generated clarification questions with human ones for the dialogues in *CQ*. We find that in this case the increase in both size accuracy and similarity score is even less pronounced than what we observe in Table 1, resulting in non-significant results for both question-generation methods.

⁸An exploratory study on extending our approach to other attributes reveals that using both *size* and *orientation* uncertainty further increases the similarity score on the test set up to 3.46, indicating that our proposed method can scale up.

θ	% dialogues with at least 1 CQ	avg. number of CQ per dialogue
0.3	96.6	3.63
0.7	84.7	2.69
1.1	57.1	1.84

Table 2: The effect of different θ values on the number of questions asked by *QDrawer*. We report the percentage of dialogues with at least one clarification question and, within this subset, the average number of clarification questions per dialogue.




Figure 4: Effect of the average number of questions per dialogue (considering all dialogues in the test set) on the *size* accuracy. We compare the uncertainty-guided *QDrawer* with a version that asks questions in random turns and one that asks questions in the turns selected by an external *Decider* model.

We further investigate the impact of clarification exchanges by analyzing the calibration of the model uncertainty estimates. We compare the probability distribution over the size attribute against the ground-truth size of each object by using two metrics: the Expected Calibration Error (ECE, [Naeni et al., 2015](#); [Guo et al., 2017](#)) and the Brier score ([Brier, 1950](#)). For both metrics, a lower score indicates better calibration.⁹ We compare the *Silent Drawer* and *QDrawer* ($\theta = 0.3$) and find that the *Silent Drawer* reaches an ECE value of 0.10 while the ECE for *QDrawer* goes down to 0.05 (the difference is significant: t-test, $p < 0.001$). Similarly, the Brier score significantly improves (t-test, $p < 0.001$) from 0.36 (*Silent Drawer*) to 0.28 (*QDrawer*). These results indicate that asking clarification questions improves the calibration of the model uncertainty estimates.

⁹We refer to [Ovadia et al. \(2019\)](#) for a broader discussion on calibration metrics.

6.5 How Many Questions to Ask

Asking clarification questions carries a cost, both from the perspective of the agent asking the question and the agent processing it (Clark, 1996; Purver, 2004). We take this partially into account by adopting a simple approach: we evaluate *QDrawer* by using different entropy threshold values θ to control for the *number* of questions generated by the model. As reported in Table 2, in *QDrawer* the value of the entropy threshold θ has, by design, a direct impact on the number of questions asked: i.e., a model that only asks for clarification in the face of high uncertainty ($\theta = 1.1$) gives rise to fewer dialogues with questions and to a lower average number of questions per dialogue.

We compare the uncertainty-driven *QDrawer* against (1) a version that asks questions in random dialogue turns with different average number of questions per dialogue and (2) the *QDrawer* paired with the *Decider* module as described in Section 6.2. For the latter, the maximum possible number of questions per dialogue (average of 1.5) is defined by the *Decider* model output. Figure 4 shows the effect of the number of questions on size accuracy for different model configurations. While the performance of the *Decider* model almost overlaps with the *QDrawer* asking questions in random dialogue turns, using model uncertainty significantly improves accuracy, even when a few questions are asked. This result confirms that even when controlling for the number of questions asked (which may be desirable in terms of cost control), the uncertainty-driven *QDrawer* shows the most promising results.

7 Conclusions

Different kinds of uncertainties characterize language use, and asking informative and effective questions is a crucial tool to address these uncertainties. While humans are able to ask strategic questions since early childhood to resolve a state of uncertainty, equipping modern dialogue systems with the ability to ask effective questions has proved challenging. Some existing approaches propose employing dedicated neural network models trained on human behaviour to decide when to ask for clarification in a dialogue. However, the high variability in human clarification strategies casts doubt on the effectiveness of this approach. Other proposed methods involve using model uncertainty to decide *when* and/or *what* to ask. Yet, estimating

and exploiting model uncertainty is not trivial, and the relationship between model and human uncertainty is still unclear.

In our work, we considered the CoDraw collaborative dialogue task as a test bed. We proposed different ways of extracting model uncertainty and investigated its relationship to human uncertainty, as measured by the clarification decisions of human players. We found that human and model uncertainty are poorly correlated, possibly due to the wide diversity of human clarification strategies. We then presented an approach that makes use of model uncertainty to generate clarification questions, and compared it to other methods for deciding when clarification is needed. Our approach allows us to vary the level of uncertainty that elicits the generation of clarification questions to control for the number of questions generated. The results of our experiments show that our proposal outperforms a number of baselines and other methods trained on human clarification decisions. More generally, our findings pave the way for further research on integrating model uncertainty into natural language generation approaches as an effective way to resolve ambiguities and address underspecification in language use (Pezzelle, 2023).

