

Conversational Machine Comprehension: a Literature Review

Somil Gupta
CICS, UMass Amherst
somilgupta@umass.edu

Bhanu Pratap Singh Rawat
CICS, UMass Amherst
brawat@umass.edu

Hong Yu
CICS, UMass Lowell
Hong_Yu@uml.edu

Abstract

Conversational Machine Comprehension (CMC), a research track in conversational AI, expects the machine to understand an open-domain natural language text and thereafter engage in a multi-turn conversation to answer questions related to the text. While most of the research in Machine Reading Comprehension (MRC) revolves around single-turn question answering (QA), multi-turn CMC has recently gained prominence, thanks to the advancement in natural language understanding via neural language models such as BERT and the introduction of large-scale conversational datasets such as CoQA and QuAC. The rise in interest has, however, led to a flurry of concurrent publications, each with a different yet structurally similar modeling approach and an inconsistent view of the surrounding literature. With the volume of model submissions to conversational datasets increasing every year, there exists a need to consolidate the scattered knowledge in this domain to streamline future research. This literature review attempts at providing a holistic overview of CMC with an emphasis on the common trends across recently published models, specifically in their approach to tackling conversational history. The review synthesizes a generic framework for CMC models while highlighting the differences in recent approaches and intends to serve as a compendium of CMC for future researchers.

1 Introduction

Developing open-domain, intelligent dialog systems that can satisfactorily interact like humans, perform complex tasks and/or answer on a range of topics has been one of the most ambitious and difficult goals in Artificial Intelligence (AI). The study of such systems, called **Conversational AI** (ConvAI), is at the confluence of Natural language Processing (NLP), Information Retrieval (IR), and Machine Learning (ML), attracting significant research from both academia and industry. The recent developments in Deep Learning (DL) (Du and Black, 2019; Hatua et al., 2019) and Reinforcement Learning (RL) (Lipton et al., 2016; Peng et al., 2018) have further boosted research in the domain, making it one of the most sought after research topics in AI.

Based on the nature of problems, a ConvAI system is expected to solve three major research problems (Gao et al., 2018). *Question Answering* (QA) involves providing answers to user queries through conversation, using the knowledge drawn from various data sources like a snippet from a text, a collection of web documents, or an entire knowledge base. *Task completion* expects the conversational agent to accomplish task/s for the user, using the information acquired through conversation. Finally, *Social Chat* makes the agent emulate humans and converse seamlessly and appropriately with users, as in the Turing test (Saygin et al., 2000). Each of these fields has its own set of challenges to tackle.

Challenges in QA can vary depending on the source of knowledge, the answer extraction strategy employed, and the domain of the question. **Machine Reading Comprehension (MRC)** is one such challenge in QA, that requires the conversational QA (ConvQA) agent to understand a given open-domain text and thereafter answer question/s in conversation about it. These questions are often not paraphrased

This work is licensed under a Creative Commons Attribution 4.0 International License. License details: <http://creativecommons.org/licenses/by/4.0/>.

and may co-reference previous queries. The required answer may be a span of the given text or free-form. When the machine comprehension dialog involves multiple co-referenced questions such that a latter question may be a logical successor of the former, the challenge is termed as **Conversational Machine Comprehension (CMC)**. A lot of research in MRC revolves around single-turn QA, but multi-turn CMC also holds major relevance because humans seek information conversationally by asking follow-up questions for additional information based on what they have already learned. Still, the inherent complexity involved in dealing with text comprehension and reasoning over dialogs and context had kept CMC as a far-fetched goal. However, the recent success in achieving at-par-with-human performance on single-turn MRC models (Rajpurkar et al., 2018) due to the advancement in natural language understanding and modeling (Devlin et al., 2019; Lan et al., 2019; Liu et al., 2019a), and the introduction of large-scale conversational datasets CoQA (Reddy et al., 2019) and QuAC (Choi et al., 2018) have made information-seeking dialogs possible.

As a consequence, CMC has seen a significant surge in research in recent years. In less than 2 years since the introduction of these datasets, there have been 40 submissions¹ to CoQA leaderboard² and 22 submissions¹ to QuAC leaderboard³. Many of these models are unpublished, indicating active ongoing research on these datasets. Besides, the current state-of-the-art in QuAC lags behind human performance F1 benchmark by a margin of 6.7¹, suggesting significant scope for improvement. Almost simultaneously, there have been breakthroughs in NLP (Devlin et al., 2019; Radford et al., 2018; Liu et al., 2019a) which the researchers have tried to leverage in their upcoming models (Qu et al., 2019b; Yeh and Chen, 2019; Chen et al., 2019). Since many models are being published concurrently, there have been inconsistencies and/or overlap in their methodologies and research directions. This makes it difficult to compare different approaches against each other and weigh their pros and cons. This prevailing scenario has blurred the bigger picture and made it difficult for researchers to attend to novel research in this field. Moreover, there is no singular summarized view on CMC models, except the individual literature studies of these publications which can be highly localized and inconsistent with the global view. Thus, the current scenario motivates the need for organizing the scattered knowledge across these publications into a consolidated overview, so that future research in this field can be streamlined.

This literature review, therefore, provides a bird-eye overview of Conversational Machine Comprehension. We commence with an introduction to CMC, acquainting the reader with the challenges that make CMC unique, and the large-scale conversational datasets that spurred research in this field. *To develop a general understanding of the CMC approaches, we shift the focus from comprehending individual models to observing the common trends that mark these models, synthesizing a generic framework for a CMC model in the process.* We finally end our review with a discussion on the current trends and suggest advancements in the future.

2 Related Work

There have been several published literature reviews on MRC in recent years. Gao et al. (2018) provides an extensive review of Conversational AI with a detailed account of the neural approaches being employed in each of its dialog systems (QA, Task completion, and social chat). It briefly discusses the problem of CMC and its datasets but does not comment upon the recent advancements and prevalent approaches in this domain. Zhang et al. (2019) provides a summary of all the recent single-turn MRC datasets and approaches, however, it briefly discusses CoQA but does not touch upon any approaches for CMC. Qiu et al. (2019) summarizes the classic models of single-turn MRC with a focus on deriving a common architecture and suggesting improvements based on the analysis. CMC is mentioned as an emerging research direction in this survey. The latest review by Baradaran et al. (2020) provides an overview of MRC along with the statistical analysis of datasets and the various problems in this domain. It mentions CMC as an MRC challenge but does not provide any further details.

¹Recorded as of July 1, 2020. Please note that many of these submissions are either ensemble versions of single models, or hyper-parameter variants of their pre-published models or are simply unpublished. Therefore, unique published models' count is 13 in CoQA and 7 in QuAC.

²<https://stanfordnlp.github.io/coqa/>

³<http://quac.ai/>

This review differs from its predecessors as it focuses primarily on Conversational (multi-turn) Machine Comprehension which has not been detailed in the previous literature. CMC has its own set of challenges and an active research community around it. This calls for considering CMC as a separate research direction from single-turn MRC and review its rapid developments in terms of its general trends.

