

PUBLIC SELF-CONSCIOUSNESS FOR ENDOWING CONSISTENT PERSONA TO DIALOGUE AGENTS

Hyunwoo Kim Byeongchang Kim Gunhee Kim

Department of Computer Science and Engineering
Seoul National University, Seoul, Korea

{hyunw.kim, byeongchang.kim}@vision.snu.ac.kr gunhee@snu.ac.kr

ABSTRACT

Although consistency has been a long-standing issue in dialogue agents, we show best-performing persona-conditioned generative models still suffer from high insensitivity to contradiction. Current approaches for improving consistency rely on supervised external models and labels which are demanding. Inspired by social cognition and pragmatics, we model *public self-consciousness* in dialogue agents through an imaginary listener to improve consistency. Our approach, based on the Rational Speech Acts framework (Frank & Goodman, 2012), attempts to maintain consistency in an unsupervised manner requiring neither additional annotations nor pretrained external models. We further extend the framework by learning the distractor supply for the first time. Experimental results show that our approach effectively reduces contradiction and improves consistency on Dialogue NLI (Welleck et al., 2019) and PersonaChat (Zhang et al., 2018).

1 INTRODUCTION




Figure 1: Illustration of the consistency issue. While a literal dialogue agent (S_0) fails to deliver a consistent persona to the interlocutor, our self-conscious agent (S_1) does so, by modeling an imaginary listener. Icons are designed by Nhor Phai and Vincent Le Moign.

In the study of dialogue agents, *consistency* has been a long-standing issue (Li et al., 2016; Zhang et al., 2018; Welleck et al., 2019). To resolve this, research has been conducted to endow *personas* to dialogue agents (Li et al., 2016; Zhang et al., 2018). Recently, Welleck et al. (2019); Li et al. (2019) exploit Natural Language Inference (NLI) annotations for ranking consistent utterances or training to discourage inconsistent utterances.

In spite of such recent significant progress, there is much room for improvement. First, we observe that even best performing persona-based generative models (Tselousov & Golovanov, 2019; Wolf et al., 2019) are highly insensitive to contradictory utterances, and thus fail to deliver consistent persona to the interlocutor (Figure 1). Second, the NLI method (Welleck et al., 2019; Dziri et al., 2019; Song et al., 2019) and unlikelihood training (Li et al., 2019) require direct supervision under additional NLI annotations on the target dataset, which could be highly demanding.

In this work, we step back from supervision and ponder: *how do humans maintain consistency?* Humans encounter personal and social pressures for consistency everyday (Schlenker, 1975). Inconsistency is recognized irrational, and may evoke reactions of confusion and even anger, which could further lead to social punishment (Cialdini, 1993). However, we do not ask others whether we sound consistent or not. We ask ourselves, by predicting how we will be perceived by others. *Public self-consciousness* is this awareness of the self as a social object that can be observed and evaluated

by others (Fenigstein et al., 1975). Major findings in cognitive science argue that predicting other’s reactions is essential in social interactions; humans rely on abstract models of others (Gopnik & Wellman, 1992), and simulate their reactions by imagination (Hassabis et al., 2013). Thus, our behavior is affected by role-taking and imagination (Schlenker & Weigold, 1990). Such self-regulation of our behavior rely on a range of different mechanisms. Our work focuses on modeling the public self-consciousness mechanism through an imaginary listener, and show that it can help dialogue agents improve consistency.

Modeling a listener has been one of the main topics in pragmatics. We extend this long line of work in cognitive science by making use of the Rational Speech Acts (RSA) framework (Frank & Goodman, 2012), which has shown promising results in a number of NLP tasks (Andreas & Klein, 2016; Mao et al., 2016; Vedantam et al., 2017; Cohn-Gordon et al., 2018; Fried et al., 2017; 2018; Cohn-Gordon & Goodman, 2019; Shen et al., 2019; Zarri   & Schlangen, 2019). However, its usage was limited to improving informativeness and its application to the dialogue domain also remains understudied. The RSA framework operates In this work, we explore how the RSA framework can be adopted in dialogue agents to alleviate the *consistency* problem. We further extend the framework by learning the supply of distractors, which are negative samples of the given target. We also propose a different update for the listener’s prior.

