

Gunrock: A Social Bot for Complex and Engaging Long Conversations

Dian Yu, Michelle Cohn, Yi Mang Yang, Chun-Yen Chen, Weiming Wen
Jiaping Zhang, Mingyang Zhou, Kevin Jesse, Austin Chau, Antara Bhowmick
Shreenath Iyer, Girithaja Sreenivasulu, Sam Davidson, Ashwin Bhandare, Zhou Yu

University of California, Davis

{dianyu, mdcohn, yimyang, abtchen, wmwen, jpzhang, minzhou, krjesse}@ucdavis.edu
{amchou, abhowmick, shriyer, s.girithaja, ssdavidson, asbhandar, joyu}@ucdavis.edu

Abstract

Gunrock is the winner of the 2018 Amazon Alexa Prize, as evaluated by coherence and engagement from both real users and Amazon-selected expert conversationalists. We focus on understanding complex sentences and having in-depth conversations in open domains. In this paper, we introduce some innovative system designs and related validation analysis. Overall, we found that users produce longer sentences to Gunrock, which are directly related to users' engagement (e.g., ratings, number of turns). Additionally, users' backstory queries about Gunrock are positively correlated to user satisfaction. Finally, we found dialog flows that interleave facts and personal opinions and stories lead to better user satisfaction.

1 Introduction

Amazon Alexa Prize (Ram et al., 2018) provides a platform to collect real human-machine conversation data and evaluate performance on speech-based social conversational systems. Our system, Gunrock (Chen et al., 2018)¹ addresses several limitations of prior chatbots (Vinyals and Le, 2015; Zhang et al., 2018; Fang et al., 2018) including inconsistency and difficulty in complex sentence understanding (e.g., long utterances) and provides several contributions: First, Gunrock's multi-step language understanding modules enable the system to provide more useful information to the dialog manager, including a novel dialog act scheme. Additionally, the natural language understanding (NLU) module can handle more complex sentences, including those with coreference. Second, Gunrock interleaves actions to elicit users' opinions and provide responses to create an in-depth, engaging conversation; while a related

strategy to interleave task- and non-task functions in chatbots has been proposed (Rudnicky, 2019), no chatbots to our knowledge have employed a fact/opinion interleaving strategy. Finally, we use an extensive persona database to provide coherent profile information, a critical challenge in building social chatbots (Zhang et al., 2018). Compared to previous systems (Fang et al., 2018), Gunrock generates more balanced conversations between human and machine by encouraging and understanding more human inputs (see Table 1 for an example).

User_1:	let's chat
Gunrock_1:	Hi, this is an Alexa Prize Socialbot. It's nice to talk to you again, George. I think last time we were talking about movies. Would you like to talk about it now?
User_2:	sure [SPL] that would be great
Gunrock_2:	By the way, I love to discover new movies. I recently watched Crazy Rich Asians because someone told me it was hilarious. What movie would you recommend to me?
User_3:	ha [SPL] it's a tough question [SPL] i don't think i have a good one to recommend [SPL] wait [SPL] i think that <i>stars born</i> is good
Gunrock_3:	Yes I remember watching <i>a star is born</i> . Hmm. If you had to rate this movie from 1 to 10, what would it be?
User_4:	ten
Gunrock_4:	A perfect 10! You have to tell me more. What was so exceptional? When I watched it, the music design really stood out to me.
User_5:	when i watched it the music was amazing [SPL] and bradley cooper was super talented in the movie [SPL] i really like him
Gunrock_5:	You've got me wanting to watch the movie again. Ouu, Maybe you will find this interesting. Bradley Cooper claimed that his dream was to be a director before he got into acting and this was his directorial debut ...

Table 1: An example conversation between Gunrock and a human user (User). Text in red represents ASR errors while text in green represents ASR correction. [SPL] represents a split of the sentences from a sequence of words from ASR.