3 What is Conversational Machine Comprehension?

The task of CMC is defined as: *Given a passage P , the conversation history in the form of question-answer pairs $\{Q_1, A_1, Q_2, A_2, \dots, Q_{i-1}, A_{i-1}\}$ and a question Q_i , the model needs to predict the answer A_i .* The answer A_i can either be a text span (s_i, e_i) (Choi et al., 2018) or a free-form text $\{a_{i,1}, a_{i,2}, \dots, a_{i,j}\}$ with evidence R_i (Reddy et al., 2019). Single-turn MRC models cannot directly cater to CMC, as the latter is much more challenging to address. The major challenges being:

- The encoding module needs to encode not only P and A_i but also the conversational history.
- General observation about information-seeking dialog in humans suggests that the starting dialog-turns tend to focus on the beginning chunks of the passage and shift focus to the later chunks as the conversation progresses (Choi et al., 2018). The model is thus expected to capture these focal shifts during a conversation and reason pragmatically, instead of only matching lexically or via paraphrasing.
- Multi-turn conversations are generally incremental and co-referential. These conversational dialogs are either *drilling down* (the current question is a request for more information about the topic), *shifting topic* (the current question is not immediately relevant to something previously discussed), *returning topic* (the current question is asking about a topic again after it had previously been shifted away from), *clarification* of topic, or *definition* of an entity (Yatskar, 2019). The model should, therefore, be able to take context from history which may or may not be immediate.

4 Multi-Turn Conversational Datasets

The surge in CMC research is credited to the emergence of large-scale multi-turn conversational datasets: CoQA (Reddy et al., 2019) and QuAC (Choi et al., 2018).

4.1 CoQA

Conversational QA (CoQA) dataset consists of 126k questions sourced from 8k conversations.

- **Dataset preparation:** Conversations are prepared over passages collected across 7 different domains, each with its source dataset, such as news articles derived from CNN (Hermann et al., 2015). Amongst the 7 domains, two are used for out-of-domain evaluation (only for evaluation, not training), while the other five aid in-domain evaluation (both training and evaluation). The dialog is prepared in a two annotator setting with one questioning and another answering, both referring to the entire context.
- **Questions:** Questions are factoid but require sufficient co-referencing and pragmatic reasoning (Bell, 1999).
- **Answers:** Answers are free-form, with their corresponding rationale highlighted in the passage. However, Yatskar (2019) identified that the answers are slightly modified versions of the rationale, and therefore optimizing an extractive model to predict the answer span with maximum F1 overlap to the gold answer can achieve up to 97.8 F1.
- **Dialog features:** The dialogs mostly involve drilling-down for details (about 60% of all questions) but lack other dialog features like topic-shift, clarification, or definition.
- **Evaluation:** Macro-average F1 score of word overlap is used as an evaluation metric and is computed separately for in-domain and out-of-domain.

4.2 QuAC

Question Answering in Context (QuAC) contains 100K questions obtained from 14K information-seeking dialogs.

- **Dataset preparation:** Dialogs are prepared over sections from Wikipedia articles about people from different genres such as culture and wildlife. The dataset is prepared using an asymmetric setting, with a student exposed only to the title of the article and a summary while the teacher is exposed to the entire section of the article on which the dialog is to be based. The student, therefore, tries to seek information about the hidden questions based on the limited information it gets from the dialog, and the teacher answers by providing short excerpts from the section (or ‘No Answer’ if not possible).
- **Questions:** Questions are descriptive, highly-contextual, and open-ended due to the asymmetric nature of the dataset that prevents paraphrasing. They require sufficient co-referencing and pragmatic reasoning.
- **Dialog features:** Besides drilling down, dialogs switch to new topics more frequently than CoQA. The dataset though lacks definition or clarification dialogs.
- **Answers:** Answers are extractive and can be either Yes/No or ‘No Answer’. Besides extractive span, the response also includes additional signals called dialog acts like continuation (follow up, maybe follow up, or don’t follow up) and affirmation (yes, no, or neither), which provides additional useful dialog flow information to train on, as used by Qu et al. (2019b) and Ju et al. (2019). Further, an analysis of the answer token lengths in Table 1 shows that QuAC answers are longer, which can be attributed to its asymmetric nature thereby motivating the seeker to ask open-ended questions to gauge hidden text.
- **Evaluation:** Besides the macro-averaged F1 score on the entire set, QuAC also evaluates Human Equivalence Score (HEQ) to judge system performance relative to an average human, by finding the percentage of instances for which the system’s F1 matches or exceeds human F1. HEQ-Q and HEQ-D are thus HEQ scores with the instances as questions and dialogs respectively.

General dataset characteristics and an example from each of the datasets are provided in Appendix A.

5 Generic Framework of a CMC Model

Gao et al. (2018) defined the steps for performing reading comprehension in a typical neural MRC model as (1) *encoding* the questions and context into a set of embeddings in a neural space; (2) *reasoning* in the neural space to identify the answer vector and (3) *decoding* the answer vector into a natural language output. Huang et al. (2018a) adapted these steps in CMC by adding conversational history modeling. Qu et al. (2019c) proposed a ConvQA model with separate modules for history selection and modeling. Based on these prior works, we synthesize a generic framework for a CMC model. A typical CMC model is provided with context C , current question Q_i and the conversation history $H_i = [\{Q_k, A_k\}]_{k=1}^{i-1}$, and needs to generate an output set O_i . The CMC framework is provided in Fig. 1. There are *four* major components of the framework, based on their contribution to the overall CMC flow.

1. **History Selection module:** With complicated dialog behaviors like topic shift or topic return (Yatskar, 2019), simply selecting immediate turns may not work well. A history selection module, therefore, chooses a subset H_i^l of the history turns H_i based on a policy (dynamic or static) that is expected to be more helpful than the others. If the history selection module is based on a dynamic learned policy (e.g. Qu et al. (2019b)), then feedback from the other modules can guide its update.
2. **Encoder:** The lexical tokens of the context passage C , selected conversational turns H_i^l , and the current question Q_i need to be transformed into input embeddings for the reasoning module. Encoder facilitates this transition. The encoder steps may vary with every approach and reasoning inputs, at a high level, encoding involves transformation and combination of context-independent word embeddings called *lexical embeddings* such as GloVe (Pennington et al., 2014), *intra-sequence contextual*

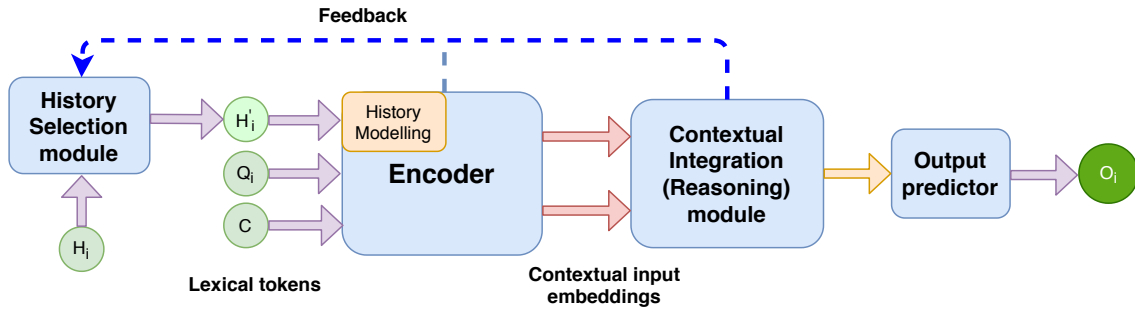


Figure 1: Generic framework of a CMC model. A typical CMC model would consist of (1) History selection module, that selects a subset H_i^l of conversational history H_i relevant to the current question Q_i ; (2) Encoder, that encodes the lexical tokens of context C , Q_i and H_i^l into input embeddings for contextual integration layer; (3) Reasoning module, that performs contextual integration of input embeddings into contextualized embeddings; and finally, (4) Output predictor, that predicts the output set O_i based on contextualized embeddings.