The objective of this work is to propose a self-conscious dialogue agent that alleviates the consistency problem in an unsupervised manner. We take inspiration from social cognition and pragmatics to build an agent who imagines the listener’s reaction and then incorporates it to her utterance.

2 DATASETS & ANALYSIS OF INSENSITIVITY

PersonaChat Dialogue Dataset. Zhang et al. (2018) release the *PersonaChat*, which is a chitchat dataset involving two interlocutors. Each given a persona profile in sentences, they are asked to play the role while getting to know each other. This was the task of ConvAI2 competition (Dinan et al., 2019) at NeurIPS 2018.

Dialogue NLI Evaluation Set. Based on PersonaChat, Welleck et al. (2019) introduce the *Dialogue NLI* dataset to address the consistency issue. They collect entailing and contradictory utterances to the given persona, and release an evaluation set comprised of dialogues each with 31 utterance candidates: 10 entailing, 10 neutral, and 10 contradictory utterances with 1 ground-truth (GT) utterance. The task is to rank the appropriate candidates higher than inappropriate ones.

Insensitivity to Contradictory Utterances. From quantitative evaluation, we reveal existing generative models for dialogues are highly insensitive to contradictory utterances. On the Dialogue NLI evaluation set, we run two recent models (Wolf et al., 2019; Tselousov & Golovanov, 2019) that achieve best performance on the PersonaChat. We report four ranking metrics following Welleck et al. (2019): Hits@1, Entail@1, Neutral@1 and Contradict@1. Each is the proportion of the GT, entailing, neutral and contradictory utterances in the top-1 candidate chosen by the model, respectively. The models rank the candidates by perplexity scores. The first row in Table 1 shows that both models select contradictory candidates (54.1, 46.5) much more often than GT (8.5, 11.1). Contradict@1 is even higher than the sum of Hits@1 and Entail@1 (32.9, 37.5).

3 APPROACH

To resolve such insensitivity to contradiction, we introduce a self-conscious dialogue agent that keeps the consistency of every token generation by reflecting the imaginary listener’s reactions. Given that modeling the interactions between a listener and a speaker is a main topic in pragmatics, we make use of the RSA framework (Frank & Goodman, 2012). It treats language use as a recursive process where probabilistic speaker and listener reason about each other’s intentions in Bayesian fashion. To apply the RSA framework to sequence generation, we extend the incremental approach proposed by Cohn-Gordon et al. (2018). To generate an utterance, the agent computes the distribution of token u_t at timestep t as follows.

Base Speaker S_0 . We assume persona i is given to the base speaker S_0^t , along with the dialogue history h and partial utterance $u_{<t}$, as shown in Figure 2. S_0^t returns a distribution over the next token at timestep t : $S_0^t(u_t|i, h, u_{<t})$. Any conditional language model can be used as a base speaker.

Imaginary Listener L_0 . While the base speaker generates each token one at a time, the imaginary listener reasons about the speaker’s persona. The imaginary listener L_0^t is the posterior distribution of the speaker’s persona in terms of the base speaker and the world prior $p_t(i)$ over personas:

$$L_0^t(i|h, u_{\leq t}, p_t) \propto \frac{S_0^t(u_t|i, h, u_{<t})^\beta * p_t(i)}{\sum_{i' \in \mathcal{I}} S_0^t(u_t|i', h, u_{<t})^\beta * p_t(i')}. \quad (1)$$

where $\beta (\leq 1)$ on S_0^t is the listener rationality coefficient that controls the amount of information from the current timestep compared to the cumulative prior $p_t(i)$. L_0 returns a probability distribution over the personas in world \mathcal{I} , which is a finite set of personas including the given persona i and distractor personas. We decide the world \mathcal{I} per dialogue instance through learning, which is described below.

Self-Conscious Speaker S_1 . With S_0^t and L_0^t , the self-conscious speaker S_1^t is defined as

$$S_1^t(u_t|i, h, u_{<t}) \propto L_0^t(i|h, u_{\leq t}, p_t)^\alpha * S_0^t(u_t|i, h, u_{<t}), \quad (2)$$

where α is the speaker rationality coefficient that determines how much the likelihood is considered. By taking the listener’s distribution into account, the speaker is now self-conscious about what persona she sounds like. Especially, the agent seeks to be perceived as the given persona i rather than other persona i' . The likelihood of each token being identified as persona i acts as a bonus added to the S_0^t ’s token scores. Hence, tokens consistent to the given persona are preferred to others. The token with the highest score is added to the $u_{<t}$, giving us the next input $u_{<t+1}$ for the speaker.

