

EVALUATING VISUAL CONVERSATIONAL AGENTS VIA COOPERATIVE HUMAN-AI GAMES

A Thesis
Presented to
The Academic Faculty

By

Prithvijit Chattopadhyay

In Partial Fulfillment
of the Requirements for the Degree
Master of Science in the
College of Computing

Georgia Institute of Technology

May 2019

Copyright © Prithvijit Chattopadhyay 2019

EVALUATING VISUAL CONVERSATIONAL AGENTS VIA COOPERATIVE HUMAN-AI GAMES

Approved by:

Dr. Devi Parikh, Advisor
School of Interactive Computing
Georgia Institute of Technology

Dr. Dhruv Batra
School of Interactive Computing
Georgia Institute of Technology

Dr. Stefan Lee
School of Interactive Computing
Georgia Institute of Technology

Date Approved: April 24, 2019

It doesn't make a difference how beautiful your guess is.

It doesn't make a difference how smart you are,
who made the guess, or what his name is.

If it disagrees with experiment, it's wrong.

Richard P. Feynman

To Ma, Baba and Boopka, for always being there for me.

To Devi and Dhruv, for believing in me.

ACKNOWLEDGEMENTS

I would like to thank my advisor, Prof. Devi Parikh. Thank you for believing in and providing a naive undergraduate student (who used to romanticise research for the most part) an opportunity to conduct actual research at Virginia Tech four years ago and then later at Georgia Tech. Devi has my immense gratitude for not only supporting and teaching me the right set of skills – ranging from brainstorming to research to day-to-day planning to execution – but also for cultivating an environment where I had the opportunity to explore, learn and freely pursue problems. Thanks for being patient with me, bearing with all the ridiculous ideas I might have pitched and teaching me what the right set of questions to ask are. Throughout my interactions with Devi (across project meetings, group meetings and otherwise), I have learned how important it is to look at a problem holistically, rather than being stuck at a minute detail or aspect that I might find exciting. I have found Devi extremely convenient to look upto and reach out when experiments are not working out and seek help to correct the course of projects under time constraints. I always have and will look forward to any feedback Devi provides me at any point in my career.

I would also like to thank Prof. Dhruv Batra, for providing me with valuable feedback, guidance and motivation over the years. As a part of the lab, I have always felt lucky that I have had not one, but two advisors. Over multiple interactions, either directly or indirectly, one important lesson I have learned from Dhruv is to reason about the most important bits of a contribution in research. Thanks for correcting the course of my thought process several times about a problem and showing me the overall picture whenever I was stuck at a seemingly minute detail. There are innumerable lessons Devi and Dhruv have taught me along the way that I will always be grateful for.

I have had the fortune to work with extremely talented peers and mentors in grad school by being a part of the Visual Intelligence Lab. I would like to thank Dr. Stefan Lee for being a constant source of motivation and for being around for all the fighting in the trenches

around deadlines. Thanks for always making me feel welcome for all the boardgame sessions outside work – they were indeed necessary doses of normalcy. I have thoroughly enjoyed all our discussions both professionally and outside work. I would also like to thank Dr. Ramakrishna Vedantam for being an amazing mentor, peer and friend. I have found Rama to be the person I could consistently reach out to discuss ideas irrespective of them being in a nebulous or a well-defined state. Thanks for teaching me to not accept ideas at face value but to ask “why” and convince myself of them. I would like to thank Deshraj for being an awesome flatmate, friend and being a constant presence through the good and bad times. I would like to thank Abhishek for the constant support and encouragement; Harsh for making me feel I had never left Delhi; Ram for taking care of me on multiple occasions and demonstrating what having faith in an idea means and, Arjun, Ashwin and Samyak for their friendship, advice, fun conversations and always being there to listen to me. A special shoutout to - Viraj, Nirbhay, Karan, Mohit, Ayush, Rishabh, Aishwarya, Yash, Nilaksh and Sanya for making grad school fun through trips, celebrations and providing a superb environment.

Finally, I would like to thank the people without whose support and countless sacrifices none of this would have been possible and I would not be in the position I am now. I thank my parents, Sanjukta Chattopadhyay and Kalyanjit Chattopadhyay, for always believing in me, supporting me in my endeavours both financially and emotionally, and providing an ideal to strive towards. Thank you for guiding me as a person and instilling in me a love for mathematics and the joy of reading. I thank my brother, Riddhijit Chattopadhyay, for his love, unwavering support and for reminding me in bad times that things can change. Without such a team behind me, I doubt that I would be in this place today.

TABLE OF CONTENTS

Acknowledgments	v
List of Tables	ix
List of Figures	x
Chapter 1: Introduction	1
1.1 Motivation	1
1.2 GuessWhich	2
1.3 Contributions	4
1.4 Thesis Outline	4
Chapter 2: Related Work	6
2.1 Visual Conversational Agents	6
2.2 Human Computation Games	6
2.3 Evaluating Conversational Agents.	7
2.4 Turing Test	7
Chapter 3: The AI: Alice	8
Chapter 4: Our GuessWhich Game	10
4.1 Game Description	10

4.2	Pool Selection	11
4.3	Data Collection and Player Reward Structure	12
4.4	Evaluation	13
Chapter 5: Infrastructure		14
Chapter 6: Results		16
6.1	ALICE _{SL} vs. ALICE _{RL}	16
6.2	Human perception of AI teammate	19
6.3	Questioning Strategies	20
Chapter 7: Challenges		23
7.1	Knowledge Leak	23
7.2	Engagement vs. Fairness	23
Chapter 8: Conclusion		25
Appendix A: Publications		27
References		30

LIST OF TABLES

6.1	Performance of Human-ALICE teams with ALICE_{SL} and ALICE_{RL} measured by MR (lower is better) and MRR (higher is better). Error bars are 95% CIs from 1000 bootstrap samples. Unlike (Das et al., 2017b), we find no significant difference between ALICE_{SL} and ALICE_{RL}	16
6.2	Performance of Human-ALICE and QBOT-ALICE teams measured by MR (lower is better). We observe that AI-AI teams outperform human-AI teams.	18