¹<https://gunrock-ucdavis.weebly.com/2018-system-demonstration.html>

2 System Architecture

Figure 1 provides an overview of Gunrock’s architecture. We extend the Amazon Conversational Bot Toolkit (CoBot) (Khatrı et al., 2018) which is a flexible event-driven framework. CoBot provides ASR results and natural language processing pipelines through the Alexa Skills Kit (ASK) (Kumar et al., 2017). Gunrock corrects ASR according to the context (§2.1) and creates a natural language understanding (NLU) (§2.2) module where multiple components analyze the user utterances. A dialog manager (DM) (§2.3) uses features from NLU to select topic dialog modules and defines an individual dialog flow. Each dialog module leverages several knowledge bases (§2.4). Then a natural language generation (NLG) (§2.5) module generates a corresponding response. Finally, we markup the synthesized responses and return to the users through text to speech (TTS) (§2.6). While we provide an overview of the system in the following sections, for detailed system implementation details, please see the technical report (Chen et al., 2018).

2.1 Automatic Speech Recognition

Gunrock receives ASR results with the raw text and timestep information for each word in the sequence (without case information and punctuation). Keywords, especially named entities such as movie names, are prone to generate ASR errors without contextual information, but are essential for NLU and NLG. Therefore, Gunrock uses domain knowledge to correct these errors by comparing noun phrases to a knowledge base (e.g. a list of the most popular movies names) based on their phonetic information. We extract the primary and secondary code using The Double Metaphone Search Algorithm (Philips, 2000) for noun phrases (extracted by noun trunks) and the selected knowledge base, and suggest a potential fix by code matching. An example can be seen in User_3 and Gunrock_3 in Table 1.

2.2 Natural Language Understanding

Gunrock is designed to engage users in deeper conversation; accordingly, a user utterance can consist of multiple units with complete semantic meanings. We first split the corrected raw ASR text into sentences by inserting break tokens. An example is shown in User_3 in Table 1. Meanwhile, we mask named entities before segmenta-

tion so that a named entity will not be segmented into multiple parts and an utterance with a complete meaning is maintained (e.g., “i like the movie a star is born”). We also leverage timestep information to filter out false positive corrections. After segmentation, our coreference implementation leverages entity knowledge (such as person versus event) and replaces nouns with their actual reference by entity ranking. We implement coreference resolution on entities both within segments in a single turn as well as across multiple turns. For instance, “him” in the last segment in User_5 is replaced with “bradley cooper” in Table 1. Next, we use a constituency parser to generate noun phrases from each modified segment. Within the sequence pipeline to generate complete segments, Gunrock detects (1) topic, (2) named entities, and (3) sentiment using ASK in parallel. The NLU module uses knowledge graphs including Google Knowledge Graph² to call for a detailed description of each noun phrase for understanding.

In order to extract the intent for each segment, we designed MIDAS, a human-machine dialog act scheme with 23 tags and implemented a multi-label dialog act classification model using contextual information (Yu and Yu, 2019). Next, the NLU components analyzed on each segment in a user utterance are sent to the DM and NLG module for state tracking and generation, respectively.

2.3 Dialog Manager

We implemented a hierarchical dialog manager, consisting of a high level and low level DMs. The former leverages NLU outputs for each segment and selects the most important segment for the system as the central element using heuristics. For example, “i just finished *reading* harry potter;” triggers Sub-DM: Books. Utilizing the central element and features extracted from NLU, input utterances are mapped onto 11 possible topic dialog modules (e.g., movies, books, animals, etc.), including a backup module, retrieval.

Low level dialog management is handled by the separate topic dialog modules, which use modular finite state transducers to execute various dialog segments processed by the NLU. Using topic-specific modules enables deeper conversations that maintain the context. We design dialog flows in each of the finite state machines, as well. Dia-

