

A Unified Approach to Entity-Centric Context Tracking in Social Conversations

Ulrich Rückert, Srinivas Sunkara, Abhinav Rastogi, Sushant Prakash, Pranav Khaitan

Google Research

1600 Amphitheatre Pkwy, Mountain View, CA 94043, USA

{rueckert, srinivasksun, abhirast, sush, pranavkhaitan}@google.com

Abstract

In human-human conversations, Context Tracking deals with identifying important entities and keeping track of their properties and relationships. This is a challenging problem that encompasses several subtasks such as slot tagging, coreference resolution, resolving plural mentions and entity linking. We approach this problem as an end-to-end modeling task where the conversational context is represented by an entity repository containing the entity references mentioned so far, their properties and the relationships between them. The repository is updated turn-by-turn, thus making training and inference computationally efficient even for long conversations. This paper lays the groundwork for an investigation of this framework in two ways. First, we release *Contrack*, a large scale human-human conversation corpus for context tracking with people and location annotations. It contains over 7000 conversations with an average of 11.8 turns, 5.8 entities and 15.2 references per conversation. Second, we open-source a neural network architecture for context tracking. Finally we compare this network to state-of-the-art approaches for the subtasks it subsumes and report results on the involved tradeoffs.

Keywords: Context Tracking, Coreference Resolution, Entity Linking

1. Introduction

Computers and mobile phones have changed how people communicate. A large amount of today’s interpersonal communication happens in messaging apps on mobile devices or on chat and discussion services on the internet. Consequently, this has piqued the interest of the research community in developing assistive technologies for human-human conversations. Representing the current status of a conversation in a succinct and semantically complete way is a central component of such technologies. At the core of this endeavor lies the task of tracking the entities mentioned in a conversation, their properties and the relationships that are being expressed about them. In this paper, we frame this task, which we call *Context Tracking*, as an online machine learning problem, where the model is expected to track the current status of the conversation at any time. This formulation extends and complements existing research in three key areas.

First of all, in this framework the model ingests the messages of a conversation turn by turn and updates a growing repository of detected entity references in each turn. This allows for fast inference because the model only needs to ingest a single message and the repository instead of the entire conversation history. That is important particularly for long conversations, where ingesting the entire dialogue history is not feasible.

Second, the repository of entity references serves as a condensed storage of the semantic knowledge conveyed in the conversation. In this sense, context tracking serves the same purpose for open-domain conversations that dialog state tracking does for task-oriented conversations. While the current formulation only deals with tracking entity references, the same frame-

work can be used to track relationships between entities or to build personal knowledge graphs. Once it can capture a large enough share of semantic information, the repository could be used to address higher-level tasks such as reasoning (Chen et al., 2020), question answering (Huang et al., 2019) or summarization (Huang et al., 2020).

Third, since the framework aims at extracting a rich set of semantic information, it unifies existing NLP tasks such as entity recognition, slot tagging, coreference resolution and resolving plural mentions in one formulation. This means it is more parameter-efficient, because it shares parameters across multiple tasks. Recent research on large neural networks (Brown et al., 2020; Raffel et al., 2020) has raised the prospect that larger networks with multiple endpoints may even improve predictive accuracy over single-task approaches.

In this paper, we take the first steps towards a systematic investigation of Context Tracking and lay the groundwork for further research. The most important part of this effort is the new *Contrack* dataset, which contains over 7000 social open-domain conversations with annotations for people and location entities with an average of 11.8 turns, 5.8 entities and 15.2 references per conversation. It is publicly available and can be downloaded at <https://github.com/google-research-datasets/contrack>.

In the second part of the paper we present a simple baseline model for Context Tracking. The model is a straightforward adaptation of a Transformer-based (Vaswani et al., 2017) neural network to the peculiarities of Context Tracking. Its conceptual simplicity makes it well suited to serve as a benchmark for further work. We release the code for this baseline

implementation as Open Source. It is available at <https://github.com/google-research/google-research/tree/master/ctrack>. The third part of the paper reports the results of experiments we performed with the baseline model. These experiments give some insights on which parts of Context Tracking are solved easily and which parts require more research. We also compare the dataset and baseline model to other approaches for related tasks to get a sense of the tradeoffs involved with modeling multiple endpoints concurrently in turn-by-turn fashion.