LIST OF FIGURES

1.1	A human and an AI (a visual conversation agent called ALICE) play the proposed GuessWhich game. At the start of the game (top), ALICE is provided an image (shown above ALICE) which is unknown to the human. Both ALICE and the human are then provided a brief description of the image. The human then attempts to identify the secret image. In each subsequent round of dialog, the human asks a question about the unknown image, receives an answer from ALICE, and makes a best guess of the secret image from a fixed pool of images. After 9 rounds of dialog, the human makes consecutive guesses until the secret image is identified. The fewer guesses the human needs to identify the secret image, the better the human-AI team performance.	3
4.1	GuessWhich Interface: A user asks a question to ALICE in each round and ALICE responds with an answer. The user then selects an appropriate image which they think is the secret image after each round of conversation. At the end of the dialog, user successively clicks on their best guesses until they correctly identify the secret image.	11
5.1	We outline the backend architecture of our implementation of GuessWhich. Since GuessWhich requires a live interaction between the human and the AI, we design a workflow that can handle multiple queues and can quickly pair a human with an AI agent.	14
6.1	Mean rank (MR) of secret image across (a) number of games and (b) rounds of dialog. Lower is better. Error bars are 95% confidence intervals from 1000 bootstrap samples.	17
6.2	Worker ratings for ALICE _{SL} and ALICE _{RL} on 6 metrics. Higher is better. Error bars are 95% confidence intervals from 1000 bootstrap samples. Humans perceive no significant differences between ALICE _{SL} and ALICE _{RL} across the 6 feedback metrics.	19

6.3	Distribution of first n-grams for questions asked to ALICE. Word ordering starts from the center and radiates outwards. Arc length is proportional to the number of questions containing the word. The most common question-types are binary – followed by ‘What color..’ questions.	21
6.4	We contrast two games played by different workers with ALICE _{SL} and ALICE _{RL} on the same pool (secret image outlined in green). In both cases, the workers are able to find the secret image within three guesses. It is also interesting to note how the answers provided by ALICE are different in the two cases. . . .	22

SUMMARY

As AI continues to advance, human-AI teams are inevitable. However, progress in AI is routinely measured in isolation, without a human in the loop. It is crucial to benchmark progress in AI, not just in isolation, but also in terms of how it translates to helping humans perform certain tasks, *i.e.*, the performance of human-AI teams.

This thesis introduces a cooperative game – GuessWhich – to measure human-AI team performance in the specific context of the AI being a visual conversational agent. GuessWhich involves live interaction between the human and the AI. The AI, which we call ALICE, is provided an image which is unseen by the human. Following a brief description of the image, the human questions ALICE about this secret image to identify it from a fixed pool of images.

We measure performance of the human-ALICE team by the number of guesses it takes the human to correctly identify the secret image after a fixed number of dialog rounds with ALICE. We compare performance of the human-ALICE teams for two versions of ALICE. Our human studies suggest a counter-intuitive trend – that while AI literature shows that one version outperforms the other when paired with an AI questioner bot, we find that this improvement in AI-AI performance does not translate to improved human-AI performance. As this implies a mismatch between benchmarking of AI in isolation and in the context of human-AI teams, this thesis further motivates the need to evaluate AI additionally in the latter setting to effectively leverage the progress in AI for efficient human-AI teams.

CHAPTER 1

INTRODUCTION

1.1 Motivation

As Artificial Intelligence (AI) systems become increasingly accurate and interactive (*e.g.* Alexa, Siri, Cortana, Google Assistant), human-AI teams are inevitably going to become more commonplace. To be an effective teammate, an AI must overcome the challenges involved with adapting to humans; however, progress in AI is routinely measured in isolation, without a human in the loop. In this work, we focus specifically on the evaluation of visual conversational agents and develop a human computation game to benchmark their performance as members of human-AI teams.

Visual conversational agents [1, 2, 3, 4] are AI agents trained to understand and communicate about the contents of a scene in natural language. For example, in Fig. 1.1, the visual conversational agent (shown on the right) replies to answers questions about a scene while inferring context from the dialog history – Human: “What is he doing?” Agent: “Playing frisbee”. These agents are typically trained to mimic large corpora of human-human dialogs and are evaluated automatically on how well they retrieve actual human responses (ground truth) in novel dialogs.

Recent work has evaluated these models more pragmatically by evaluating how well pairs of visual conversational agents perform on goal-based conversational tasks rather than response retrieval from fixed dialogs. Specifically, [2] train two visual conversational agents – a questioning bot QBOT, and an answering bot ABOT – for an image-guessing task. Starting from a description of the scene, QBOT and ABOT converse over multiple rounds of questions (QBOT) and answers (ABOT) in order to improve QBOT’s understanding of a secret image known only to ABOT. After a fixed number of rounds, QBOT must guess the secret image

from a large pool and both QBOT and ABOT are evaluated based on this guess.

[2] compare supervised baseline models with QBOT-ABOT teams trained through reinforcement learning based self-talk on this image-guessing task. They find that the AI-AI teams improve significantly at guessing the correct image after self-talk updates compared to the supervised pretraining. While these results indicate that the self-talk fine-tuned agents are better visual conversational agents, crucially, it remains unclear if these agents are indeed better at this task when *interacting with humans*.

1.2 GuessWhich

In this work, we propose to evaluate if and how this progress in AI-AI evaluation translates to the performance of human-AI teams. Inspired by the popular GuessWhat or 20-Questions game, we design a human computation game – GuessWhich – which requires collaboration between human and visual conversational AI agents. Mirroring the setting of [2], GuessWhich is an image-guessing game that consists of 2 participants – *questioner* and *answerer*. At the start of the game, the answerer is provided an image that is unknown to the questioner and both questioner and answerer are given a brief description of the image content. The questioner interacts with the answerer for a fixed number of rounds of question-answer (dialog) to identify the secret image from a fixed pool of images (see Fig. 1.1).

We evaluate human-AI team performance in GuessWhich, for the setting where the questioner is a human and the answerer is an AI (that we denote ALICE). Specifically, we evaluate two versions of ALICE for GuessWhich:

1. ALICE_{SL} which is trained in a supervised manner on the Visual Dialog dataset [1] to mimic the answers given by humans when engaged in a conversation with other humans about an image, and
2. ALICE_{RL} which is pre-trained with supervised learning and fine-tuned via reinforcement learning for an image-guessing task as in [2].