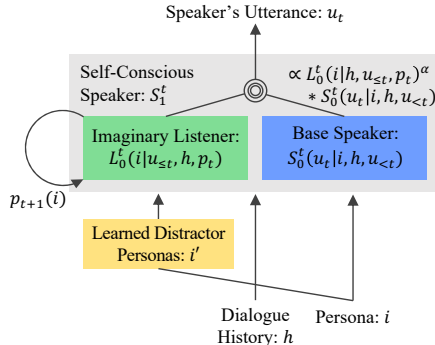


Figure 2: The self-conscious agent S_1 consists of a base speaker S_0 and an imaginary listener L_0 . It recursively generates next token u_t at every time t .

Updating the world prior with L_0 . Starting from a uniform distribution as the initial prior $p_0(i)$, we update the prior $p_{t+1}(i)$ according to S_1 ’s output u_t at every time step: $p_{t+1}(i) = L_0^t(i|h, u_{\leq t}, p_t)$. Hence, $p_t(i)$ represents the cumulative state of the partial utterance up to t . Reportedly, the prior update with $L_1 \propto S_0^t(u_t|i, h, u_{<t}) * L_0^t(i|h, u_{\leq t}, p_t)$ makes little difference compared to a uniform prior (Cohn-Gordon et al., 2018). We find that updating the prior with L_0 can alleviate this issue.

Learning to Provide Distractors. In previous works of RSA, the distractors in world \mathcal{I} are supplied manually. However, we find that this is rather impractical and the performance varies largely according to distractors. We thus propose to learn the distractor supply, especially based on the life-long memory (LLM) network (Kaiser et al., 2017). The number of possible conversational contexts per dialogue instance can be infinite as similar semantics can be uttered in many different ways. LLM can efficiently memorize and retrieve distractor personas for each dialogue context by clustering similar ones with associated persona. In order to train LLM, we find the best distractor persona per training dialogue that helps the self-conscious agent represent the GT utterance (*i.e.* distractor showing the lowest perplexity for GT). We regard this best distractor as the distractor label of the dialogue instance and perform supervised learning. At inference, given a test example, we obtain a query by encoding the dialogue context and the given persona using BERT (Devlin et al., 2019). With this query, we find k nearest keys from the memory, and use their values (*i.e.* persona indices) as the distractor personas.

4 EXPERIMENTS

Base Speakers. We experiment on GPT-based two winning models, LostInConv (Tselousov & Golovanov, 2019) and Transfer-T (Wolf et al., 2019), as base speakers (S_0) for our self-conscious agents (S_1). We improve these baselines by granting them the sense of self-consciousness.

Quantitative Results. Table 1 reports the performance of models on Dialogue NLI evaluation set and PersonaChat. In Dialogue NLI, our self-conscious S_1 significantly reduces Contradict@1 score and increases the Entail@1 along with Hits@1 of S_0 . Since each entailing candidate shares the same

Table 1: Comparison of our approach (S_1) with base speakers (S_0) on the Dialogue NLI evaluation set (Welleck et al., 2019) and PersonaChat (Zhang et al., 2018). +DM is the *Distractor Memory*. H@1, E@1, C@1, PPL denotes Hits@1, Entail@1, Contradict@1 and perplexity, respectively. C is a metric for dialogue consistency evaluated by a pretrained NLI model (Madotto et al., 2019).

Dialogue NLI		LostInConv			Transfer-T		
Model	H@1 ↑	E@1 ↑	C@1 ↓	H@1 ↑	E@1 ↑	C@1 ↓	
S_0	8.5	24.4	54.1	11.1	26.4	46.5	
S_1	11.4	40.6	30.8	16.4	38.8	28.8	
S_1 +DM	12.4	47.1	24.5	18.6	43.9	18.4	

PersonaChat		LostInConv				Transfer-T			
Model	H@1 ↑	F1 ↑	PPL ↓	C ↑	H@1 ↑	F1 ↑	PPL ↓	C ↑	
S_0	19.4	21.1	18.6	0.41	16.7	19.2	17.8	0.84	
S_1	21.2	20.5	23.1	0.50	19.2	19.5	22.6	0.98	
S_1 +DM	21.6	20.6	23.3	0.50	19.2	19.6	22.5	0.99	

annotated triple as the GT utterance, Entail@1 is a lenient version of Hits@1 (Welleck et al., 2019). Our *Distractor Memory* significantly improves the vanilla S_1 models across all metrics.