²<https://developers.google.com/knowledge-graph/>

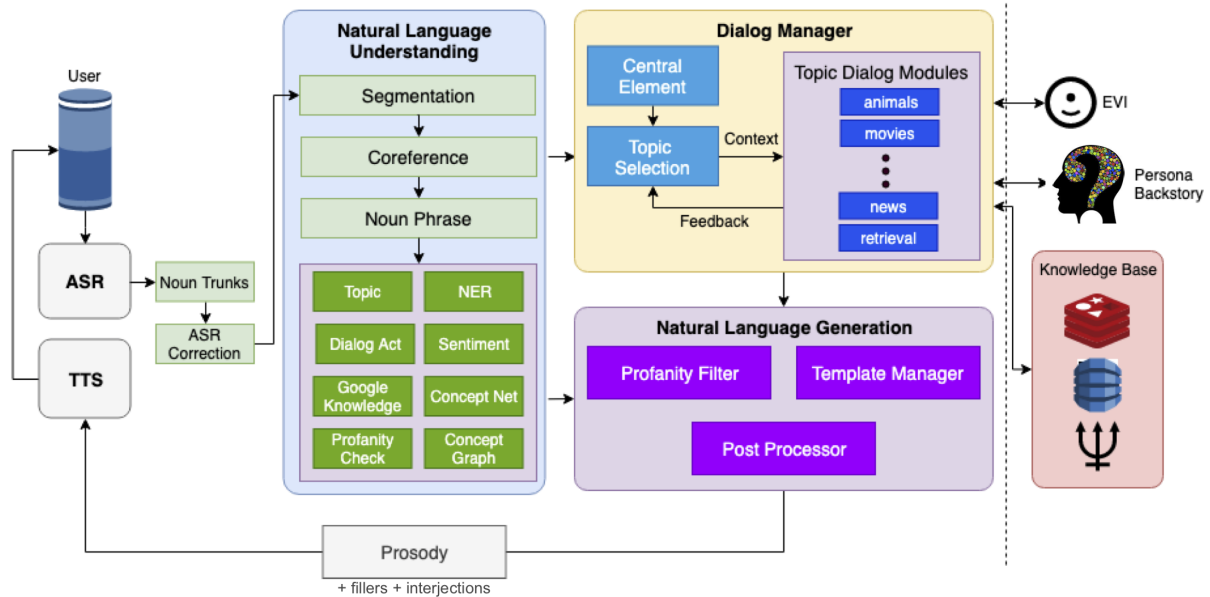


Figure 1: Gunrock system architecture

log flow is determined by rule-based transitions between a specified fixed set of dialog states. To ensure that our states and transitions are effective, we leverage large scale user data to find high probability responses and high priority responses to handle in different contexts. Meanwhile, dialog flow is customized to each user by tracking user attributes as dialog context. In addition, each dialog flow is adaptive to user responses to show acknowledgement and understanding (e.g., talking about pet ownership in the animal module). Based on the user responses, many dialog flow variations exist to provide a fresh experience each time. This reduces the feeling of dialogs being scripted and repetitive. Our dialog flows additionally interleave facts, opinions, experiences, and questions to make the conversation flexible and interesting.

In the meantime, we consider feedback signals such as “continue” and “stop” from the current topic dialog module, indicating whether it is able to respond to the following request in the dialog flow, in order to select the best response module. Additionally, in all modules we allow mixed-initiative interactions; users can trigger a new dialog module when they want to switch topics while in any state. For example, users can start a new conversation about movies from any other topic module.

2.4 Knowledge Databases

All topic dialog modules query knowledge bases to provide information to the user. To respond

to general factual questions, Gunrock queries the EVI factual database³, as well as other up-to-date scraped information appropriate for the submodule, such as news and current showing movies in a specific location from databases including IMDB. One contribution of Gunrock is the extensive Gunrock Persona Backstory database, consisting of over 1,000 responses to possible questions for Gunrock as well as reasoning for her responses for roughly 250 questions (see Table 2). We designed the system responses to elicit a consistent personality within and across modules, modeled as a female individual who is positive, outgoing, and is interested in science and technology.

User1:	do you like jeopardy?
Gunrock1:	Sure I like Jeopardy, especially when Watson competed.
User2:	why?
Gunrock2:	I’m so impressed with the capabilities of a supercomputer.

Table 2: Example interaction between Gunrock and a human user (User) querying Gunrock’s backstory.

2.5 Natural Language Generation

In order to avoid repetitive and non-specific responses commonly seen in dialog systems (Li et al., 2015), Gunrock uses a template manager to select from a handcrafted response templates based on the dialog state. One dialog state can map to multiple response templates with simi-

³<https://www.evi.com/>

lar semantic or functional content but differing surface forms. Among these response templates for the same dialog state, one is randomly selected without repetition to provide variety unless all have been exhausted. When a response template is selected, any slots are substituted with actual contents, including queried information for news and specific data for weather. For example, to ground a movie name due to ASR errors or multiple versions, one template is “Are you talking about {movie_title} released in {release_year} starring {actor_name} as {actor_role}?”. Module-specific templates were generated for each topic (e.g., animals), but some of the templates are generalizable across different modules (e.g., “Whats your favorite [movie | book | place to visit]?”)