Figure 1.1: A human and an AI (a visual conversation agent called ALICE) play the proposed GuessWhich game. At the start of the game (top), ALICE is provided an image (shown above ALICE) which is unknown to the human. Both ALICE and the human are then provided a brief description of the image. The human then attempts to identify the secret image. In each subsequent round of dialog, the human asks a question about the unknown image, receives an answer from ALICE, and makes a best guess of the secret image from a fixed pool of images. After 9 rounds of dialog, the human makes consecutive guesses until the secret image is identified. The fewer guesses the human needs to identify the secret image, the better the human-AI team performance.

It is important to appreciate the difficulty and sensitivity of the GuessWhich game as an evaluation tool – agents have to understand human questions and respond with accurate, consistent, fluent and informative answers for the human-AI team to do well. Furthermore, they have to be robust to their own mistakes, *i.e.*, if an agent makes an error at a particular round, that error is now part of its conversation history, and it must be able to correct itself rather than be consistently inaccurate. Similarly, human players must also learn to adapt to ALICE’s sometime noisy and inaccurate responses.

At its core, GuessWhich is a game-with-a-purpose (GWAP) that leverages human computation to evaluate visual conversational agents. Traditionally, GWAP [5] have focused on *human-human collaboration*, *i.e.* collecting data by making humans play games to label im-

ages [6], music [7] and movies [8]. We extend this to human-AI teams and to the best of our knowledge, our work is the first to evaluate visual conversational agents in an interactive setting where humans are continuously engaging with agents to succeed at a cooperative game.

1.3 Contributions

More concretely, we make the following contributions in this work:

- We design an interactive image-guessing game (GuessWhich) for evaluating human-AI team performance in the specific context of the AIs being visual conversational agents. GuessWhich pairs humans with ALICE, an AI capable of answering a sequence of questions about images. ALICE is assigned a secret image and answers questions asked about that image from a human for 9 rounds to help them identify the secret image.
- We evaluate human-AI team performance on this game for both supervised learning (SL) and reinforcement learning (RL) versions of ALICE. Our main experimental finding is that despite significant differences between SL and RL agents reported in previous work [2], we find *no significant difference* in performance between ALICE_{SL} or ALICE_{RL} when paired with human partners. This suggests that while self-talk and RL are interesting directions to pursue for building better visual conversational agents, there appears to be a disconnect between AI-AI and human-AI evaluations – progress on former does not seem predictive of progress on latter. This is an important finding to guide future research.

1.4 Thesis Outline

The structure of the remainder of this work is as follows:

- In Chapter 2, we draw connections to relevant work on Visual Conversational Agents, Human Computation and Dialog Evaluation.
- In Chapter 3, we describe ALICE– the conversational agent being evaluated through the GuessWhich game. Specifically, we discuss the backend neural architecture of ALICE and how it is pretrained using supervised learning (SL) and finetuned using reinforcement learning (RL).
- In Chapter 4, we describe the GuessWhich game in detail. We start by describing the players and gameplay in GuessWhich. Following this, we describe how we construct the pool of images and discuss the player reward structure and evaluation protocols.
- In Chapter 5, we briefly describe the backend software architecture associated with GuessWhich.
- In Chapter 6, we discuss and summarize the results obtained from our experiments. We first discuss the performance of the human-AI teams in terms of retrieving the secret image after 9 rounds of dialog, when the human is paired with ALICE_{SL} and ALICE_{RL}. Following this, we discuss how retrieval performance changes across successive rounds of dialog and multiple games for both human-ALICE_{SL} and human-ALICE_{RL} teams. Next, we discuss the human subjects’ perception of the AI teammate and what kind of questioning strategies do humans adopt in GuessWhich.
- In Chapter 7, we highlight some of the challenges associated with evaluation of human-AI teams in GuessWhich– specifically, preventing knowledge leak and ensuring an appropriate trade-off between engagement and fairness.
- In Chapter 8, we briefly conclude and summarize our findings to motivate evaluating AI in the context of human-AI teams in addition to measuring progress in isolation.

CHAPTER 2

RELATED WORK

Given that our goal is to evaluate visual conversational agents through a human computation game, we draw connections to relevant work on visual conversational agents, human computation games, and dialog evaluation below.

2.1 Visual Conversational Agents

Our AI agents are visual conversational models, which have recently emerged as a popular research area in visually-grounded language modeling [1, 2, 3, 4]. [1] introduced the task of Visual Dialog and collected the VisDial dataset by pairing subjects on Amazon Mechanical Turk (AMT) to chat about an image (with assigned roles of questioner and answerer). [2] pre-trained questioner and answerer agents on this VisDial dataset via supervised learning and fine-tuned them via self-talk (reinforcement learning), observing that RL-fine-tuned QBOT-ABOT are better at image-guessing after interacting with each other. However, they do not evaluate if this change in QBOT-ABOT performance translates to human-AI teams.

2.2 Human Computation Games

Human computation games have been shown to be time- and cost-efficient, reliable, intrinsically engaging for participants [9, 10], and hence an effective method to collect data annotations. There is a long line of work on designing such Games with a Purpose (GWAP) [5] for data labeling purposes across various domains including images [6, 11, 12, 13], audio [14, 7], language [15, 16], movies [8] *etc.* While such games have traditionally focused on human-human collaboration, we extend these ideas to human-AI teams. Rather than collecting labeled data, our game is designed to measure the effectiveness of the AI in the

context of human-AI teams.

2.3 Evaluating Conversational Agents.

Goal-driven (non-visual) conversational models have typically been evaluated on task-completion rate or time-to-task-completion [17], so shorter conversations are better. At the other end of the spectrum, free-form conversation models are often evaluated by metrics that rely on n-gram overlaps, such as BLEU, METEOR, ROUGE, but these have been shown to correlate poorly with human judgment [18]. Human evaluation of conversations is typically in the format where humans rate the quality of machine utterances given context, without actually taking part in the conversation, as in [2] and [19]. To the best of our knowledge, we are the first to evaluate conversational models via team performance where humans are continuously interacting with agents to succeed at a downstream task.