embeddings e.g. ELMo (Peters et al., 2018), BERT (Devlin et al., 2019) or RNN, *question-aware embeddings*, and *additional feature embeddings* like POS tags (Zhu et al., 2018), history embedding (Qu et al., 2019c) or conversation count. Conversational history H_i^l is generally integrated with this module into any or all of the contextual input embeddings. This process is called *History modeling* and is the most significant aspect of a CMC encoder.

3. **Contextual Integration layer:** Contextual information accumulated in the passage, query, and/or history embeddings individually must be fused to generate query-aware and/or history-aware contextualized output embeddings. This process may involve a single layer (*single-step reasoning*) or repetition across multiple layers (*multi-step reasoning*). Input for this module generally consists of two (or more) sequence sets for every history turn, or aggregated across all turns, which are then fused in each layer and often inter-weaved (Huang et al., 2018b) with attention.
4. **Output Predictor:** The model output may be in the form of a text span, signals like dialog acts (Choi et al., 2018) or a free-form (abstractive) answer (Reddy et al., 2019). Contextual embeddings generated by the reasoning module have all the latent information about the question, context passage, and conversational history. To get the token-level output, a fully-connected network followed by a softmax layer is generally used for per-token probability (abstractive) or start/end probability (extractive). Besides, a linear neural network may be used to find the aggregated result of the sequence.

6 Common Trends across CMC models

Instead of describing each CMC model separately, we categorized them under the approaches they employ in their components (section 5) or other model characteristics. This will help in developing a high-level understanding of the CMC models. A model-wise summary of the CMC models is provided in Appendix C.

6.1 Trends in History Selection

Almost all of the current CMC models select conversational history based on a heuristic of considering k immediate turns, often decided by performance such as BiDAF++ (Choi et al., 2018; Yatskar, 2019), SDNet (Zhu et al., 2018), BiDAF++ w/ 2-ctx (Ohsugi et al., 2019) use last two turns as including the third turn degrades performance. History Attention Mechanism (HAM) based model Qu et al. (2019b) uses a dynamic history selection policy by attending over contextualized representations of all the previous history turns at word-level or sequence-level and combining with current turn’s representation as shown in Fig. 3a.

6.2 Trends in History Modeling

How conversational history is integrated or used in the encoding process of contextual input embeddings can be used to classify CMC models. Different trends observed in this respect are described below. Some models may use a combination of these approaches.

1. **Appending selected history questions and/or answers** (in raw form or text span indices) **to the current question** before encoding. QA tokens across turns should be distinguishable or separated when appending. Models DrQA+PGNet (Reddy et al., 2019), SDNet (Zhu et al., 2018) and RoBERTa + AT + KD (Ju et al., 2019) append all history QA pairs separated by tokens like symbols $[Q]$ or $[A]$ such that new $Q_k^* = \{[Q], Q_1, [A], A_1, \dots, [Q], Q_{k-1}, [A], A_{k-1}, [Q], Q_k\}$. On the other hand, QuAC baseline model BiDAF++ w/ 2-ctx (Ohsugi et al., 2019) and GraphFlow (Chen et al., 2019) append only the history questions to the current question and encode relative dialog-turn number within each question embedding to differentiate. Choi et al. (2018) validate that this dialog-turn encoding strategy performs better in practice.
2. **Encoding context tokens with history answer marker embeddings (HAE)** before passing on for reasoning. These embeddings indicate if the context token is present in any conversational history answer or not, such as in BiDAF++ w/ 2-ctx (Choi et al., 2018), GraphFlow (Chen et al., 2019), BERT+HAE (Qu et al., 2019a) and HAM (Qu et al., 2019b). HAM encodes a dialog-turn encoded variant of HAE called *Positional HAE*. It maintains a lookup table of history embeddings for every relative position from the current conversation and embeds the corresponding embedding if the token is found in that history answer, e.g. for the current question q_k if a token is found in history answer a_{k-2} then Positional HAE embedding at index 2 is encoded, otherwise embedding at index 0 is encoded. This setting is illustrated in Fig. 3b.
3. **Integrating intermediate representations generated in the reasoning modules** of selected history conversation turns to grasp the deep latent semantics of the history, rather than acting on raw inputs. This approach is also called the *FLOW based approach*. The models that follow this approach are FlowQA (Huang et al., 2018a), FlowDelta (Yeh and Chen, 2019), and GraphFlow (Chen et al., 2019). GraphFlow encodes conversational histories into context graphs which are used by the reasoning module for contextual analysis.

For contextual encoding, most of the models utilize one of the two types of encoders: (a.) Bidirectional sequential language models such as BiDAF (Seo et al., 2017) or ELMo (Peters et al., 2018) (b.) Deep bidirectional transformer-based models such as BERT (Devlin et al., 2019) or RoBERTa (Liu et al., 2019b).

6.3 Trends in Contextual Reasoning

While every CMC model has its unique flavor in integrating encoded representations of the query, history, and text contextually, some recurrent themes in reasoning can still be drawn. It is important to note that some of these themes will reflect state-of-the-art techniques around their release, which may now be obsolete. However, having their knowledge would prevent the re-exploration of those ideas. Following are the commonly observed themes:

A. Attention-based Reasoning with Sequence Models

This was a common theme across MRC models until transformers (Vaswani et al., 2017) were introduced and got rid of sequence modeling. Consequently, initial baseline models were based on this approach.

CoQA baseline (Reddy et al., 2019) first involves DrQA (Chen et al., 2017), which performs BiLSTM based contextual integration over encoded tokens for extractive span, and later PGNet, that uses attention-based neural machine translation (Bahdanau et al., 2015) for abstractive answer reasoning.

QuAC baseline (Choi et al., 2018) combines self-attention with BiDAF (Seo et al., 2017) that performs reasoning via multi-layered bidirectional attention followed by multi-layered BiLSTM (BiDAF++).

SDNet (Zhu et al., 2018) applies both inter-attention and self-attention in multiple layers, interleaved with BiLSTM, to comprehend conversation context.

B. FLOW based approaches

Analogous to recurrent models which propagate contextual information through the sequence, FLOW is a sequence of latent representations that propagates reasoning in direction of the dialog progression by feeding intermediate latent representations, generated during reasoning in previous conversations, into contextual reasoning for the current question. This helps to leverage the reasoning effort of previous conversations as compared to using shallow history, such as directly appending history question-answers, where important contextual information in conversations may be lost due to the overwhelming input. There are *two* major flow-based approaches based on the propagated latent representation.