In many cases, response templates corresponding to different dialog acts are dynamically composed to give the final response. For example, an appropriate acknowledgement for the users response can be combined with a predetermined follow-up question.

2.6 Text To Speech

After NLG, we adjust the TTS of the system to improve the expressiveness of the voice to convey that the system is an engaged and active participant in the conversation. We use a rule-based system to systematically add interjections, specifically Alexa Speechcons, and fillers to approximate human-like cognitive-emotional expression (Tokuhisa and Terashima, 2006). For more on the framework and analysis of the TTS modifications, see (Cohn et al., 2019).

3 Analysis

From January 5, 2019 to March 5, 2019, we collected conversational data for Gunrock. During this time, no other code updates occurred. We analyzed conversations for Gunrock with at least 3 user turns to avoid conversations triggered by accident. Overall, this resulted in a total of 34,432 user conversations. Together, these users gave Gunrock an average rating of 3.65 (median: 4.0), which was elicited at the end of the conversation (“On a scale from 1 to 5 stars, how do you feel about talking to this socialbot again?”). Users engaged with Gunrock for an average of 20.92 overall turns (median 13.0), with an average of 6.98 words per utterance, and had an average conversation time of 7.33 minutes (median: 2.87 min.). We conducted three

principal analyses: users’ response depth (§3.1), backstory queries (§3.2), and interleaving of personal and factual responses (§3.3).

3.1 Response Depth: Mean Word Count

Two unique features of Gunrock are its ability to dissect longer, complex sentences, and its methods to encourage users to be active conversationalists, elaborating on their responses. In prior work, even if users are able to drive the conversation, often bots use simple yes/no questions to control the conversational flow to improve understanding; as a result, users are more passive interlocutors in the conversation. We aimed to improve user engagement by designing the conversation to have more open-ended opinion/personal questions, and show that the system can understand the users’ complex utterances (See §2.2 for details on NLU). Accordingly, we ask if users’ speech behavior will reflect Gunrock’s technical capability and conversational strategy, producing longer sentences.

We assessed the degree of conversational depth by measuring users’ mean word count. Prior work has found that an increase in word count has been linked to improved user engagement (e.g., in a social dialog system (Yu, 2016)). For each user conversation, we extracted the overall rating, the number of turns of the interaction, and the user’s per-utterance word count (averaged across all utterances). We modeled the relationship between word count and the two metrics of user engagement (overall rating, mean number of turns) in separate linear regressions.

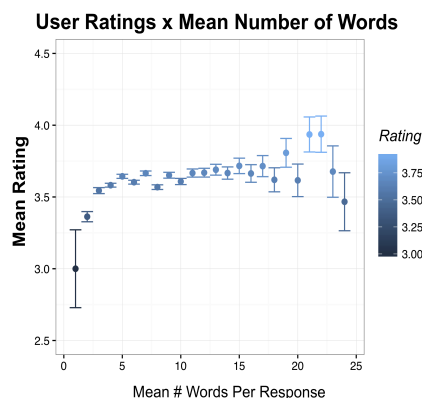


Figure 2: Mean user rating by mean number of words. Error bars show standard error.

Results showed that users who, on average, produced utterances with more words gave significantly higher ratings ($\beta=0.01$, $SE=0.002$, $t=4.79$,

$p < 0.001$)⁴(see Figure 2) and engaged with Gunrock for significantly greater number of turns ($\beta=1.85$, $SE=0.05$, $t=35.58$, $p < 0.001$) (see Figure 2). These results can be interpreted as evidence for Gunrock’s ability to handle complex sentences, where users are not constrained to simple responses to be understood and feel engaged in the conversation – and evidence that individuals are more satisfied with the conversation when they take a more active role, rather than the system dominating the dialog. On the other hand, another interpretation is that users who are more talkative may enjoy talking to the bot in general, and thus give higher ratings in tandem with higher average word counts.