2.4 Turing Test

Finally, our GuessWhich game is in line with ideas in [20], re-imagining the traditional Turing Test for state-of-the-art AI systems, taking the pragmatic view that an effective AI teammate need not appear human-like, act or be mistaken for one, provided its behavior does not feel jarring or baffle teammates, leaving them wondering not about what it is thinking but whether it is.

In the following chapter, we formally define the AI agent ALICE, describe the GuessWhich game setup, and present results and analysis from human studies.

CHAPTER 3

THE AI: ALICE

Recall from Chapter 1 that our goal is to evaluate how progress in AI measured through automatic evaluation translates to performance of human-AI teams in the context of visual conversational agents. Specifically, we are considering the question-answering agent ABOT from [2] as ABOT is the agent more likely to be deployed with a human partner in real applications (*e.g.* to answer questions about visual content to aid a visually impaired user). For completeness, we will review this work in this section.

[2] formulate a self-supervised image-guessing task between a questioner bot (QBOT) and an answerer bot (ABOT) which plays out over multiple rounds of dialog. At the start of the task, QBOT and ABOT are shown a one sentence description (*i.e.* a caption) of an image (unknown to QBOT). The pair can then engage in question and answer based dialog for a fixed number of iterations after which QBOT must try to select the secret image from a pool. The goal of the QBOT-ABOT team is two-fold, QBOT should: 1) build a mental model of the unseen image purely from the dialog and 2) be able to retrieve that image from a line-up of images.

Both QBOT and ABOT are modeled as Hierarchical Recurrent Encoder-Decoder neural networks [1, 21] which encode each round of dialog independently via a recurrent neural network (RNN) before accumulating this information through time with an additional RNN (resulting in hierarchical encoding). This representation (and a convolutional neural network based image encoding in ABOT’s case) are used as input to a decoder RNN which produces an agent’s utterance (question for QBOT and answer for ABOT) based on the dialog (and image for ABOT). In addition, QBOT includes an image feature regression network that predicts a representation of the secret image based on dialog history. We refer to [2] for complete model details.

These agents are pre-trained with supervised dialog data from the VisDial dataset [1] with a Maximum Likelihood Estimation objective. This pre-training ensures that agents can generally recognize objects/scenes and utter English. Following this, the models are fine-tuned by ‘smoothly’ transitioning to a deep reinforcement learning framework to directly improve image-guessing performance. This annealed transition avoids abrupt divergence of the dialog in face of an incorrect question-answer pair in the QBOT-ABOT exchange. During RL based self-talk, the agents’ parameters are updated by gradients corresponding to rewards depending on individual good or bad exchanges. We refer to the baseline supervised learning based ABOT as $ALICE_{SL}$ and the RL fine-tuned bot as $ALICE_{RL}$. [2] found that the AI-AI pair succeeds in retrieving the correct image more often after being fine-tuned with RL. In the following section, we outline our GuessWhich game designed to evaluate whether this improvement between $ALICE_{SL}$ and $ALICE_{RL}$ in automatic metrics translates to human-AI collaborations.

CHAPTER 4

OUR GUESSWHICH GAME

4.1 Game Description

We begin by describing our game setting; outlining the players and gameplay mechanics. A video of an example game being played can be found at <https://vimeo.com/229488160>.

Players. We replace QBOT in the AI-AI dialog with humans to perform a collaborative task of identifying a secret image from a pool. In the following, we will refer to ABOT as ALICE and the human player as H. We evaluate two versions of ALICE – ALICE_{SL} and ALICE_{RL}, where SL and RL correspond to agents *trained in a supervised setting* and *fine-tuned with reinforcement learning* respectively.

Gameplay. In our game setting, ALICE is assigned a secret image I^c (unknown to H) from a pool of images $\mathbb{I} = \{I_1, I_2, \dots, I_n\}$ taken from the COCO dataset [22]. Prior to beginning the dialog, both ALICE and H are provided a brief description (*i.e.* a caption) of I^c generated by Neuraltalk2 [23], an open-source implementation of [24]. H then makes a guess about the secret image by selecting one from the pool \mathbb{I} based only on the caption, *i.e.* before the dialog begins.

In each of the following rounds, H asks ALICE a question q_t about the secret image I^c in order to better identify it from the pool and ALICE responds with an answer a_t . After each round, H must select an image I^t that they feel is most likely the secret image I^c from pool \mathbb{I} based on the dialog so far. At the end of $k = 9$ rounds of dialog, H is asked to successively click on their best guess. At each click, the interface gives H feedback on whether their guess is correct or not and this continues until H guesses the true secret image. In this way, H induces a partial ranking of the pool up to the secret image based on their mental model of I^c from the dialog.



Figure 4.1: GuessWhich Interface: A user asks a question to ALICE in each round and ALICE responds with an answer. The user then selects an appropriate image which they think is the secret image after each round of conversation. At the end of the dialog, user successively clicks on their best guesses until they correctly identify the secret image.

4.2 Pool Selection

When creating a pool of images, our aim is to ensure that the game is challenging and engaging, and not too easy or too hard. Thus, we construct each pool of images \mathbb{I} in two steps – first, we choose the secret image I^c , and then sample similar images as distractors for I^c . Fig. 4.1 shows a screenshot of our game interface including a sample image pool and chat.

Secret Image Selection. VisDial v0.5 is constructed on 68k COCO images which contain complex everyday scenes with 80 object categories. ABOT is trained and validated on VisDial v0.5 *train* and *val* splits respectively. As the images for both these splits come from COCO-train, we sample secret images and pools from COCO-validation to avoid overlap.

To select representative secret images and diverse image pools, we do the following. For each image in the COCO validation set, we extract the penultimate layer (‘fc7’) activations

of a standard deep convolutional neural network (VGG-19 from [25]). For each of the 80 categories, we average the embedding vector of all images containing that category. We then pick those images closest to the mean embeddings, yielding 80 candidates.