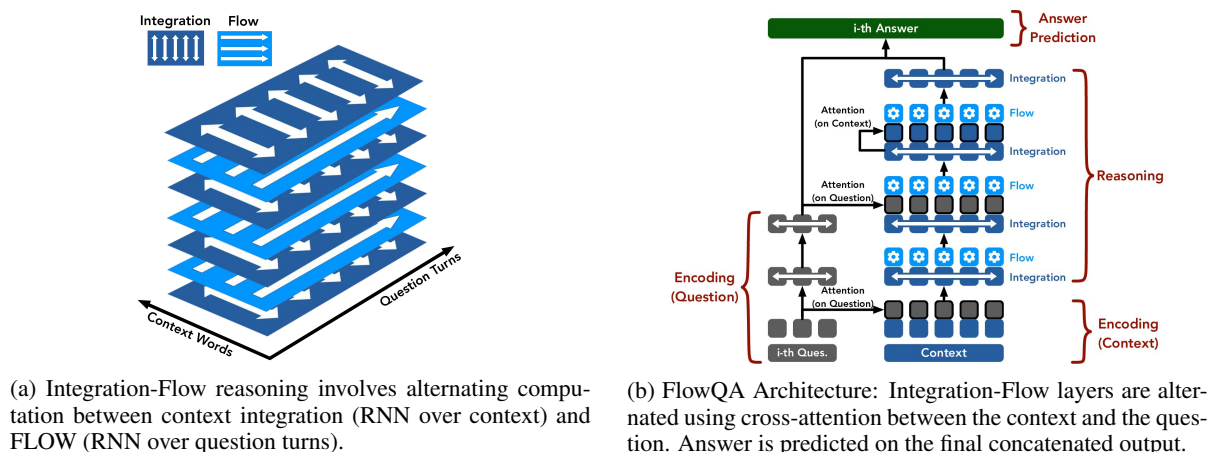


Figure 2: Illustration of the Integration-Flow based reasoning in FlowQA. Source : (Huang et al., 2018a)

1. **Integration-Flow (IF)**: This mechanism uses contextualized embeddings as the propagated latent representation. FlowQA (Huang et al., 2018a) which also introduced the idea of FLOW, involves sequential processing along context tokens in parallel to the question turns followed by sequential processing in direction of the question turns (Flow), in parallel to context tokens as illustrated in Fig. 2a. FlowQA employs multiple IF layers interleaved with self and cross attentions to reason over encoded embeddings (Fig. 2b). Recently released FlowDelta (Yeh and Chen, 2019) is an improvement on the IF approach that uses the similar FlowQA architecture and achieves better results. Instead of passing the latent representation directly, as in FlowQA, FlowDelta passes the information gain (the difference between the latent representation of previous 2 layers) with the intuition that information gain would allow the model to focus on more informative cues in context.
2. **Integration-GraphFlow (IG)**: GraphFlow (Chen et al., 2019) claims that the IF mechanism does not mimic human reasoning, as it first performs reasoning in parallel for each question, and then refines the reasoning results across different turns. Therefore, they use dynamically constructed, question-aware context graphs for each turn as the propagated latent representation. Processing through this flow (called GraphFlow) is facilitated by applying GNNs (Li et al., 2016) on the current context graph and previous context. To capture local interactions among consecutive words in context before feeding to a GNN, a BiLSTM is applied for contextual Integration. GraphFlow architecture alternates this mechanism with co-attention over the question and GNN output. This is illustrated in the figure provided in Appendix B.

C. Contextual Integration using Pre-trained Language Models

Large-scale pre-trained LMs such as BERT (Devlin et al., 2019), GPT (Radford et al., 2018) and RoBERTa (Liu et al., 2019b), have become the current state-of-the-art approaches for contextual reasoning in CMC models, with leaderboards of both datasets stacked with these models or their variants. The approach is based on the fine-tune BERT-based MRC modeling outlined by Devlin et al. (2019), in which question and context are packed together (with marker embeddings to distinguish) in an input sequence to BERT that outputs contextualized question-aware embeddings for each input token.

Using pre-trained models for reasoning is advantageous in two aspects: *Firstly*, it simplifies the architecture by fusing encoding and reasoning modules into a single module. *Secondly*, it provides a ready-to-tune architecture that abstracts out complex contextual interactions between query and context while providing sufficient flexibility to control interactivity via *augmentation of input embeddings* i.e. concatenation of special embeddings to input tokens that signal the model to incorporate a desirable characteristic in contextualization.

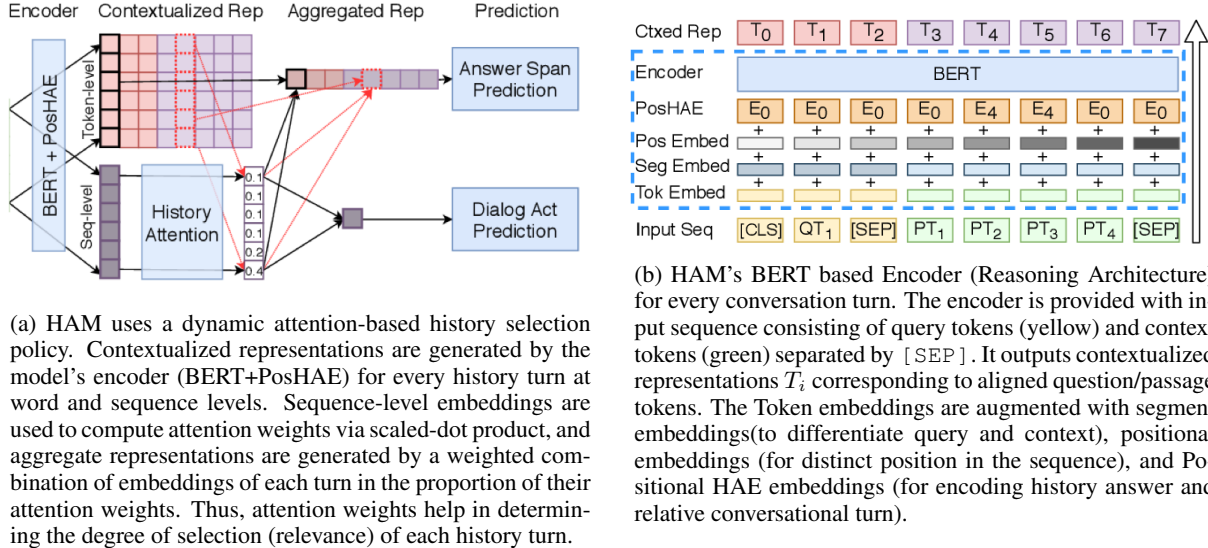


Figure 3: Illustration of (a) history selection module and (b) encoder/reasoning module of History Attention Mechanism (HAM) model (Qu et al., 2019b).