In PersonaChat, our model S_1 outperforms all other generative dialogue agents in terms of consistency related metrics, *i.e.* Hits@1 and C score. Since the posterior update of our self-conscious agent revises the distribution learned by the base speaker, the increase in perplexity is natural. For Transfer-T, our approach also improves the F1 score. The *Distractor Memory* further improves in both consistency metrics and accuracy metrics like F1 score and perplexity.

Human Evaluation. We sample 250 test examples from Transfer-T, each of which is rated by three human judges in terms of (*c*)*onsistency* and (*e*)*ngagingness*. In the test, we show the model’s given persona, dialogue context, and model’s generated utterance. To evaluate consistency, we follow Madotto et al. (2019) and ask judges to assign 1, 0, -1 to the utterance for consistency, neutrality, and contradiction, respectively. Following See et al. (2019), we evaluate the engagingness of utterances in a 4-point scale, where higher scores are better. Human judges rate our self-conscious agent S_1 (*c*: 0.61 (0.02), *e*: 2.55 (0.03)) as more consistent and engaging than the base agent S_0 (*c*: 0.53 (0.02), *e*: 2.48 (0.03)). Numbers in parentheses are the standard errors.

World Prior Update. Our experiments confirm that updating the world prior with L_1^t makes no difference to the performance compared with using a uniform distribution; this was first reported in Cohn-Gordon et al. (2018). However, our approach with L_0^t makes significant difference, as shown in Figure 3. The pragmatic listener $L_1^t \propto S_0^t(u_t|i, h, u_{<t}) * L_0^t(i|h, u_{\leq t}, p_t)$ reflects the *current* S_0^t twice per timestep; in L_0^t and in itself. Hence, the world prior updated with L_1^t becomes more of an instantaneous prior than a cumulative one. On the other hand, L_0^t moderately combines the information from both S_0^t and $p_t(i)$, preserving better cumulative information.

5 CONCLUSION

This work introduced how we can model public self-consciousness to endow consistent persona to dialogue agents. We proposed an unsupervised method using the RSA framework (Frank & Goodman, 2012). We extend the framework by proposing a learning method for distractor provision and a different update for the listener’s world prior. Our self-conscious agents outperformed other agents on Dialogue NLI and PersonaChat, without additional labels and pretrained external models.

ACKNOWLEDGEMENTS

This work is supported by Brain Research Program through the NRF of Korea (2017M3C7A1047860) and Creative-Pioneering Researchers Program through Seoul National University. Gunhee Kim is the corresponding author.




Figure 3: Performance variation of the self-conscious agents according to β and different updates for world prior $p_t(i)$.