Generating Distractor Images. The distractor images are designed to be semantically similar to the secret image I^c . For each candidate secret image, we created 3 concentric hyper-spheres as euclidean balls (of radii increasing in arithmetic progression) centered on the candidate secret image in fc7 embedding space, and sampled images from each sphere in a fixed proportion to generate a pool corresponding to the secret image. The radius of the largest sphere was varied and manually validated to ensure pool difficulty. The sampling proportion can be varied to generate pools of varying difficulty. Of the 80 candidate pools, we picked 10 that were of medium difficulty based on manual inspection.

4.3 Data Collection and Player Reward Structure

We use AMT to solicit human players for our game. Each Human Intelligence Task (HIT) consists of 10 games (each game corresponds to one pool) and we find that overall 76.7% of users who started a HIT completed it *i.e.* played all 10 games. We note that incomplete game data was discarded and does not contribute to the analysis presented in subsequent chapters.

We published HITs until 28 games with both $ALICE_{SL}$ and $ALICE_{RL}$ were completed. This results in a total of 560 games split between the agents, with each game consisting of 9 rounds of dialog and 10 rounds of guessing. Workers are paid a base pay of \$5 per HIT (\sim \$10/hour).

To incentivize workers to try their best at guessing the secret image, workers are paid a two-part bonus – (1) based on the number of times their best guess matched the true secret image after each round (up to \$1 per HIT), and (2) based on the rank of the true secret image in their final sorting at the end of dialog (up to \$2 per HIT).

This final ranking explicitly captures the workers’ mental model of the secret image

(unlike the per-round, best-guess estimates), and is closer to the overall purpose of the game (identifying the secret image at the end of the dialog). As such, this final sorting is given a higher potential bonus.

4.4 Evaluation

Since the game is structured as a retrieval task, we evaluate the human-AI collaborative performance using standard retrieval metrics. Note that the successive selection of images by H at the end of the dialog tells us the rank of the true secret image in a sorting of the image pool based on H’s mental model. For example, if H makes 4 guesses before correctly selecting the secret image, then H’s mental model ranked the secret image 5th within the pool.

To evaluate human-AI collaboration, we use the following metrics: (1) Mean Rank (MR), which is the mean rank of the secret image (*i.e.* number of guesses it takes to identify the secret image). Lower values indicate better performance. (2) Mean Reciprocal Rank (MRR), which is the mean of the reciprocal of the rank of the secret image. MRR penalizes differences in lower ranks (e.g., between 1 and 2) greater than those in higher ranks (e.g., between 19 and 20). Higher values indicate better performance.

At the end of each round, H makes their best guess of the secret image. To get a coarse estimate of the rank of the secret image in each round, we sort the image pool based on distance in fc7 embedding space from H’s best guess. This can be used to assess accuracy of H’s mental model of the secret image after each round of dialog (e.g., Fig. 6.1b).

to the client browser via websockets such that whenever an inference job is completed, a websocket polls the AI response and delivers it to the human in real-time. We store and fetch data efficiently to and from a PostgreSQL database. Fig. 5.1 shows a schematic diagram of the backend architecture. Our complete backend infrastructure and code is available at <https://bit.ly/2KUCpBr> for others to easily make use of our human-AI game interface.

CHAPTER 6

RESULTS

6.1 Alice_{SL} vs. Alice_{RL}

We compare the performance of the two agents ALICE_{SL} and ALICE_{RL} in the GuessWhich game. These bots are competitive with state-of-the-art visual dialog agents with respect to emulating human responses and generating visually discriminative responses in AI-AI dialog. [2] evaluate these agents against strong baselines and report AI-AI team results that are significantly better than chance on a pool of $\sim 10\text{k}$ images (rank ~ 1000 for SL, rank ~ 500 for RL). In addition to evaluating them in the context of human-AI teams we also report QBOT-ALICE team performances for reference.

Table 6.1: Performance of Human-ALICE teams with ALICE_{SL} and ALICE_{RL} measured by MR (lower is better) and MRR (higher is better). Error bars are 95% CIs from 1000 bootstrap samples. Unlike (Das et al., 2017b), we find no significant difference between ALICE_{SL} and ALICE_{RL} .

Team	MR	MRR
Human- ALICE_{SL}	6.86 ± 0.53	0.27 ± 0.03
Human- ALICE_{RL}	7.19 ± 0.55	0.25 ± 0.03

In Table 6.1, we compare the performances of human- ALICE_{SL} and human- ALICE_{RL} teams according to Mean Rank (MR) and Mean Reciprocal Rank (MRR) of the secret image based on the guesses H makes at the end of dialog. We observe that at the end of each game (9 rounds of dialog), human subjects correctly guessed the secret image on their 6.86th attempt (Mean Rank) when ALICE_{SL} was their teammate. With ALICE_{RL} as their teammate, the average number of guesses required was 7.19. We also observe that ALICE_{RL} outperforms ALICE_{SL} on the MRR metric. On both metrics, however, the differences are within the stan-

dard error margins (reported in the table) and not statistically significant. As we collected additional data, the error margins became smaller but the means also became closer. This interesting finding stands in stark contrast to the results reported by [2], where ALICE_{RL} was found to be significantly more accurate than ALICE_{SL} when evaluated in an AI-AI team. Our results suggest that the improvements of RL over SL (in AI-AI teams) do not seem to translate to when the agents are paired with a human in a similar setting.