Limitations

Even though our experimental setup (which focuses on the *size* attribute using a template-based generation approach) allows us to isolate the effect of clarification exchanges in a controlled setting, we acknowledge that further work is required to confirm our findings in more complex collaborative tasks and model architectures. In our work, we focused on questions about the *size* attribute, and we mentioned some results about expanding our approach to the *orientation* attribute to assess the robustness of our approach. Assessing the contribution of clarification exchanges is particularly challenging given the wide range of factors involved, both on the appropriateness of the question asked and the correctness of the answer received. Without some restrictions, it is difficult to disentangle the role of these factors and this may threaten a fair model evaluation. For this reason, we focus on the *size* attribute, as it applies to all cliparts and it is often underspecified in the Teller’s instructions (see Section 5). An additional advantage of focusing on the *size* attribute is that we can ensure the correctness of the answer (see the templates in

Section 6.1), while for other attributes such as position there is a less direct mapping between natural language expressions and space coordinates.

Regarding the generalization of our setup to other settings, we would like to highlight that, although simple, the CoDraw task shares many features with a wide variety of different tasks and datasets. In particular, in Vision-and-Language Navigation tasks an agent has to reach a goal by moving around an environment and performing specific tasks. We believe this setting, where the agent has a set of actions at their disposal, shares many similarities with the CoDraw setup. On the other hand, in more complex tasks the role of the visual input may play a greater role compared to CoDraw, leading to new challenges in evaluating the grounding of the generated question-answer exchanges. At the same time, we acknowledge the fact that a template-based approach may very well not scale up to other settings. However, in our work, this choice is motivated by the focus on the effect of clarifying rather than the question generation problem.

Finally, we only used a simple approach to control for the “cost” of asking clarification questions by limiting the overall number of possible questions. We acknowledge that a more fine-grained approach is required to consider the wide variety of different factors that influence the cost of generating clarification requests. Given the simplicity of the CoDraw task and the structure of the questions our method generates, we do not analyze how clarification questions may influence the perceived naturalness of the generated dialogues. We acknowledge that this aspect should be taken into account when dealing with more complex tasks and generative strategies.

Acknowledgments

We would like to thank Brielen Madureira for discussions about the CoDraw task and dataset, as well as for sharing insights on CoDraw-iCR. We are grateful to Patrick Kahardipraja for sharing an updated version of the original CoDraw codebase. Finally, we thank the Dialogue Modelling Group (DMG) at the University of Amsterdam for their feedback and support at the different stages of this work. This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 819455).

References

- Mohammad Aliannejadi, Julia Kiseleva, Aleksandr Chuklin, Jeff Dalton, and Mikhail Burtsev. 2021. [Building and evaluating open-domain dialogue corpora with clarifying questions](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4473–4484, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Joris Baan, Nico Daheim, Evgenia Ilia, Dennis Ulmer, Haau-Sing Li, Raquel Fernández, Barbara Plank, Rico Sennrich, Chrysoula Zerva, and Wilker Aziz. 2023. Uncertainty in natural language generation: From theory to applications. *arXiv preprint arXiv:2307.15703*.
- Luciana Benotti. 2009. Clarification potential of instructions. In *The 10th Annual SIGDIAL Meeting on Discourse and Dialogue*.
- Luciana Benotti and Patrick Blackburn. 2017. Modeling the clarification potential of instructions: Predicting clarification requests and other reactions. *Computer Speech & Language*, 45:536–551.
- Luciana Benotti and Patrick Blackburn. 2021. [A recipe for annotating grounded clarifications](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4065–4077, Online. Association for Computational Linguistics.
- Glenn W Brier. 1950. Verification of forecasts expressed in terms of probability. *Monthly weather review*, 78(1):1–3.
- Khyathi Raghavi Chandu, Yonatan Bisk, and Alan W Black. 2021. [Grounding ‘grounding’ in NLP](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 4283–4305, Online. Association for Computational Linguistics.
- Ta-Chung Chi, Minmin Shen, Mihail Eric, Seokhwan Kim, and Dilek Hakkani-tur. 2020. Just ask: An interactive learning framework for vision and language navigation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34-03, pages 2459–2466.
- Javier Chiyah-Garcia, Alessandro Suglia, Arash Eshghi, and Helen Hastie. 2023. ‘What are you referring to?’ Evaluating the ability of multi-modal dialogue models to process clarificational exchanges. In *Proceedings of the 24rd Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL’23)*.
- Michelle M Chouinard, Paul L Harris, and Michael P Maratsos. 2007. Children’s questions: A mechanism for cognitive development. *Monographs of the society for research in child development*, pages i–129.
- Herbert H Clark. 1996. *Using language*. Cambridge university press.