REFERENCES

- Jacob Andreas and Dan Klein. Reasoning about Pragmatics with Neural Listeners and Speakers. In *EMNLP*, 2016.
- Robert B Cialdini. *Influence: The Psychology of Persuasion*. Morrow New York, 1993.
- Reuben Cohn-Gordon and Noah Goodman. Lost in Machine Translation: A Method to Reduce Meaning Loss. In *NAACL-HLT*, 2019.
- Reuben Cohn-Gordon, Noah Goodman, and Christopher Potts. Pragmatically Informative Image Captioning With Character-level Inference. In *NAACL-HLT*, 2018.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL-HLT*, 2019.
- Emily Dinan, Varvara Logacheva, Valentin Malykh, Alexander Miller, Kurt Shuster, Jack Urbanek, Douwe Kiela, Arthur Szlam, Iulian Serban, Ryan Lowe, et al. The Second Conversational Intelligence Challenge (ConvAI2). *arXiv:1902.00098*, 2019.
- Nouha Dziri, Ehsan Kamaloo, Kory W Mathewson, and Osmar Zaiane. Evaluating Coherence in Dialogue Systems Using Entailment. In *NAACL-HLT*, 2019.
- Allan Fenigstein, Michael F Scheier, and Arnold H Buss. Public and Private Self-Consciousness: Assessment and Theory. *Journal of Consulting and Clinical Psychology*, 43(4):522, 1975.
- Michael C Frank and Noah D Goodman. Predicting Pragmatic Reasoning in Language Games. *Science*, 336(6084):998–998, 2012.
- Daniel Fried, Jacob Andreas, and Dan Klein. Unified Pragmatic Models for Generating and Following Instructions. In *NAACL-HLT*, 2017.
- Daniel Fried, Ronghang Hu, Volkan Cirik, Anna Rohrbach, Jacob Andreas, Louis-Philippe Morency, Taylor Berg-Kirkpatrick, Kate Saenko, Dan Klein, and Trevor Darrell. Speaker-follower models for vision-and-language navigation. In *NeurIPS*, 2018.
- Alison Gopnik and Henry M Wellman. Why the Child’s Theory of Mind Really is a Theory. *Mind & Language*, 7(1-2):145–171, 1992.
- Demis Hassabis, R Nathan Spreng, Andrei A Rusu, Clifford A Robbins, Raymond A Mar, and Daniel L Schacter. Imagine All the People: How the Brain Creates and Uses Personality Models to Predict Behavior. *Cerebral Cortex*, 24(8):1979–1987, 2013.
- Łukasz Kaiser, Ofir Nachum, Aurko Roy, and Samy Bengio. Learning to Remember Rare Events. In *ICLR*, 2017.
- Jiwei Li, Michel Galley, Chris Brockett, Georgios P Spithourakis, Jianfeng Gao, and Bill Dolan. A Persona-Based Neural Conversation Model. In *ACL*, 2016.
- Margaret Li, Stephen Roller, Ilia Kulikov, Sean Welleck, Y-Lan Boureau, Kyunghyun Cho, and Jason Weston. Don’t Say That! Making Inconsistent Dialogue Unlikely with Unlikelihood Training. *arXiv:1911.03860*, 2019.
- Andrea Madotto, Zhaojiang Lin, Chien-Sheng Wu, and Pascale Fung. Personalizing Dialogue Agents via Meta-Learning. In *ACL*, 2019.
- Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, Alan L Yuille, and Kevin Murphy. Generation and Comprehension of Unambiguous Object Descriptions. In *CVPR*, 2016.
- Barry R Schlenker. Self-Presentation: Managing the Impression of Consistency When Reality Interferes With Self-Enhancement. *Journal of Personality and Social Psychology*, 32(6):1030, 1975.
- Barry R Schlenker and Michael F Weigold. Self-Consciousness and Self-Presentation: Being Autonomous Versus Appearing Autonomous. *Journal of Personality and Social Psychology*, 59(4): 820, 1990.

- Abigail See, Stephen Roller, Douwe Kiela, and Jason Weston. What Makes a Good Conversation? How Controllable Attributes Affect Human Judgments. In *NAACL-HLT*, 2019.
- Sheng Shen, Daniel Fried, Jacob Andreas, and Dan Klein. Pragmatically Informative Text Generation. In *NAACL-HLT*, 2019.
- Haoyu Song, Wei-Nan Zhang, Jingwen Hu, and Tiu Liu. Generating Persona Consistent Dialogues by Exploiting Natural Language Inference. *arXiv:1911.05889*, 2019.
- Alexander Tselousov and Sergey Golovanov. Lost In Conversation. https://github.com/atseousov/transformer_chatbot, 2019.
- Ramakrishna Vedantam, Samy Bengio, Kevin Murphy, Devi Parikh, and Gal Chechik. Context-Aware Captions from Context-Agnostic Supervision. In *CVPR*, 2017.
- Sean Welleck, Jason Weston, Arthur Szlam, and Kyunghyun Cho. Dialogue Natural Language Inference. In *ACL*, 2019.
- Thomas Wolf, Victor Sanh, Julien Chaumond, and Clement Delangue. TransferTransfo: A Transfer Learning Approach for Neural Network based Conversational Agents. *arXiv:1901.08149*, 2019.
- Sina Zarrieß and David Schlangen. Know What You Don’t Know: Modeling a Pragmatic Speaker that Refers to Objects of Unknown Categories. In *ACL*, 2019.
- Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. Personalizing Dialogue Agents: I Have a Dog, Do You Have Pets Too? In *ACL*, 2018.