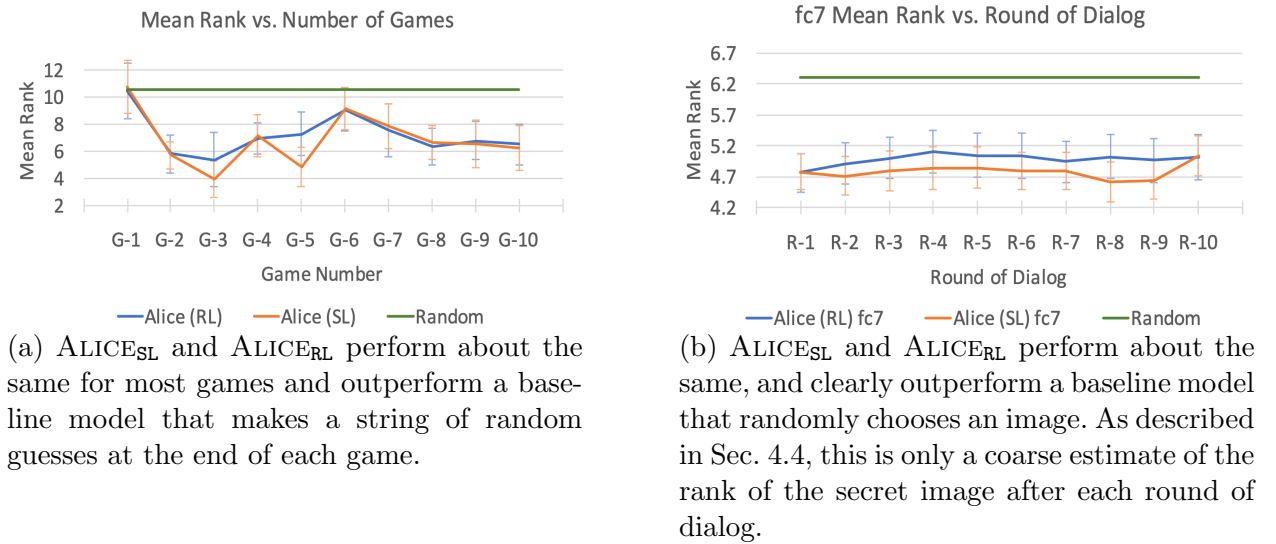


Figure 6.1: Mean rank (MR) of secret image across (a) number of games and (b) rounds of dialog. Lower is better. Error bars are 95% confidence intervals from 1000 bootstrap samples.

MR with varying number of games. In Fig. 6.1a, we plot the mean rank (MR) of the secret image across different games. We see that the human-ALICE team performs about the same for both ALICE_{SL} and ALICE_{RL} except Game 5, where ALICE_{SL} seems to marginally outperform ALICE_{RL} . We compare the performance of these teams against a baseline model that makes a string of random guesses at the end of the game. The human-ALICE teams outperform this random baseline with a relative improvement of about 25%.

AI-Alice teams versus human-Alice teams. In Table 6.2, we compare team performances by pairing three kinds of questioners – human, QBOT (SL) and QBOT (RL) with ALICE_{SL} and ALICE_{RL} (6 teams in total) to gain insights about how the questioner and ALICE influence team performances. Interestingly, we observe that AI-ALICE teams outperform

Table 6.2: Performance of Human-ALICE and QBOT-ALICE teams measured by MR (lower is better). We observe that AI-AI teams outperform human-AI teams.

Team	Alice _{SL}	Alice _{RL}
Human	6.9	7.2
QBOT (SL)	5.6	5.3
QBOT (RL)	4.7	4.7

human-ALICE teams. On average, a QBOT (SL)-ALICE_{SL} team takes about 5.6 guesses to arrive at the correct secret image (as opposed to 6.86 guesses for a human-ALICE_{SL} team). Similarly, a QBOT (RL)-ALICE_{RL} team takes 4.7 guesses as opposed to a human-ALICE_{RL} team which takes 7.19 guesses. When we compare AI-AI teams (see Row 2 and 3) under different settings, we observe that teams having QBOT (RL) as the questioner outperform those with QBOT (SL). Qualitatively, we found that QBOT (SL) tends to ask repeating questions in a dialog and that questions from QBOT (RL) tend to be more visually grounded compared to QBOT (SL). Also, note that among the four teams ALICE does not seem to affect performance across SL and RL.

Since we observe that QBOT (RL) tends to be a better questioner on average compared to QBOT (SL), as future work, it will be interesting to explore a setting where we evaluate QBOT via a similar game with the human playing the role of answerer in a QBOT-human team.

MR with varying rounds of dialog. Fig. 6.1b shows a coarse estimate of the mean rank of the secret image across rounds of a dialog, averaged across games and workers. As explained in Sec. 4.4 of Chapter 4, image ranks are computed via distance in embedding space from the guessed image (and hence, are only an estimate). We see that the human-ALICE team performs about the same for both ALICE_{SL} and ALICE_{RL} across rounds of dialog in a game. When compared with a baseline agent that makes random guesses after every round of dialog, the human-ALICE team clearly performs better.

Statistical tests. Observe that on both the metrics (MR and MRR), the differences between performances of ALICE_{SL} and ALICE_{RL} are within error margins. Since both standard error and bootstrap based 95% confidence intervals overlap significantly, we ran further statistical tests. We find no significant difference between the mean ranks of ALICE_{SL} and ALICE_{RL} under a Mann-Whitney U test ($p = 0.44$).

6.2 Human perception of AI teammate

At the end of each HIT, we asked workers for feedback on ALICE. Specifically, we asked workers to rate ALICE on a 5-point scale (where 1=Strongly disagree, 5=Strongly agree), along 6 dimensions. As shown in Fig. 6.2, ALICE was rated on – how accurate they thought it was (accuracy), how consistent its answers were with its previous answers (consistency), how well it understood the secret image (image understanding), how detailed its answers were (detail), how well it seemed to understand their questions (question understanding) and how fluent its answers were (fluency).




Figure 6.2: Worker ratings for ALICE_{SL} and ALICE_{RL} on 6 metrics. Higher is better. Error bars are 95% confidence intervals from 1000 bootstrap samples. Humans perceive no significant differences between ALICE_{SL} and ALICE_{RL} across the 6 feedback metrics.

We see in Fig. 6.2 that humans perceive both ALICE_{SL} and ALICE_{RL} as comparable in

terms of all metrics. The small differences in perception are not statistically significant.

6.3 Questioning Strategies

Fig. 6.3 shows the distribution of questions that human subjects ask ALICE in GuessWhich. Akin to the format of the human-human GuessWhat game, we observe that binary (yes/no) questions are overwhelmingly the most common question type, for instance, “Is there/the/he ...?” (region shaded yellow in the figure), “Are there ...?” (region shaded red), etc. The next most frequent question is “What color ...?”. These questions may be those that help the human discriminate the secret image the best. It could also be that humans are attempting to play to the perceived strengths of ALICE. As people play multiple games with ALICE, it is possible that they discover ALICE’s strengths and learn to ask questions that play to its strengths. Another common question type is counting questions, such as “How many ...?”. Interestingly, some workers adopt the strategy of querying ALICE with a single word (e.g., nouns such as “people”, “pictures”, etc.) or a phrase (e.g., “no people”, “any cars”, etc.). This strategy, while minimizing human effort, does not appear to change ALICE’s performance. Fig. 6.4 shows a game played by two different subjects.