- Herbert H Clark and Susan E Brennan. 1991. Grounding in communication. In *Perspectives on socially shared cognition.*, pages 127–149. American Psychological Association.
- Herbert H Clark and Edward F Schaefer. 1987. Collaborating on contributions to conversations. *Language and cognitive processes*, 2(1):19–41.
- Marty J Demetras, Kathryn Nolan Post, and Catherine E Snow. 1986. Feedback to first language learners: The role of repetitions and clarification questions. *Journal of child language*, 13(2):275–292.
- Yang Deng, Wenqiang Lei, Lizi Liao, and Tat-Seng Chua. 2023. Prompting and evaluating large language models for proactive dialogues: Clarification, target-guided, and non-collaboration. *arXiv preprint arXiv:2305.13626*.
- Stefan Depeweg, Jose-Miguel Hernandez-Lobato, Fina Doshi-Velez, and Steffen Udluft. 2018. Decomposition of uncertainty in bayesian deep learning for efficient and risk-sensitive learning. In *International Conference on Machine Learning*, pages 1184–1193. PMLR.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. **BERT: Pre-training of deep bidirectional transformers for language understanding**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Aysu Ezen-Can. 2020. A comparison of LSTM and BERT for small corpus. *arXiv preprint arXiv:2009.05451*.
- Malte Gabsdil. 2003. Clarification in spoken dialogue systems. In *Proceedings of the 2003 AAI Spring Symposium. Workshop on Natural Language Generation in Spoken and Written Dialogue*, pages 28–35.
- Xiaofeng Gao, Qiaozi Gao, Ran Gong, Kaixiang Lin, Govind Thattai, and Gaurav S Sukhatme. 2022. Dialfred: Dialogue-enabled agents for embodied instruction following. *IEEE Robotics and Automation Letters*, 7(4):10049–10056.
- Spandana Gella, Aishwarya Padmakumar, Patrick Lange, and Dilek Hakkani-Tur. 2022. **Dialog acts for task driven embodied agents**. In *Proceedings of the 23rd Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 111–123, Edinburgh, UK. Association for Computational Linguistics.
- 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 22nd Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 353–359, Singapore and Online. Association for Computational Linguistics.
- Claudio Greco, Alberto Testoni, Raffaella Bernardi, and Stella Frank. 2022. **A small but informed and diverse model: The case of the multimodal Guess-What!? guessing game**. In *Proceedings of the 2022 CLASP Conference on (Dis)embodiment*, pages 1–10, Gothenburg, Sweden. Association for Computational Linguistics.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. 2017. On calibration of modern neural networks. In *International conference on machine learning*, pages 1321–1330. PMLR.
- Eyke Hüllermeier and Willem Waegeman. 2021. Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods. *Machine Learning*, 110:457–506.
- Alex Kendall and Yarin Gal. 2017. What uncertainties do we need in bayesian deep learning for computer vision? *Advances in neural information processing systems*, 30.
- Baber Khalid, Malihe Alikhani, and Matthew Stone. 2020. **Combining cognitive modeling and reinforcement learning for clarification in dialogue**. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 4417–4428, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Jin-Hwa Kim, Nikita Kitaev, Xinlei Chen, Marcus Rohrbach, Byoung-Tak Zhang, Yuandong Tian, Dhruv Batra, and Devi Parikh. 2019. **CoDraw: Collaborative drawing as a testbed for grounded goal-driven communication**. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6495–6513, Florence, Italy. Association for Computational Linguistics.
- Julia Kiseleva, Alexey Skrynnik, Artem Zholus, Shrestha Mohanty, Negar Arabzadeh, Marc-Alexandre Côté, Mohammad Aliannejadi, Milagro Teruel, Ziming Li, Mikhail Burtsev, et al. 2022. Iglu 2022: Interactive grounded language understanding in a collaborative environment at neurips 2022. *arXiv preprint arXiv:2205.13771*.
- Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. 2017. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30.
- Thao Minh Le, Vuong Le, Svetha Venkatesh, and Truyen Tran. 2021. Hierarchical conditional relation networks for multimodal video question answering. *International Journal of Computer Vision*, 129(11):3027–3050.
- Kathryn A Leech, Virginia C Salo, Meredith L Rowe, and Natasha J Cabrera. 2013. Father input and child vocabulary development: The importance of wh questions and clarification requests. In *Seminars in speech and language*, volume 34-04, pages 249–259. Thieme Medical Publishers.

- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Brielen Madureira and David Schlangen. 2023a. "Are you telling me to put glasses on the dog?" Content-grounded annotation of instruction clarification requests in the codraw dataset. *arXiv preprint arXiv:2306.02377*.
- Brielen Madureira and David Schlangen. 2023b. [Instruction clarification requests in multimodal collaborative dialogue games: Tasks, and an analysis of the CoDraw dataset](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2303–2319, Dubrovnik, Croatia. Association for Computational Linguistics.
- Shrestha Mohanty, Negar Arabzadeh, Milagro Teruel, Yuxuan Sun, Artem Zhohus, Alexey Skrynnik, Mikhail Burtsev, Kavya Srinet, Aleksandr Panov, Arthur Szlam, et al. 2022. Collecting interactive multi-modal datasets for grounded language understanding. *arXiv preprint arXiv:2211.06552*.
- Mahdi Pakdaman Naeni, Gregory Cooper, and Milos Hauskrecht. 2015. Obtaining well calibrated probabilities using bayesian binning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 29.
- Anjali Narayan-Chen, Prashant Jayannavar, and Julia Hockenmaier. 2019. [Collaborative dialogue in Minecraft](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5405–5415, Florence, Italy. Association for Computational Linguistics.
- Kata Naszadi, Putra Manggala, and Christof Monz. 2023. [Aligning predictive uncertainty with clarification questions in grounded dialog](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 14988–14998, Singapore. Association for Computational Linguistics.
- Khanh Nguyen and Hal Daumé III. 2019. [Help, anna! visual navigation with natural multimodal assistance via retrospective curiosity-encouraging imitation learning](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 684–695, Hong Kong, China. Association for Computational Linguistics.
- Yaniv Ovadia, Emily Fertig, Jie Ren, Zachary Nado, David Sculley, Sebastian Nowozin, Joshua Dillon, Balaji Lakshminarayanan, and Jasper Snoek. 2019. Can you trust your model’s uncertainty? evaluating predictive uncertainty under dataset shift. *Advances in neural information processing systems*, 32.
- Aishwarya Padmakumar, Jesse Thomason, Ayush Shrivastava, Patrick Lange, Anjali Narayan-Chen, Spandana Gella, Robinson Piraamuthu, Gokhan Tur, and Dilek Hakkani-Tur. 2022. Teach: Task-driven embodied agents that chat. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36-2, pages 2017–2025.
- Tim Paek and Eric Horvitz. 2000. Conversation as action under uncertainty. In *Proceedings of the 16th Conference on Uncertainty in Artificial Intelligence*, pages 455–464.
- Sandro Pezzelle. 2023. [Dealing with semantic under-specification in multimodal NLP](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12098–12112, Toronto, Canada. Association for Computational Linguistics.
- Sandro Pezzelle and Raquel Fernández. 2023. Semantic adaptation to the interpretation of gradable adjectives via active linguistic interaction. *Cognitive Science*, 47(2):e13248.
- J Piaget. 1954. The construction of reality in the child (vol. xiii). *New York: Basic Books*. <https://doi.org/10.1037:11168-000>.
- Matthew Purver, Jonathan Ginzburg, and Patrick Healey. 2001. [On the means for clarification in dialogue](#). In *Proceedings of the Second SIGdial Workshop on Discourse and Dialogue*.
- Matthew Richard John Purver. 2004. *The theory and use of clarification requests in dialogue*. Ph.D. thesis, University of London.
- Verena Rieser and Oliver Lemon. 2006. [Using machine learning to explore human multimodal clarification strategies](#). In *Proceedings of the COLING/ACL 2006 Main Conference Poster Sessions*, pages 659–666, Sydney, Australia. Association for Computational Linguistics.
- 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)*.
- Anselm Rothe, Brenden M Lake, and Todd M Gureckis. 2018. Do people ask good questions? *Computational Brain & Behavior*, 1:69–89.
- Azzurra Ruggeri, Tania Lombrozo, Thomas L Griffiths, and Fei Xu. 2016. Sources of developmental change in the efficiency of information search. *Developmental psychology*, 52(12):2159.