3.2 Gunrock’s Backstory and Persona

We assessed the user’s interest in Gunrock by tagging instances where the user triggered Gunrock’s backstory (e.g., “What’s your favorite color?”). For users with at least one backstory question, we modeled overall (log) Rating with a linear regression by the (log) ‘Number of Backstory Questions Asked’ (log transformed due to the variables’ non-linear relationship). We hypothesized that users who show greater curiosity about Gunrock will display higher overall ratings for the conversation based on her responses. Overall, the number of times users queried Gunrock’s backstory was strongly related to the rating they gave at the end of the interaction ($\log:\beta=0.10$, $SE=0.002$, $t=58.4$, $p < 0.001$)(see Figure 3). This suggests that maintaining a consistent personality — and having enough responses to questions the users are interested in — may improve user satisfaction.

3.3 Interleaving Personal and Factual Information: Animal Module

Gunrock includes a specific topic module on animals, which includes a factual component where the system provides animal facts, as well as a more personalized component about pets. Our system is designed to engage users about animals in a more casual conversational style (Ventola, 1979), eliciting follow-up questions if the user indicates they have a pet; if we are able to extract the pet’s name, we refer to it in the conversation (e.g., “Oliver is a great name for a cat!”, “How long have you had Oliver?”). In cases where the user does not indi-

⁴ β = Unstandardized Beta Coefficient, SE = Standard Error, t = T value, p = P value

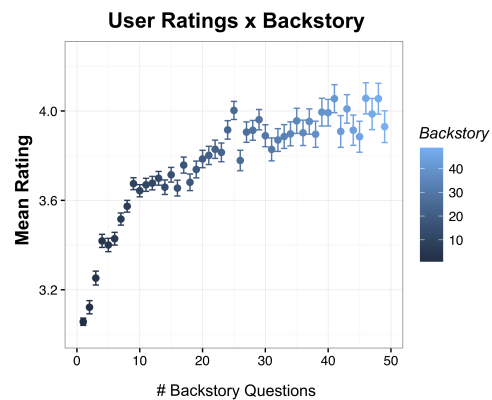


Figure 3: Mean user rating based on number of queries to Gunrock’s backstory. Error bars show standard error.

cate that they have a pet, the system solely provides animal facts. Therefore, the animal module can serve as a test of our interleaving strategy: we hypothesized that combining facts and personal questions — in this case about the user’s pet — would lead to greater user satisfaction overall.

We extracted conversations where Gunrock asked the user if they had ever had a pet and categorized responses as “Yes”, “No”, or “NA” (if users did not respond with an affirmative or negative response). We modeled user rating with a linear regression model, with predictor of “Has Pet” (2 levels: Yes, No). We found that users who talked to Gunrock about their pet showed significantly higher overall ratings of the conversation ($\beta=0.15$, $SE=0.06$, $t=2.53$, $p=0.016$) (see Figure 4). One interpretation is that interleaving factual information with more in-depth questions about their pet result in improved user experience. Yet, another interpretation is that pet owners may be more friendly and amenable to a socialbot; for example, prior research has linked differences in personality to pet ownership (Kidd and Kidds, 1980).

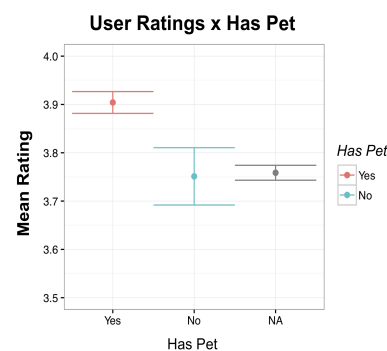


Figure 4: Mean user rating based ‘Has Pet’. Error bars show standard error.

4 Conclusion

Gunrock is a social chatbot that focuses on having long and engaging speech-based conversations with thousands of real users. Accordingly, our architecture employs specific modules to handle longer and complex utterances and encourages users to be more active in a conversation. Analysis shows that users' speech behavior reflects these capabilities. Longer sentences and more questions about Gunrock's backstory positively correlate with user experience. Additionally, we find evidence for interleaved dialog flow, where combining factual information with personal opinions and stories improve user satisfaction. Overall, this work has practical applications, in applying these design principles to other social chatbots, as well as theoretical implications, in terms of the nature of human-computer interaction (cf. 'Computers are Social Actors' (Nass et al., 1994)). Our results suggest that users are engaging with Gunrock in similar ways to other humans: in chitchat about general topics (e.g., animals, movies, etc.), taking interest in Gunrock's backstory and persona, and even producing more information about themselves in return.