2. Related Work

Context tracking or dialogue state tracking has been a well studied problem in task-oriented systems (Mrkšić et al., 2017; Rastogi et al., 2017), where the context is tracked in terms of the slots of the underlying task or API. However, a slot based representation is not suitable for human-human conversations because they tend to cover multiple domains making enumeration of slots impractical. Models like Meena (Adiwardana et al., 2020) and DialoGPT (Zhang et al., 2020) model human-human conversation using a latent representation of the dialogue context. Such a representation is not suitable for applications where an explicit or interpretable representation of context is desired.

The availability of public datasets has played a central role in driving this area of research. Many datasets for coreference resolution have focused on either document based coreferences, like OntoNotes (Ralph Weischedel et al., 2013) (discussed by Pradhan et al. (2012)), ACE (Doddington et al., 2004) or conversations about specific tasks like MuDoCo (Martin et al., 2020b) and CoQA (Reddy et al., 2019). Despite their scale and popularity, these datasets are not similar to human-human conversations as it is known that a large portion of human-human conversation centers on chit-chat, socialization and personal interests (Dunbar et al., 1997). Datasets like Persona Chat (Zhang et al., 2018), Daily Dialog (Li et al., 2017) and others contain social conversations in different settings, but they are not suitable to be used for context tracking as they are not focused on references to entities in the conversations.

A variety of data driven techniques have been studied for semantic parsing and resolving ambiguities arising commonly in natural language (Yadav and Bethard, 2018; Lee et al., 2017; Sevgili et al., 2020). With growing interest in Conversational AI and owing to the availability of public datasets (Chen and Choi, 2016; Gopalakrishnan et al., 2019; Martin et al., 2020a), some of these techniques have been successfully applied to human-human conversations. Chen and Choi (2016) studied the problem of coreference resolution and linking mentions to a fixed set of entities, whereas Zhou and Choi (2018) investigated the problem of resolving plural mentions, which refer to multiple entities. Obtaining structured information from conver-

sations has also been explored by El-Assady et al. (2017) for generating a graphical representation of entities mentioned in a conversation through pairwise co-occurrence, and by Li et al. (2014) for constructing a user centric personal knowledge graph from conversations.

These techniques separately address various important aspects of context tracking and representation of the entities mentioned in the conversation. Our framework combines the several important pieces of these works including coreference resolution, plural mentions, named-entity recognition, and attribute classification.

3. The Context Tracking Modeling Task

In the context tracking task the model processes conversations incrementally, one turn at a time. The model keeps track of the context within a repository of entity references. In order to disambiguate references in later turns, the repository also needs to store data about grammatical gender, group membership and other context data. In this section, we first describe the elements that are stored in the repository. Afterwards, we discuss the dataset we collected for this task.

3.1. Context Representation

The conversation context is represented by a repository containing all *entity references* seen so far in the conversation. An entity reference is a span of tokens that refers to an entity. When adding an entity references to the repository we store the following information:

- A unique *entity ID*. This numerical ID uniquely identifies the entity or group of entities present in the entity repository. If two references refer to the same entity, they have the same entity ID.
- A list of explicit *entity properties*. Properties specify relevant information about the entity such as the type of entity (e.g. people or location), grammatical gender, and whether the reference is a plural form. One could include other signals relevant for disambiguation or for downstream tasks that rely on this information.
- *Group membership*. Some references are plural in nature, for example when referring to a group of people with *they*. Such plural references (which we call *groups* in the following) are added in the repository, and the group membership field enumerates the known entities which are members of a group.
- *Implicit context data*. Additionally, the repository can optionally contain signals encoding information that is relevant to an entity but which cannot be easily captured by an explicit property. These signals could either be distributed (e.g. an embedding of surrounding tokens to the most recent