- Azzurra Ruggeri, Zi Lin Sim, and Fei Xu. 2017. “why is toma late to school again?” preschoolers identify the most informative questions. *Developmental psychology*, 53(9):1620.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.
- Emanuel A. Schegloff. 1987. Some sources of misunderstanding in talk-in-interaction. *Linguistics*, 25(287):201–218.
- 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 SIGdial Workshop on Discourse and Dialogue at HLT-NAACL 2004*, pages 136–143.
- Julian J Schlöder and Raquel Fernández. 2014. Clarification requests on the level of uptake. In *Proceedings of the 18th Workshop on the Semantics and Pragmatics of Dialogue (SemDial 2014, ‘Dial-Watt’)*, pages 237–239. Citeseer.
- Ivan Sekulić, Mohammad Aliannejadi, and Fabio Crestani. 2021. Towards facet-driven generation of clarifying questions for conversational search. In *Proceedings of the 2021 ACM SIGIR international conference on theory of information retrieval*, pages 167–175.
- Ying Shen and Ismini Lourentzou. 2023. Learning by asking for embodied visual navigation and task completion. *arXiv preprint arXiv:2302.04865*.
- Zhengxiang Shi, Yue Feng, and Aldo Lipani. 2022. [Learning to execute actions or ask clarification questions](#). In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 2060–2070, Seattle, United States. Association for Computational Linguistics.
- Gabriel Skantze. 2008. Galatea: A discourse modeller supporting concept-level error handling in spoken dialogue systems. *Recent Trends in Discourse and Dialogue*, pages 155–189.
- Svetlana Stoyanchev, Alex Liu, and Julia Hirschberg. 2014. Towards natural clarification questions in dialogue systems. In *AISB symposium on questions, discourse and dialogue*, volume 20.
- Alberto Testoni and Raffaella Bernardi. 2021. [Looking for confirmations: An effective and human-like visual dialogue strategy](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 9330–9338, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Michael Tomasello. 2009. *The cultural origins of human cognition*. Harvard university press.
- Xin Wei, Jian Wu, Kehinde Ajayi, and Diane Oyen. 2022. Visual descriptor extraction from patent figure captions: A case study of data efficiency between bilstm and transformer. In *Proceedings of the 22nd ACM/IEEE Joint Conference on Digital Libraries*, pages 1–5.
- Albert Zeyer, Parnia Bahar, Kazuki Irie, Ralf Schlüter, and Hermann Ney. 2019. A comparison of Transformer and LSTM encoder decoder models for asr. In *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 8–15. IEEE.
- Zhiling Zhang and Kenny Zhu. 2021. Diverse and specific clarification question generation with keywords. In *Proceedings of the Web Conference 2021*, pages 3501–3511.
- 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1594–1603.
- C Lawrence Zitnick and Devi Parikh. 2013. Bringing semantics into focus using visual abstraction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3009–3016.
- C Lawrence Zitnick, Devi Parikh, and Lucy Vanderwende. 2013. Learning the visual interpretation of sentences. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1681–1688.

Appendix

A Similarity Score (v2)

In our work, we apply some changes to the similarity score metric proposed in [Kim et al. \(2019\)](#) to fix some flaws in its design and better capture the properties under analysis in the experimental setup. More specifically, (a) we assign the same weights to all components of the similarity score; (b) we consider the *orientation* attribute only for those cliparts that are not perfectly symmetrical; (c) we change the denominator of the *facial expression* and *body pose* attributes so that they apply only to cliparts representing people. Using the new metric, scenes drawn by human players reach a similarity score of 4.42 (sd=0.62).

B Model vs. Human Uncertainty

Logistic Regression results and analysis Table 3 shows the Average Precision of the logistic regression model trained to predict when human players

ask clarification questions in CoDraw based on model uncertainty. We extract uncertainty from Drawer models trained for different training epochs and initialization seeds. We also report the performance of a random baseline and the CoDraw similarity score for each configuration. Similarly, Table 4 reports the F1 score of the same model. Table 5 shows the correlation coefficients of the logistic regression model for each attribute-specific uncertainty analyzed (clipart selection, size, orientation, position). We report the coefficients for each model configuration (initialization seed + training epoch). Finally, Table 6 illustrates some insights and examples of the “clusters” analyzed in Section 5 to investigate the variability of human confirmation-seeking strategies.

Epistemic and Aleatoric uncertainty We also experimented with including Epistemic and Aleatoric uncertainty in the Logistic Regression analysis. We obtain the same results as using the other uncertainty measures described above, without any statistically significant predictive power. Similarly to Greco et al. (2022), we use an Ensemble of models trained with different initialization seeds to compute *data* (aleatoric) and *model* (epistemic) uncertainty (Depeweg et al., 2018; Hüllermeier and Waegeman, 2021) for the size and orientation attributes. The ensemble *total uncertainty* is computed as the entropy of its predictive distribution $[p_{\theta_n}(y|x)]$, which averages the distributions of the single ensemble components:

$$H[p(y|x)] = H[1/N \sum_{n=1}^N p_{\theta_n}(y|x)]. \quad (1)$$

Data uncertainty is measured as the average uncertainty of each ensemble component: $1/N \sum_{n=1}^N H[p_{\theta_n}(y|x)]$. The difference between total uncertainty and data uncertainty results in *model uncertainty*. As mentioned above, including these uncertainties in the logistic regression model described in Section 5 does not improve its Average Precision, which remains significantly above chance level but low in absolute values (0.19 AP, similar to the values reported in Figure 2).