Acknowledgments

We would like to acknowledge the help from Amazon in terms of financial and technical support.

References

- Chun-Yen Chen, Dian Yu, Weiming Wen, Yi Mang Yang, Jiaping Zhang, Mingyang Zhou, Kevin Jesse, Austin Chau, Antara Bhowmick, Shreenath Iyer, Giritha Sreenivasulu, Runxiang Cheng, Ashwin Bhandare, and Zhou Yu. 2018. [Gunrock: Building a human-like social bot by leveraging large scale real user data](#). In *2nd Proceedings of Alexa Prize*.
- Michelle Cohn, Chun-Yen Chen, and Zhou Yu. 2019. A large-scale user study of an alexa prize chatbot: Effect of TTS dynamism on perceived quality of social dialog. In *Proceedings of the 20th SIGdial Workshop on Discourse and Dialogue*.
- Hao Fang, Hao Cheng, Maarten Sap, Elizabeth Clark, Ari Holtzman, Yejin Choi, Noah A. Smith, and Mari Ostendorf. 2018. [Sounding board: A user-centric and content-driven social chatbot](#). *CoRR*, abs/1804.10202.
- Chandra Khatri, Behnam Hedayatnia, Anu Venkatesh, Jeff Nunn, Yi Pan, Qing Liu, Han Song, Anna Gottardi, Sanjeev Kwatra, Sanju Pancholi, Ming Cheng, Qinglang Chen, Lauren Stubel, Karthik Gopalakrishnan, Kate Bland, Raefer Gabriel, Arindam Mandal, Dilek Hakkani-Tür, Gene Hwang, Nate Michel, Eric King, and Rohit Prasad. 2018. [Advancing the state of the art in open domain dialog systems through the alexa prize](#). *CoRR*, abs/1812.10757.
- Aline H Kidd and Robert M Kidds. 1980. [Personality characteristics and preferences in pet ownership](#). *Psychological Reports*, 46(3):939–949.
- Anjishnu Kumar, Arpit Gupta, Julian Chan, Sam Tucker, Björn Hoffmeister, and Markus Dreyer. 2017. [Just ASK: building an architecture for extensible self-service spoken language understanding](#). *CoRR*, abs/1711.00549.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2015. [A diversity-promoting objective function for neural conversation models](#). *CoRR*, abs/1510.03055.
- Clifford Nass, Jonathan Steuer, and Ellen R Tauber. 1994. Computers are social actors. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 72–78. ACM.
- Lawrence Philips. 2000. [The double metaphone search algorithm](#). *C/C++ Users J.*, 18(6):38–43.
- Ashwin Ram, Rohit Prasad, Chandra Khatri, Anu Venkatesh, Raefer Gabriel, Qing Liu, Jeff Nunn, Behnam Hedayatnia, Ming Cheng, Ashish Nagar, Eric King, Kate Bland, Amanda Wartick, Yi Pan, Han Song, Sk Jayadevan, Gene Hwang, and Art Pettigrew. 2018. [Conversational AI: the science behind the alexa prize](#). *CoRR*, abs/1801.03604.
- Alexander I Rudnicky. 2019. C ha d: Chat-oriented dialog systems. In *Advanced Social Interaction with Agents*, pages 57–60. Springer.
- Ryoko Tokuhisa and Ryuta Terashima. 2006. [Relationship between utterances and "enthusiasm" in non-task-oriented conversational dialogue](#). In *Proceedings of the 7th SIGdial Workshop on Discourse and Dialogue*.
- Eija Ventola. 1979. The structure of casual conversation in english. *Journal of pragmatics*, 3(3-4):267–298.
- Oriol Vinyals and Quoc V. Le. 2015. [A neural conversational model](#). *CoRR*, abs/1506.05869.
- Dian Yu and Zhou Yu. 2019. [Midas: A dialog act annotation scheme for open domain human machine spoken conversations](#). *arXiv preprint arXiv:1908.10023*.
- Zhou Yu. 2016. *Situated Intelligent Interactive Systems*. Ph.D. thesis, Carnegie Mellon University.
- Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018. [Personalizing dialogue agents: I have a dog, do you have pets too?](#) *CoRR*, abs/1801.07243.