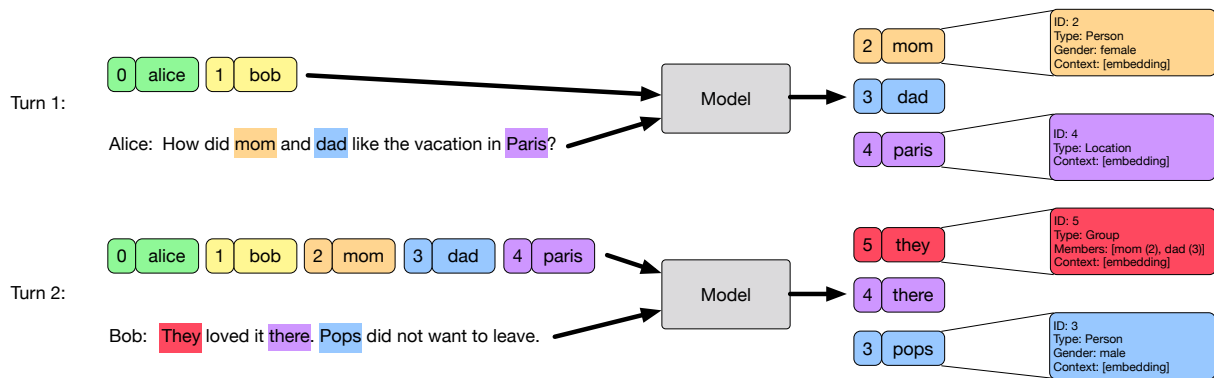


Figure 1: Inspecting the contrack representation for two turns of an example conversation. The model is expected to identify tokens which refer to entities and outputs semantic data about them. The entity references output in turn one are passed as input to the model in the second turn.

reference) or discrete (e.g. a bag of salient words co-occurring with mentions of the entity).

Figure 1 illustrates this representation on two turns of an example conversation between Alice and Bob about their parents’ vacation. Before the conversation starts, the repository contains two entity references referencing the conversation participants Alice and Bob. When Alice sends the first message “How did mom and dad like the vacation in Paris?”, the model identifies three new entities, one each for the person entities *mom* and *dad* and one for the location *paris*. It serially assigns unique numerical IDs two, three and four to these new entities respectively, since the IDs zero and one are already taken. Note that the references also contain properties such as type and grammatical gender. In the next turn the repository contains five entity references.

For the next message “They loved it there. Pops did not want to leave.”, the model is expected to identify three new entity references. The first one is *they*, which is a group containing the two members *mom* (ID two) and *dad* (ID three). The references *there* and *pops* refer to existing entities and are thus labeled with the corresponding IDs. This annotation conveys that *pops* and *dad* refer to the same entity, as do *there* and *paris*.

3.2. Dataset Collection and Details

Real human to human casual conversations contain personal data and are hard to obtain. Hence, there is a need to collect a synthetic dataset which offers the same challenges as a real human-human conversation corpus. Our data collection procedure, outlined below, ensures that the conversations are natural, cover a variety of topics and mention different entities while minimizing annotation errors.

3.2.1. Scenario Generation

Our conversations are seeded by manually created scenarios which are short summaries of the content of the conversation to be generated. Here is a sample scenario:

Jeffrey’s siblings, Nate and Marie are going to visit him and he asks Beverly for suggestions on where to take them. She asks him what they like to do and what they had done the last time they visited. She comes up with a few suggestions and they finally decide on a plan.

The scenarios are created by a separate group of expert crowd workers. They were asked to focus on variety and to leave certain details unspecified to facilitate creation of multiple conversations from each scenario. Scenarios enable us to control the broad topic of the conversation, to ensure that a variety of settings are captured and that relevant entities are mentioned.

3.2.2. Wizard-of-Oz Setup

We developed a web application for collecting conversations for a given scenario by pairing two crowd workers together. For the last third of the dataset, we switched to using a single crowdworker for playing both roles in a conversation (similar to findings of Byrne et al. (2019)), as we found that it improved efficiency without a loss in quality. After an initial filtering of these conversations, the crowd workers annotated them by identifying spans which represent entity references. For each entity, the workers annotated the type of the entity and resolved all pronouns or noun phrases referring to this entity.