C Asking Clarification Questions

Figure 5 shows some examples of clarification questions asked by different models for the same dialogue.

D The CoDraw Dataset and Models

The Codraw dataset (Kim et al., 2019) is available at this URL: <https://github.com/facebookresearch/CoDraw>. CoDraw is licensed under Creative Commons Attribution-NonCommercial 4.0 International Public License. The CoDraw dataset is fully anonymized. The Silent drawer model is available at this address: <https://github.com/facebookresearch/codraw-models>. This repository is licensed under Creative Commons Attribution-NonCommercial 4.0 International Public License. We use the released data and models consistently with their intended use. Codraw-iCR (v1/v2) is licensed under Creative Commons Attribution-NonCommercial 4.0 International Public License. The CoDraw dataset contains 9993 dialogues, with a total number of 138K utterances. The dataset is randomly split into training, validation, and test set with the following proportions: 80%, 10%, 10%, respectively.

E QDrawer Details

The QDrawer model has the same architecture as the *Silent Drawer* described in Section 4.1. The model has 3.347.390 parameters. For training configuration, we follow the original paper (Kim et al., 2019). We trained the model on a NVIDIA RTX A5000 GPU for 15 epochs (\approx 1 minute/epoch). In QDrawer, we compute the entropy using the SciPy library (<https://docs.scipy.org/doc/scipy/index.html>). As mentioned in Section 6, we slightly modify the loss function of the *Silent Drawer*. More specifically, the *Silent Drawer* is trained using a combination of losses, including a cross-entropy loss for categorical decisions on which clip art pieces to add to the canvas at a given dialogue turn. This loss considers which new cliparts human players added to the scene at each turn of the dialogue. If the Drawer player updates some attributes of a clipart that was added to the canvas in previous turns (for instance, changing its size or location), this clipart is not included in the loss function. This step is crucial because, at inference time, only the cliparts whose score exceeds a threshold are added and updated in the canvas. To address this issue, at each dialogue turn, we include in the loss function of QDrawer not only the new cliparts added to the canvas but also the ones already in the canvas that are modified by the Drawer at that turn. In this

way, clarification exchanges on cliparts added in previous dialogue turns can be properly addressed.

		<i>Average Precision</i>						<i>Similarity Score</i>			
		training epoch						training epoch			
		1	5	10	15			1	5	10	15
seed	0	.176 (.134)	.186 (.150)	.181 (.143)	.205 (.143)	seed	0	3.02	3.33	3.32	3.31
	1	.179 (.150)	.171 (.152)	.185 (.138)	.172 (.126)		1	3.05	3.33	3.31	3.34
	2	.168 (.142)	.181 (.139)	.176 (.137)	.196 (.148)		2	3.02	3.31	3.34	3.36
	3	.194 (.145)	.190 (.153)	.213 (.154)	.189 (.161)		3	3.01	3.30	3.35	3.34
	4	.202 (.156)	.179 (.145)	.176 (.125)	.180 (.135)		4	3.04	3.31	3.35	3.35

Table 3: On the left, Average Precision of the logistic regression model trained to predict human decisions to ask clarification questions based on model uncertainty. In parenthesis, we report the performance of a random model, which is always significantly lower than its trained counterpart ($p < 0.05$). We study the effect of different initialization seeds and training epochs. On the right, we show the corresponding average similarity score reached by the Drawer model at the end of the dialogue in the CoDraw task.

		<i>F1 Score - Logistic Regression Model</i>						<i>F1 Score - Random baseline</i>			
		training epoch						training epoch			
		1	5	10	15			1	5	10	15
seed	0	0.479	0.482	0.477	0.468	seed	0	0.469	0.492	0.509	0.513
	1	0.493	0.499	0.485	0.463		1	0.500	0.525	0.490	0.488
	2	0.475	0.487	0.472	0.481		2	0.505	0.495	0.504	0.503
	3	0.469	0.481	0.496	0.486		3	0.481	0.507	0.474	0.520
	4	0.483	0.479	0.474	0.474		4	0.493	0.514	0.490	0.519

Table 4: F1 score of the logistic regression model (left) and the performance of a random baseline (right).