However, incorporating history into these models is a key challenge in this approach as most of the transformer models such as BERT only accept 2 segments ids in the input sequence. Based on recent research in CMC, two main trends in solving the history integration issue are discussed below:

1. **Modify the input embeddings for a single-turn MRC model** to incorporate history. This is done by either appending the entire conversation to the question, such as Ju et al. (2019) which uses RoBERTa (Liu et al., 2019a) as the base model and truncates query if it exceeds the limit, or add special embeddings to highlight conversational history for the model, such as HAE (Qu et al., 2019a) embeds history answer embeddings with each context token if it is present in any of the history turns (detailed in section 6.2). This approach does not effectively use the model to capture interactions between every dialog-turn and context.
2. **Use separate model for each conversational turn** to capture one-to-one interaction between history and context, and merge the per-turn contextualized embeddings into aggregated history-aware embeddings. Two models follow this trend. Ohsugi et al. (2019) uses BERT models to capture contextual interaction for every question (history and current) and answer ($2N+1$ sequences for N turns) and concatenates all sequences together. Finally, it runs Bi-GRU (Cho et al., 2014) over the aggregated sequence to capture inter-turn interactions before sending for prediction. On the other hand, HAM (Qu et al., 2019b) ignores the history questions and uses the current question as a query with positional History Answer Embeddings (section 6.2), thus generating one output sequence per conversation turn. Fig. 3b illustrates HAM encoder. The final sequence is generated using token-level soft-attention based aggregation across all per-turn contextualized sequences.

6.4 Trends in Training Methodology

Due to the multi-output nature of both CoQA and QuAC, multi-task training is quite common amongst CMC models, e.g. HAM (Qu et al., 2019b) uses multi-task learning over QuAC to also predict dialog prediction and continuation acts, while GraphFlow (Chen et al., 2019) uses multi-task learning over

CoQA to also predict question type. Besides, recently published (Ju et al., 2019) achieved state-of-the-art results using RoBERTa, by applying multiple training techniques over CoQA. These consist of rationale tagging multi-task learning (predict if the token exists in CoQA evidence), Adversarial Training (Goodfellow et al., 2015), and Knowledge Distillation (Furlanello et al., 2018).

7 Discussion

How does the research progress in CMC, a constrained setup, benefit the more into-the-wild domain of Conversational Search? As stated by Qu et al. (2019a), Conversational QA (and CMC) is a simplified setting of Conversational Search (ConvSearch), an information-seeking, “System Ask, User Respond” paradigm (Zhang et al., 2018b), that does not focus on asking proactively. CMC, specifically, tries to address the challenges of NLU, via contextual encoding, reasoning, and handling conversational history, via history selection and modeling. In that aspect, CMC is a concrete enough setting for IR researchers to understand the change of information needs and interactivity between conversational cycles.

Could Commonsense Reasoning improve CMC? Commonsense Reasoning (CR) is based on the set of background information or world knowledge that an individual is intended to know or assume, and may be missing from context. On the other hand, Pragmatic reasoning, which the current CMC models cater to, is based on the derivation of explicit and implicit meanings within the context. The current MRC systems are nearing human performance on most datasets, however, they still perform poorly on single-turn CR based questions (Zhang et al., 2018a). While there is recently increasing interest in CR in the single-turn MRC setting (Huang et al., 2019; Ostermann et al., 2018; Lin et al., 2017), CMC remains relatively untouched. This may probably be due to the lack of foreknowledge requiring unanswerable questions (e.g. in SQuAD 2.0 (Rajpurkar et al., 2018)) in current CMC datasets (Yatskar, 2019), suggesting a need for more complex CMC datasets that incorporate CR. However, humans annotators may often apply common-sense reasoning involuntarily while answering questions or comprehending, thus leaving room for incorporating CR in models. There seems to be no recent work that invalidates, experimentally, the role of CR in CMC. QuAC, for example, is drawn from articles on personalities, and current models still lag behind the human benchmark. It may be worth experimenting if adding domain knowledge or attributes about the context, like location and gender, help improve answering these questions.

Why did the paper focus on common trends across each component rather than a single overarching classification of CMC models? The study of common trends in modeling, rather than a single overarching classification, helped in providing a multi-faceted view of CMC that can generalize on future models, and identify possible open-ended research questions, such as (a) For history selection, HAM (Qu et al., 2019b) has proved to be both effective and intuitive in selecting relevant history turns. The application of this history selection approach on previous techniques that considered immediate K turns could be experimented with. (b) As mentioned in training methodology (section 6.4), RoBERTa-based CMC model (Ju et al., 2019) that used knowledge distillation and adversarial training achieved state-of-the-art CoQA results (Reddy et al., 2019). This suggests that different training approach along with multi-task learning improves the performance of base models. These procedures could be experimented with more advanced models such as HAM (Qu et al., 2019b) and FlowDelta (Yeh and Chen, 2019).

8 Conclusion

In this paper, we provide a holistic overview of Conversational Machine Comprehension (CMC), which has seen a surge of research in recent years, owing to advancements in neural language modeling and the introduction of large-scale conversational datasets such as CoQA (Reddy et al., 2019) and QuAC (Choi et al., 2018). We discuss the challenges that make CMC different from machine reading comprehension (MRC) and compare the multi-turn conversational datasets: CoQA and QuAC, based on different CMC characteristics. To develop a high-level understanding of all the existing approaches to tackle CMC, we synthesize a general model framework and analyze the common trends across all the published models, loosely based on the components outlined in the framework. Finally, we discuss some open questions that emerged during our research and which, in our view, can be explored further. This review could serve as a compendium for researchers in this domain and help streamline research in CMC.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Razieh Baradaran, Razieh Ghiasi, and Hossein Amirkhani. 2020. A survey on machine reading comprehension systems. *CoRR*, abs/2001.01582.
- John Bell. 1999. Pragmatic reasoning: Inferring contexts. In Paolo Bouquet, Luciano Serafini, Patrick Brézillon, Massimo Benerecetti, and F. Castellani, editors, *Modeling and Using Context, Second International and Interdisciplinary Conference, CONTEXT'99, Trento, Italy, September 1999, Proceedings*, volume 1688 of *Lecture Notes in Computer Science*, pages 42–53. Springer.
- Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. Reading wikipedia to answer open-domain questions. In Regina Barzilay and Min-Yen Kan, editors, *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, pages 1870–1879. Association for Computational Linguistics.
- Yu Chen, Lingfei Wu, and Mohammed J. Zaki. 2019. Graphflow: Exploiting conversation flow with graph neural networks for conversational machine comprehension. *CoRR*, abs/1908.00059.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using RNN encoder–decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1724–1734, Doha, Qatar, October. Association for Computational Linguistics.
- Eunsol Choi, He He, Mohit Iyyer, Mark Yatskar, Wen-tau Yih, Yejin Choi, Percy Liang, and Luke Zettlemoyer. 2018. QuAC: Question answering in context. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2174–2184, Brussels, Belgium, October–November. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota, June. Association for Computational Linguistics.
- Wenchao Du and Alan W Black. 2019. Boosting dialog response generation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 38–43, Florence, Italy, July. Association for Computational Linguistics.
- Tommaso Furlanello, Zachary Lipton, Michael Tschannen, Laurent Itti, and Anima Anandkumar. 2018. Born again neural networks. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 1607–1616, Stockholm, Sweden, 10–15 Jul. PMLR.
- Jianfeng Gao, Michel Galley, and Lihong Li. 2018. Neural approaches to conversational AI. *CoRR*, abs/1809.08267.
- Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. Explaining and harnessing adversarial examples. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Amartya Hatua, Trung Tien Nguyen, and Andrew H. Sung. 2019. Dialogue generation using self-attention generative adversarial network. *2019 IEEE International Conference on Conversational Data & Knowledge Engineering (CDKE)*, pages 33–38.
- Karl Moritz Hermann, Tomáš Kočiský, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1, NIPS'15*, page 1693–1701, Cambridge, MA, USA. MIT Press.
- Minghao Hu, Furu Wei, Yuxing Peng, Zhen Huang, Nan Yang, and Dongsheng Li. 2019. Read + verify: Machine reading comprehension with unanswerable questions. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, pages 6529–6537. AAAI Press.