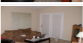











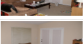












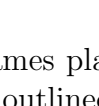
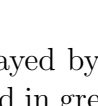
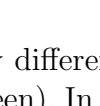
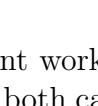


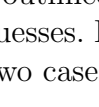
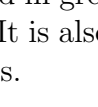
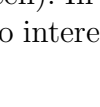
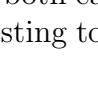


















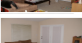
Dialog 1 (Human-Alice _{SL})		Pool				Dialog 2 (Human-Alice _{RL})	
	A man sitting on a couch with a laptop					A man sitting on a couch with a laptop	
	What color is the couch? Brown					How many people are in the picture? 1 person	
	Is there a television? Yes					What color is the man's shirt? it is black	
	Do you see any lamps? Yes					What color is the mans pants? He has black pants	
	Are there any windows? Yes					What color is the laptop? It is black	
	What color is the floor? Brown					What color are the walls? They are beige	
	What color are the walls? Beige					How big is the room? It looks pretty large	
	Are there any doors? No					What color is the floor? It is brown	
	Do you see any pictures on the walls? No					Is there a tv in the room? No I don't see TV	
	Are there curtains at the window? Yes					Is there a coffee table in the room? I can't tell	
Rank: 3						Rank: 3	

Figure 6.4: We contrast two games played by different workers with ALICE_{SL} and ALICE_{RL} on the same pool (secret image outlined in green). In both cases, the workers are able to find the secret image within three guesses. It is also interesting to note how the answers provided by ALICE are different in the two cases.

CHAPTER 7

CHALLENGES

There exist several challenges that are unique to human computation in the context of evaluating human-AI teams, for instance, making our games engaging while still ensuring fair and accurate evaluation. In this chapter, we briefly discuss some of the challenges we faced and our solutions to them.

7.1 Knowledge Leak

It has been shown that work division in crowdsourcing tasks follows a Pareto principle [28], as a small fraction of workers usually complete a majority of the work. In the context of evaluating an AI based on performance of a human-AI team, this poses a challenge.

Recently, [29] showed that human subjects can predict the responses of an AI more accurately with higher familiarity with the AI. That is, a human’s knowledge gained from familiarity with their AI teammate, can bias the performance of the human-AI team – knowledge from previous tasks might leak to later tasks. To prevent a biased evaluation of team performance due to human subjects who have differing familiarity with ALICE, every person only plays a fixed number of games (10) with ALICE. Thus, a human subject can only accept one task on AMT, which involves playing 10 games. The downside to this is that our ability to conduct a fair evaluation of an AI in an interactive, game-like setting is constrained by the number of unique workers who accept our tasks.

7.2 Engagement vs. Fairness

In order to improve user-engagement while playing our games, we offer subjects performance-based incentives that are tied to the success of the human-AI team. There is one potential

issue with this however. Owing to the inherent complexity of the visual dialog task, ALICE tends to be inaccurate at times. This increases both the difficulty and unpredictability of the game, as it tends to be more accurate for certain types of questions compared to others. We observe that this often leads to unsuccessful game-plays, sometimes due to errors accumulating from successive incorrect responses from ALICE to questions from the human. In a few other cases, the human is misled by ALICE by a single wrong answer or by the seed caption that tends to be inaccurate at times. While we would like to keep subjects engaged in the game to the best extent possible by providing performance-based incentives, issuing a performance bonus that depends on both the human and ALICE (who is imperfect), can be dissatisfying. To be fair to the subjects performing the task while still rewarding good performance, we split our overall budget for each HIT into a suitable fraction between the base pay (majority), and the performance bonus.

CHAPTER 8

CONCLUSION

In contrast to the common practice of measuring AI progress in isolation, this thesis proposes benchmarking AI agents via interactive downstream tasks (cooperative games) performed by human-AI teams. In particular, we evaluate visual conversational agents in the context of human-AI teams. We design a cooperative game – GuessWhich – that involves a human engaging in a dialog with an answerer-bot (ALICE) to identify a secret image known to ALICE but unknown to the human from a pool of images. At the end of the dialog, the human is asked to pick out the secret image from the image pool by making successive guesses.

Specifically, we studied how well progress in AI-AI (QBOT-ABOT) teams measured in isolation translates to human-AI team performance at GuessWhich when a human is paired with ALICE (ABOT). We find that $ALICE_{RL}$ (fine-tuned with reinforcement learning) that has been found to be more accurate in AI literature than its supervised learning counterpart when evaluated via a questioner bot (QBOT)-ALICE team, is not more accurate when evaluated via a human-ALICE team. This suggests that there is a disconnect between benchmarking of AI in isolation versus in the context of human-AI interaction. This disconnect implies that in order to effectively leverage the progress in AI for eventual downstream human-AI collaborative settings, it is essential to evaluate AI in the context of human-AI teams in addition to measuring progress on isolated benchmarks.

Additionally, we described the game structure and the backend architecture and discuss the unique computation and infrastructure challenges that arise when designing such live interactive settings on AMT relative to static human-labeling tasks. Our code and infrastructure are publicly available at <https://bit.ly/2KUCpBr> so as to enable others to make use or build on top of our cooperative setting.

Appendices

APPENDIX A

PUBLICATIONS

- **Evaluating Visual Conversational Agents via Cooperative Human-AI Games**

Prithvijit Chattopadhyay*, Deshraj Yadav*, Viraj Prabhu, Arjun Chandrasekaran, Abhishek Das, Stefan Lee, Dhruv Batra, Devi Parikh

5th AAAI Conference on Human Computation and Crowdsourcing (HCOMP 2017)

- **Do explanations make VQA models more predictable to a human?**

Arjun Chandrasekaran*, Viraj Prabhu*, Deshraj Yadav*, Prithvijit Chattopadhyay*, Devi Parikh

Conference on Empirical Methods in Natural Language Processing (EMNLP 2018)

- **Choose Your Neuron: Incorporating Domain Knowledge Through Neuron-Importance**

Ramprasaath R. Selvaraju*, Prithvijit Chattopadhyay*, Mohamed Elhoseiny, Dhruv Batra, Devi Parikh, Stefan Lee.