seed	epoch	uncertainty(clipart selection)	uncertainty(size)	uncertainty(orientation)	uncertainty(position)
0	1	-0.397	0.002	-0.041	-0.096
0	5	-0.392	0.006	-0.038	-0.012
0	10	-0.390	0.004	-0.044	-0.075
0	15	-0.390	0.006	-0.044	0.040
1	1	-0.397	0.002	-0.041	-0.087
1	5	-0.392	0.006	-0.038	0.016
1	10	-0.390	0.004	-0.044	-0.027
1	15	-0.390	0.006	-0.044	0.073
2	1	-0.397	0.002	-0.041	-0.070
2	5	-0.392	0.006	-0.038	0.025
2	10	-0.390	0.004	-0.044	-0.054
2	15	-0.390	0.006	-0.044	0.103
3	1	-0.397	0.002	-0.041	-0.054
3	5	-0.392	0.006	-0.038	-0.037
3	10	-0.390	0.004	-0.044	0.007
3	15	-0.390	0.006	-0.044	-0.015
4	1	-0.397	0.002	-0.041	-0.167
4	5	-0.392	0.006	-0.038	-0.019
4	10	-0.390	0.004	-0.044	-0.048
4	15	-0.390	0.006	-0.044	-0.017

Table 5: Correlation coefficients of the logistic regression model for each uncertainty type extracted from the Drawer model. We show the coefficients for each model configuration, as defined by its initialization seed and number of training epochs.

	# utts.	Clip.	Size	Loc.	Examples	% CQ	CQ type
A	3805	1	1	≥ 1	- large pine tree on right half of top and right edge cut off - middle scene swing, small size. facing right. red part above horizon	7.6%	other 48.3% loc. 43.8%
B	240	≥ 1	-	0	- little girl running towards her friend that has a pie in his hand - girl and boy, boy is on swing set	48.8%	loc. 48.3% other 35.6%
C	2274	≥ 1	0	-	- a tree to the left a few inches from the edge - left side is a bear, facing right, left arm cut from scene. bear cut in half by horizon line	23.1%	other 40.6% size 27.1%
D	76	1	-	0	- it 's at a park with a side and a sandbox - up in the tree are 4 balloons . small size	46.1%	loc. 57.1% other 22.9%
E	1388	1	0	-	- boy far lower right , facing left with happy face - girl sitting on ground , fingers touching edge of left screen , 1 2 " from bottom screen	19.5%	size 40.0% other 30.4%
F	5328	1	-	-	- small maple tree on left . almost all leaves above grass line - left top small sun left edge little bit cut	10.1%	other 40.1% loc. 34.8%

Table 6: Cluster analysis by grouping together similar instructions. # utts. refers to the number of utterances in the cluster; Clip., Size, and Loc. refer to the occurrences of clipart, size, and location keywords, respectively. % CQ refers to the percentage of clarification questions following the given set of instructions. “-” means no restrictions on that attribute.