- Hsin-Yuan Huang, Eunsol Choi, and Wen-tau Yih. 2018a. Flowqa: Grasping flow in history for conversational machine comprehension. *CoRR*, abs/1810.06683.
- Hsin-Yuan Huang, Chenguang Zhu, Yelong Shen, and Weizhu Chen. 2018b. Fusionnet: Fusing via fully-aware attention with application to machine comprehension. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Lifu Huang, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2019. Cosmos QA: Machine reading comprehension with contextual commonsense reasoning. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2391–2401, Hong Kong, China, November. Association for Computational Linguistics.
- Ying Ju, Fubang Zhao, Shijie Chen, Bowen Zheng, Xuefeng Yang, and Yunfeng Liu. 2019. Technical report on conversational question answering. *CoRR*, abs/1909.10772.
- Zhen-Zhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. Albert: A lite bert for self-supervised learning of language representations. *ArXiv*, abs/1909.11942.
- Yujia Li, Daniel Tarlow, Marc Brockschmidt, and Richard S. Zemel. 2016. Gated graph sequence neural networks. In Yoshua Bengio and Yann LeCun, editors, *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*.
- Hongyu Lin, Le Sun, and Xianpei Han. 2017. Reasoning with heterogeneous knowledge for commonsense machine comprehension. In Martha Palmer, Rebecca Hwa, and Sebastian Riedel, editors, *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, pages 2032–2043. Association for Computational Linguistics.
- Zachary C. Lipton, Jianfeng Gao, Lihong Li, Xiujun Li, Faisal Ahmed, and Li Deng. 2016. Efficient exploration for dialog policy learning with deep BBQ networks & replay buffer spiking. *CoRR*, abs/1608.05081.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019a. Roberta: A robustly optimized BERT pretraining approach. *CoRR*, abs/1907.11692.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019b. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Yasuhito Ohsugi, Itsumi Saito, Kyosuke Nishida, Hisako Asano, and Junji Tomita. 2019. A simple but effective method to incorporate multi-turn context with BERT for conversational machine comprehension. *CoRR*, abs/1905.12848.
- Simon Ostermann, Ashutosh Modi, Michael Roth, Stefan Thater, and Manfred Pinkal. 2018. MCScript: A novel dataset for assessing machine comprehension using script knowledge. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan, May. European Language Resources Association (ELRA).
- Baolin Peng, Xiujun Li, Jianfeng Gao, Jingjing Liu, and Kam-Fai Wong. 2018. Deep dyna-q: Integrating planning for task-completion dialogue policy learning. In Iryna Gurevych and Yusuke Miyao, editors, *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pages 2182–2192. Association for Computational Linguistics.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global vectors for word representation. In Alessandro Moschitti, Bo Pang, and Walter Daelemans, editors, *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 1532–1543. ACL.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In Marilyn A. Walker, Heng Ji, and Amanda Stent, editors, *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 1 (Long Papers)*, pages 2227–2237. Association for Computational Linguistics.
- Boyu Qiu, Xu Chen, Jungang Xu, and Yingfei Sun. 2019. A survey on neural machine reading comprehension. *CoRR*, abs/1906.03824.

- Chen Qu, Liu Yang, Minghui Qiu, W. Bruce Croft, Yongfeng Zhang, and Mohit Iyyer. 2019a. BERT with history answer embedding for conversational question answering. In Benjamin Piwowarski, Max Chevalier, Éric Gaussier, Yoelle Maarek, Jian-Yun Nie, and Falk Scholer, editors, *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2019, Paris, France, July 21-25, 2019*, pages 1133–1136. ACM.
- Chen Qu, Liu Yang, Minghui Qiu, Yongfeng Zhang, Cen Chen, W. Bruce Croft, and Mohit Iyyer. 2019b. Attentive history selection for conversational question answering. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management, CIKM 2019, Beijing, China, November 3-7, 2019*, pages 1391–1400. ACM.
- Chen Qu, Liu Yang, Minghui Qiu, Yongfeng Zhang, Cen Chen, W. Bruce Croft, and Mohit Iyyer. 2019c. Attentive history selection for conversational question answering. *Proceedings of the 28th ACM International Conference on Information and Knowledge Management - CIKM '19*.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding with unsupervised learning. *Technical report, OpenAI*.
- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. Know what you don't know: Unanswerable questions for squad. In Iryna Gurevych and Yusuke Miyao, editors, *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 2: Short Papers*, pages 784–789. Association for Computational Linguistics.
- Siva Reddy, Danqi Chen, and Christopher D. Manning. 2019. Coqa: A conversational question answering challenge. *Trans. Assoc. Comput. Linguistics*, 7:249–266.
- Ayse Pinar Saygin, Ilyas Cicekli, and Varol Akman. 2000. Turing test: 50 years later. *Minds and machines*, 10(4):463–518.
- Min Joon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2017. Bidirectional attention flow for machine comprehension. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015. Pointer networks. *ArXiv*, abs/1506.03134.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime G. Carbonell, Ruslan Salakhutdinov, and Quoc V. Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. *CoRR*, abs/1906.08237.
- Mark Yatskar. 2019. A qualitative comparison of coqa, squad 2.0 and quac. In Jill Burstein, Christy Doran, and Thamar Solorio, editors, *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 2318–2323. Association for Computational Linguistics.
- Yi Ting Yeh and Yun-Nung Chen. 2019. Flowdelta: Modeling flow information gain in reasoning for conversational machine comprehension. *CoRR*, abs/1908.05117.
- Sheng Zhang, Xiaodong Liu, Jingjing Liu, Jianfeng Gao, Kevin Duh, and Benjamin Van Durme. 2018a. Record: Bridging the gap between human and machine commonsense reading comprehension. *CoRR*, abs/1810.12885.
- Yongfeng Zhang, Xu Chen, Qingyao Ai, Liu Yang, and W. Bruce Croft. 2018b. Towards conversational search and recommendation: System ask, user respond. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management, CIKM '18*, page 177–186, New York, NY, USA. Association for Computing Machinery.
- Xin Zhang, An Yang, Sujian Li, and Yizhong Wang. 2019. Machine reading comprehension: a literature review. *CoRR*, abs/1907.01686.
- Chenguang Zhu, Michael Zeng, and Xuedong Huang. 2018. Sdnet: Contextualized attention-based deep network for conversational question answering. *CoRR*, abs/1812.03593.