15th European Conference on Computer Vision (ECCV 2018)

- **It Takes Two to Tango: Towards Theory of AI’s Mind**

Arjun Chandrasekaran*, Deshraj Yadav*, Prithvijit Chattopadhyay*, Viraj Prabhu*, Devi Parikh

Chalearn Looking at People Workshop, IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2017)

- **Counting Everyday Objects in Everyday Scenes**

Prithvijit Chattopadhyay*, Ramakrishna Vedantam*, Ramprasaath R. Selvaraju, Dhruv Batra, Devi Parikh

IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2017)

REFERENCES

- [1] A. Das, S. Kottur, K. Gupta, A. Singh, D. Yadav, J. M. Moura, D. Parikh, and D. Batra, “Visual Dialog,” in *CVPR*, 2017.
- [2] A. Das, S. Kottur, J. M. Moura, S. Lee, and D. Batra, “Learning Cooperative Visual Dialog Agents with Deep Reinforcement Learning,” in *ICCV*, 2017.
- [3] H. de Vries, F. Strub, S. Chandar, O. Pietquin, H. Larochelle, and A. Courville, “Guess-What?! visual object discovery through multi-modal dialogue,” in *CVPR*, 2017.
- [4] F. Strub, H. de Vries, J. Mary, B. Piot, A. C. Courville, and O. Pietquin, “End-to-end optimization of goal-driven and visually grounded dialogue systems,” *arXiv preprint arXiv:1703.05423*, 2017.
- [5] L. Von Ahn and L. Dabbish, “Designing games with a purpose,” *Communications of the ACM*, vol. 51, no. 8, pp. 58–67, 2008.
- [6] L. von Ahn and L. Dabbish, “Labeling images with a computer game,” in *CHI*, 2004.
- [7] E. L. Law, L. Von Ahn, R. B. Dannenberg, and M. Crawford, “Tagatune: A game for music and sound annotation,” in *ISMIR*, vol. 3, 2007, p. 2.
- [8] P. Michelucci, “Handbook of human computation,” in *Springer*, 2013.
- [9] S. Jain and D. C. Parkes, “A game-theoretic analysis of the esp game,” *ACM Trans. Econ. Comput.*, vol. 1, no. 1, 3:1–3:35, Jan. 2013.
- [10] M. Krause and J. Smeddinck, “Human computation games: A survey,” in *Signal Processing Conference, 2011 19th European*, IEEE, 2011, pp. 754–758.
- [11] L. Von Ahn, R. Liu, and M. Blum, “Peekaboom: A game for locating objects in images,” in *Proceedings of the SIGCHI conference on Human Factors in computing systems*, ACM, 2006, pp. 55–64.
- [12] E. Law and L. Von Ahn, “Input-agreement: A new mechanism for collecting data using human computation games,” in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM, 2009, pp. 1197–1206.
- [13] S. Kazemzadeh, V. Ordonez, M. Matten, and T. L. Berg, “ReferItGame: Referring to Objects in Photographs of Natural Scenes,” in *EMNLP*, 2014.

- [14] N. Diakopoulos, K. Luther, and I. Essa, “Audio puzzler: Piecing together time-stamped speech transcripts with a puzzle game,” in *Proceedings of the 16th ACM international conference on Multimedia*, ACM, 2008, pp. 865–868.
- [15] H. Aras, M. Krause, A. Haller, and R. Malaka, “Webpardy: Harvesting qa by hc,” in *Proceedings of the ACM SIGKDD Workshop on Human Computation*, ACM, 2010, pp. 49–52.
- [16] J. Chamberlain, M. Poesio, and U. Kruschwitz, “Phrase detectives: A web-based collaborative annotation game,” in *Proceedings of the International Conference on Semantic Systems (I-Semantics’ 08)*, 2008, pp. 42–49.
- [17] T. Paek, “Empirical methods for evaluating dialog systems,” in *Proceedings of the workshop on Evaluation for Language and Dialogue Systems-Volume 9*, 2001.
- [18] C.-W. Liu, R. Lowe, I. V. Serban, M. Noseworthy, L. Charlin, and J. Pineau, “How NOT To Evaluate Your Dialogue System: An Empirical Study of Unsupervised Evaluation Metrics for Dialogue Response Generation,” in *EMNLP*, 2016.
- [19] J. Li, W. Monroe, A. Ritter, M. Galley, J. Gao, and D. Jurafsky, “Deep Reinforcement Learning for Dialogue Generation,” in *EMNLP*, 2016.
- [20] B. Grosz, “What question would turing pose today?” *AI Magazine*, 2012.
- [21] I. V. Serban, A. Sordoni, Y. Bengio, A. Courville, and J. Pineau, “Building End-To-End Dialogue Systems Using Generative Hierarchical Neural Network Models,” in *AAAI*, 2016.
- [22] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollr, and C. L. Zitnick, “Microsoft COCO: Common Objects in Context,” in *ECCV*, 2014.
- [23] A. Karpathy, *NeuralTalk2*, github.com/karpathy/neuraltalk2, 2016.
- [24] O. Vinyals, A. Toshev, S. Bengio, and D. Erhan, “Show and tell: A neural image caption generator,” in *CVPR*, 2015.
- [25] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” in *ICLR*, 2015.
- [26] AWS, *Amazon*, aws.amazon.com/ec2/, 2017.
- [27] RabbitMQ, *RabbitMQ*, rabbitmq.com, 2017.
- [28] G. Little, *How many turkers are there.(dec 2009)*, 2009.

- [29] A. Chandrasekaran, D. Yadav, P. Chattopadhyay, V. Prabhu, and D. Parikh, “It Takes Two to Tango: Towards Theory of AI’s Mind,” *arXiv preprint arXiv:1704.00717*, 2017.