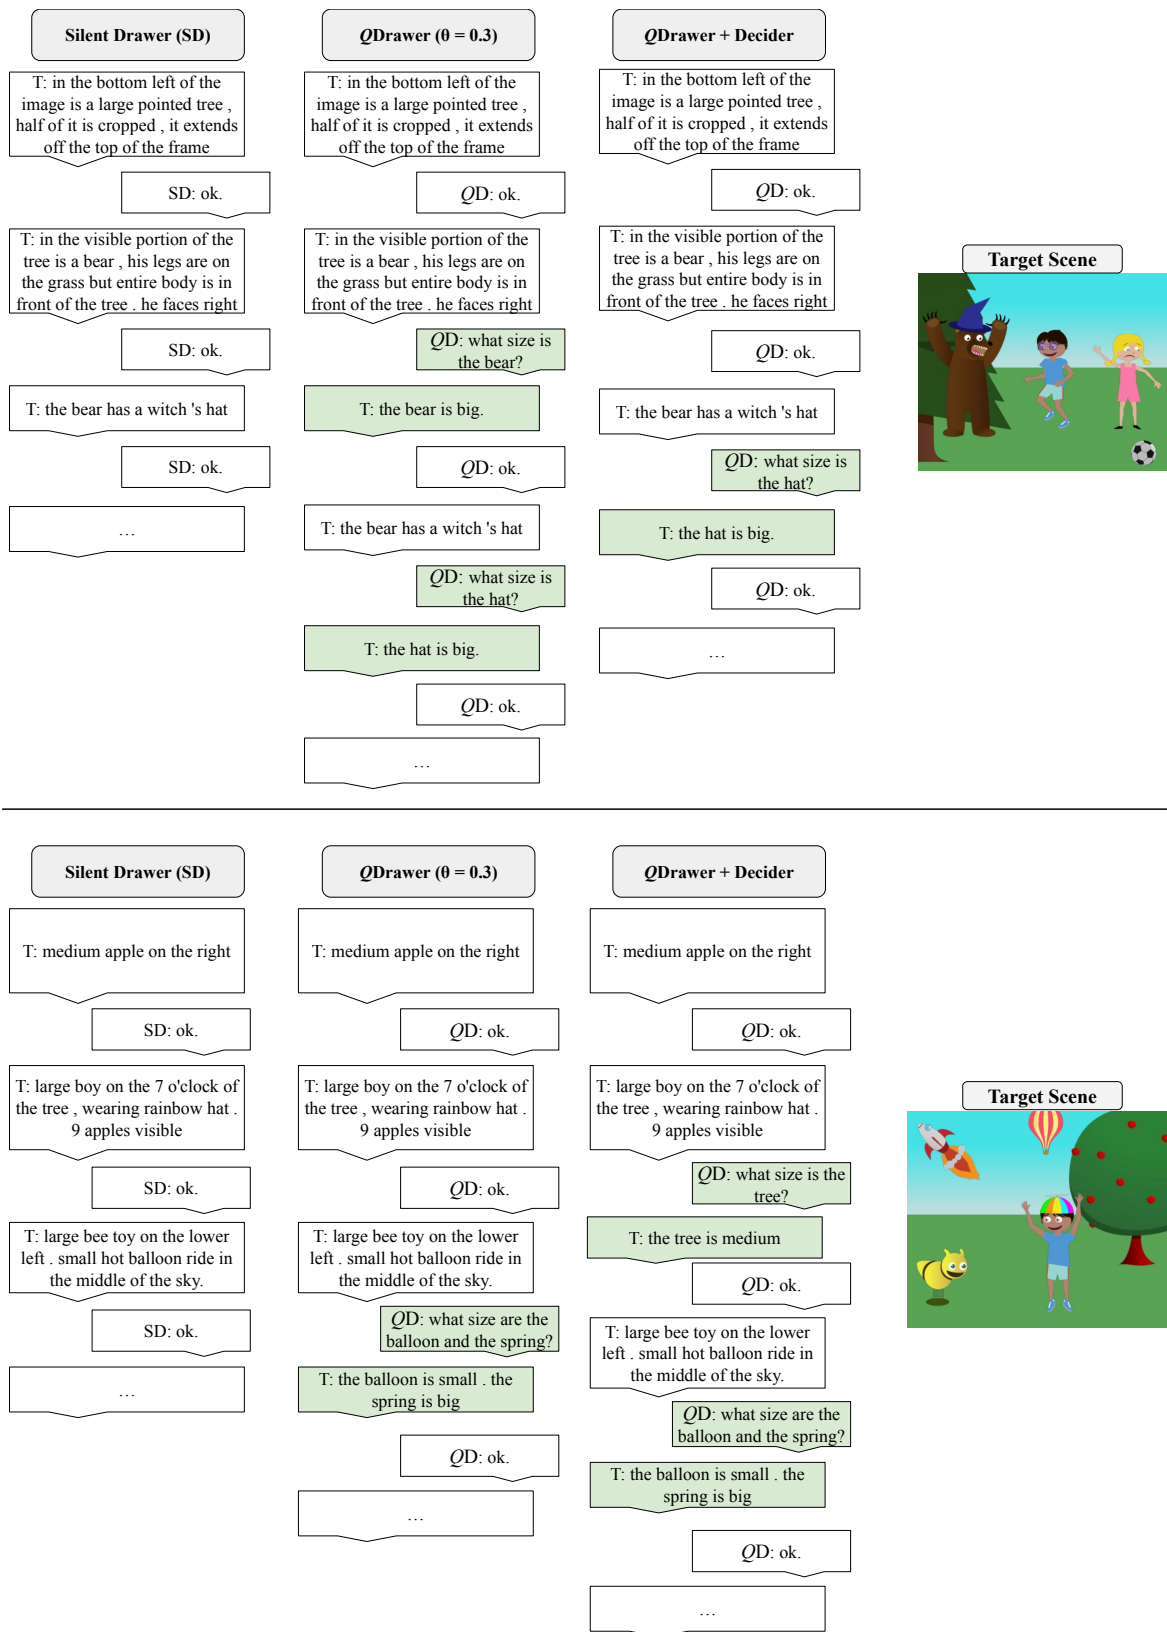


Figure 5: Examples of dialogues and clarification questions asked by different models. *T* stands for Teller, *SD* for Silent Drawer, and *QD* for QDrawer. Clarification exchanges are highlighted in green.