A General statistics for CoQA and QuAC

The dataset statistics are provided in Table 1. An example from both datasets is provided in Fig. 5

Characteristic (average)	CoQA	QuAC
<i>Dataset source</i>	Passages collected from 7 diverse domains e.g. children stories from MCTest, news articles from CNN, Wikipedia articles, etc.	Sections from Wikipedia articles filtered in the ‘‘people’’ category associated with subcategories like culture, animal, geography, etc.
<i>Conversation setting</i>	Questioner-Answerer setting where both have access to the entire context.	Teacher-Student setting where the teacher has access to the full context for answering, while the student has only the title and summary of the article.
<i>Question type</i>	factoid.	open-ended, highly contextual.
<i>Answer type</i>	free-form with an extractive rationale.	Extractive span which can be yes/no or ‘No Answer’. It also provides dialog acts.
<i>Total number of dialogs</i>	8K	14K
<i>Total number of questions</i>	126K	100K
<i>Context, Question and Answer token lengths</i>	271, 5.5, 2.7	401, 6.5, 14.6
<i>Turns per dialog</i>	15.2	7.2
<i>Unanswerable questions</i>	Very low and often erroneously marked.	Significant quantity of type ‘missing info’
<i>Evaluation metrics</i>	F1 scores for in-domain, out-of-domain and overall	F1, Human Equivalence Quotient (HEQ) scores at question and dialog levels.

Table 1: A comparison of the multi-turn conversational datasets- CoQA (Reddy et al., 2019) and QuAC (Choi et al., 2018) based on different characteristics defined in their papers and the analysis paper by Yatskar (2019).

B GraphFlow

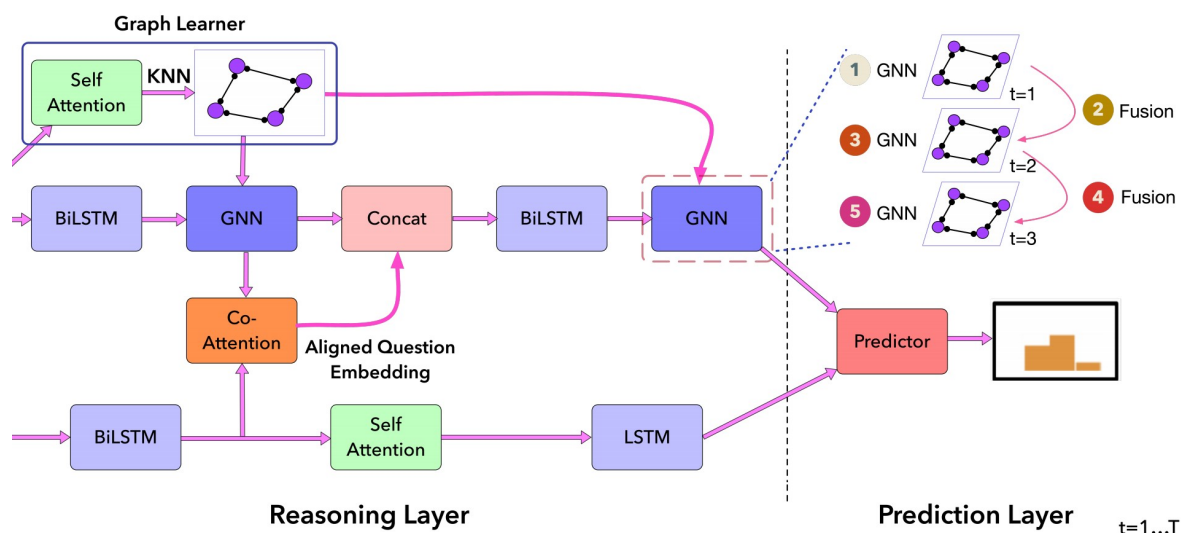


Figure 4: Architecture of the Reasoning Layer of GraphFlow. Context graph-based flow sequence is processed using GNNs and alternated with bi-LSTM and co-attention mechanisms. Source : (Chen et al., 2019)

The Virginia governor's race, billed as the marquee battle of an otherwise anticlimactic 2013 election cycle, is shaping up to be a foregone conclusion. Democrat Terry McAuliffe, the longtime political fixer and moneyman, hasn't trailed in a poll since May. Barring a political miracle, Republican Ken Cuccinelli will be delivering a concession speech on Tuesday evening in Richmond. In recent ...

Q₁: What are the candidates **running** for?

A₁: Governor

R₁: The Virginia governor's race

Q₂: **Where**?

A₂: Virginia

R₂: The Virginia governor's race

Q₃: Who is the democratic candidate?

A₃: **Terry McAuliffe**

R₃: Democrat Terry McAuliffe

Q₄: Who is **his** opponent?

A₄: **Ken Cuccinelli**

R₄: Republican Ken Cuccinelli

Q₅: What party does **he** belong to?

A₅: Republican

R₅: Republican Ken Cuccinelli

Q₆: Which of **them** is winning?

A₆: Terry McAuliffe

R₆: Democrat Terry McAuliffe, the longtime political fixer and moneyman, hasn't trailed in a poll since May

(a) A QA dialog example in the CoQA dataset. Every dialog is based on a context and each turn of the dialog contains a question (Q_i), an answer (A_i) and a rationale (R_i) that supports the answer. There is sufficient co-referencing between dialog turns as seen in this example – 'Where' in Q_2 follows on the candidature mentioned in Q_1 , 'his' in Q_4 points to A_3 , 'he' in Q_5 references A_4 , and 'them' in Q_6 refers to people mentioned in both A_3 and A_4 . Source: (Reddy et al., 2019)

TITLE: Anna Vissi

BACKGROUND: Anna Vissi, also known as Anna Vishy, is a Greek Cypriot singer, songwriter, actress, television presenter, radio personality, and businesswoman.

SECTION: 1983-1989: Collaboration with Nikos Karvelas

In May 1983, she married Nikos Karvelas, a composer, with whom she collaborated in 1975 and in November she gave birth to her daughter Sofia. After their marriage, she started a close collaboration with Karvelas. Since 1975, all her releases have become gold or platinum and have included songs by Karvelas... In March 1984, she released Na 'Hes Kardia ("If You Had a Heart"). The album was certified gold. The following year her seventh album Kati Simveni ("Something Is Happening") was released which included one of her most famous songs, titled "Dodeka" ["Twelve (O'Clock)"] and reached gold status selling 80,000 units. In 1986 I Epomeni Kinisi ("The Next Move") was released. The album included the hit Pragmata ("Things") and went platinum, becoming the best selling record of the year... (No Answer)

Q: What happened in 1983?

A: *In May 1983, she married Nikos Karvelas, a composer.*

Follow-up: Yes Affirmation: NA

Q: Did they have any children?

A: *In November she gave birth to her daughter Sofia.*

Follow-up: Maybe Affirmation: Yes

Q: Did she have any other children?

A: (No Answer)

Follow-up: No Affirmation: NA

Q: What collaborations did she do with Nikos?

A: *Since 1975, all her releases have become gold or platinum and have included songs by Karvelas.*

Follow-up: Maybe Affirmation: NA

Q: What were some of the songs?

A: *one of her most famous songs, titled "Dodeka" ["Twelve (O'Clock)"]*

Follow-up: Yes Affirmation: NA

Q: How famous was it?

A: *reached gold status selling 80,000 units.*

Follow-up: Yes Affirmation: NA

Q: Did she have any other famous songs?

A: *The album included the hit Pragmata ("Things") and went platinum.*

Follow-up: Yes Affirmation: Yes

(b) A QA dialogue example from the QuAC dataset. Dialogs are prepared in a student-teacher setting over a section of a Wikipedia article, where the student questions based on the title, background, and start of the context and the teacher responds to the questions in the form of text spans from the context and dialogue acts. Dialog acts include follow-up i.e. whether the questioner should, could or should not cross-question, and, affirmation i.e. if the question can be answered as Yes/ No, or NA if it's not a yes/no kind of question. A 'No Answer' token is appended to the context which is expected to be selected by the model in the case of unanswerable questions. The questions are open-ended due to the asymmetric nature of dataset. There is also sufficient co-referencing – 'she' in Q_3 refers to the protagonist and is a succession of Q_2 , similarly Q_7 is a follow-up on Q_5 , 'it' in Q_6 refers to song mentioned in A_5 . Source: (Choi et al., 2018)

Figure 5: Illustrative examples from the conversational datasets– CoQA and QuAC.

C A summary of the common CMC models

The following tables provide a summary of the CMC models published on the CoQA² and QuAC³ leaderboards⁴. The table also provides a link to the official code repositories for the models. Although most of the models are published on both leaderboards, some models are very specific to one of the datasets.

⁴Summarized as of July 1, 2020. Please note that the summary is only available for models that have an attached manuscript or repository.

Model	Overall test F1 scores		Approaches in			Additional Training methods used (if any)	Repository
	CoQA	QuAC	History Selection	History Modelling	Contextual Reasoning		
<i>Human Performance</i>	88.8	81.1	N/A	N/A	N/A	N/A	N/A
<i>RoBERTa + AT + KD (Jin et al., 2019)</i>	90.7 (ensemble) 90.4 (single)	74.0 (ensemble) 73.5 (single)	All history Q&A (until truncated)	Appending history Q&A to the current question (truncate query if question tokens exceed)	Contextual Integration using single-turn, pre-trained LM - RoBERTa (Liu et al., 2019a)	Rationale tagging multi-task (token prediction in CoQA evidence) along with Adversarial Training (Goodfellow et al., 2015) and Knowledge Distillation (Furlanello et al., 2018).	No official code repo is given.
<i>XLNet + Augmentation</i>	89.0	71.2	All history Q&A (until truncated)	Appending history Q&A to the current question (truncate query if question tokens exceed)	Contextual Integration using single-turn XLNet-large (Yang et al., 2019) (24L, 1024H, 16A)	None	https://github.com/sveveheng23/xlNet_extension_Lf
<i>BERT + Answer Verification</i>	82.8	N/A	All history Q&A (until truncated)	Appending history Q&A to the current question	Contextual Integration using single-turn BERT-Base (Devlin et al., 2019) (12L, 768H, 12A)	Additional answer verification multi-task as in (Hu et al., 2019).	https://github.com/sogou/SMRCToolkit
<i>SDNet</i>	79.3 (ensemble) 76.6 (single)	N/A	Last 2 Q&A	Appending history Q&A to the current question	Attention-based Reasoning with sequential model-SDNet (Zhu et al., 2018)	None	https://github.com/Microsoft/SDNet
<i>BERT w/2-context (Ohsugi et al., 2019)</i>	78.7	64.9	Last 2 Q&A	Append each Q (or A) to the current question for every BERT model.	Concatenates contextual outputs from BERT-Base models for every history question and answer (2N+1 sequences for N turns) and runs Bi-GRU over the aggregated sequence.	Answer type (Yes/No) prediction multi-task.	No official code repo is given.
<i>HAM (Qu et al., 2019b)</i>	N/A	65.4	Dynamic selection policy by attending over all the previous history turns and deriving weights based on their contextual correlation with the current turn.	Concatenate context tokens with positional history answer marker embeddings corresponding to each dialog turn at a relative position from the current turn e.g. marker embedding 2 for turn 5 when the current turn is 7. Ignores history questions.	Aggregates contextual outputs from BERT-Base models for every answer (N sequences for N turns) using attention-based selection weights to obtain word-level and seq-level representations.	Dialog act (continuation and affirmation) prediction multi-task	https://github.com/prdwb/attentive_history_selection
<i>Bert-FlowDelta Yeh and Chen, 2019)</i>	77.7	65.5	All history questions only	Integrate latent representations generated via contextual reasoning on the history turns (FLOW)	Each history turn is first passed through BERT whose last and second last layer outputs are input to separate FlowQA models. Reasoning is done using the Integration-Flow approach with information gain as the propagated FLOW representation.	None	https://github.com/Miuiab/FlowDelta
<i>GraphFlow (Chen et al., 2019)</i>	77.3	64.9	All the history Q&A	Uses a mix of all modeling techniques- encodes relative dialog turn numbers within each question embedding and appends all to the current question, then concatenates answer marker embeddings to context tokens and finally encodes the turn into N context graphs which are used for reasoning.	The encoded context graphs per turn are used as the propagated FLOW representation in the Integration-GraphFlow architecture.	Answer type (Yes/No) prediction multi-task.	https://github.com/hugochan/GraphFlow
<i>FlowQA (Huang et al., 2018a)</i>	75.0	64.1	All history questions only	Integrate latent representations generated via contextual reasoning on the history turns (FLOW)	Reasoning is done using the Integration-Flow approach with latent history turn representation directly as the propagated FLOW representation (instead of information gain).	Unanswerability multi-task.	https://github.com/momohuang/FlowQA
<i>BERT+HAE (Qu et al., 2019a)</i>	N/A	62.4	Last K history answers (max K=11)	Concatenate context tokens with history answer marker embeddings to identify if the token is present at any dialog turn. Ignores history questions.	Contextual integration using a single-turn BERT-base	None	https://github.com/prdwb/bert_hae
<i>BIDA-F++ w/2-Context (QuAC baseline) (Choi et al., 2018)</i>	67.8	60.1	Last 2 Q&A	Encodes relative dialog turn numbers within each question embedding and appends all to the current question. Concatenates answer marker embeddings to context tokens.	Combines self-attention with BiDAF (Seo et al., 2017) followed by multi-layered BiLSTM to obtain contextualized representation.	Dialog act multi-task in QuAC.	Use implementation on http://allennlp.org
<i>DrQA+PGNER (CoQA baseline) (Reddy et al., 2019)</i>	65.1	N/A	All history Q&A	Appending history Q&A to the current question	BiLSTM based contextual integration over encoded tokens for extracting evidence (DrQA), followed by PGNet (Vinyals et al., 2015) which uses attention-based neural machine translation for abstractive answer reasoning (selecting relevant tokens).	None	https://github.com/stanfordnlp/coqa-baseLines
<i>Vanilla DrQA (Chen et al., 2017)</i>	52.6	N/A	All history Q&A	Appending history Q&A to the current question	BiLSTM based contextual integration over encoded tokens for extractive span.	None	https://github.com/facebookresearch/DrQA

Table 2: Model-wise summary of all the published CMC models for QuAC and CoQA.