Over the course of the data collection for both tasks, we monitored the quality of data using heuristics like measuring the word overlap between the scenario and the chat, repeated turns for chat collection and comparing with known pronoun properties (e.g., the word “them” refers to a group etc.). We used these methods to update the list of good quality crowd workers who were used for subsequent data collection.

Finally, we had an additional round of verification on the annotated data to verify the annotations and to rate the chats on a scale of 1-5 for their naturalness and coherence. In addition chats with non-fluent English and




Figure 2: The number of tokens spans which are most frequently annotated as entity references.




Figure 3: Distribution of number of turns per conversation. The majority of conversations have 10-14 turns.

grammatical mistakes not expected in a casual conversations were also given a low score. Chats with a score of 4 or 5 were retained and chats with a score of 3 were sent for rating by another crowd worker. Finally, we removed chats with scores less than 4 and we sent chats with annotation errors back for re-annotation. To increase the diversity of entities mentioned in our data and to prevent the model from overfitting on the names in the training set, we also release an augmented set of conversations in which we replace people names in the conversations with names randomly chosen from a very large collection.




Figure 4: Distribution of number of entities per turn. Around 8% of the turns have no entities while 11% of turns have 5 or more entities.

3.2.3. Dataset Statistics

The resulting Contrack Dataset is a human to human casual conversation dataset containing 7245 conversations with annotations for entities. Currently we have annotations for entities representing people and locations. The annotations contain signals for grammatical gender and group membership for people entities, and grammatical count and group membership for location entities. In total, we have 85,538 turns in the dataset with an average of 11.5 tokens per turn. Figure 2 provides statistics of the most frequently annotated token spans in the dataset. Figures 3 and 4 show the distribution of number of turns per conversations and the distribution of number of entities per turn respectively.

4. The Contrack Baseline Model

In this section, we present a simple model for context tracking to demonstrate the feasibility of the proposed approach and to act as a baseline for future studies on this dataset. The model takes an utterance and an entity repository as input, and outputs a list of entity references which will be added to the repository. One major challenge is the *reference co-occurrence* problem, which occurs whenever a new entity is introduced and then referred to in the same turn by another token span (e.g. *karen* and *her* in “*have you heard karen totaled her car*”). It is a challenge because the entity ID for the reference *her* needs to be the same as the entity ID for *karen*, which has not been assigned yet. Our baseline model provides a simple solution to this problem, while jointly modeling all context tracking subtasks.

4.1. Model Architecture

The model solves the *reference co-occurrence* problem by dividing the computation into two stages. The first stage identifies all references to new entities that have not been mentioned in the conversation so far. These newly mentioned entities are sequentially assigned new entity IDs, so that the rest of the model can refer to them in subsequent computation. For example, The

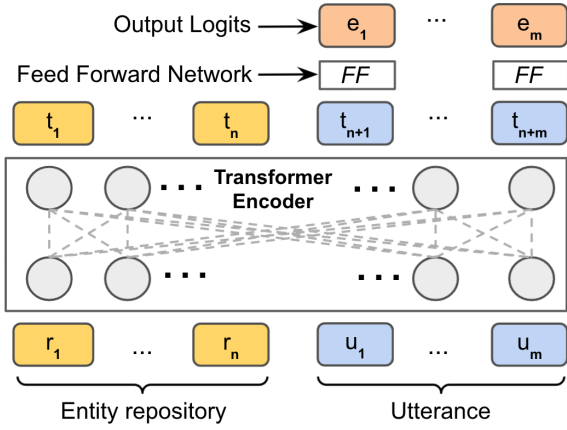


Figure 5: Architecture for each of the two computation stages in the context tracking baseline model.

second stage performs the rest of the tasks including assigning entity IDs to the other entity references, predicting properties of a reference (e.g. gender, plurality etc.) and group memberships for references referring to a group of entities (e.g. *they, us* etc.).

Both stages have the same underlying architecture, which is shown in Figure 5. The network takes two sequences as inputs. The first sequence r_1, \dots, r_n encodes the entity repository, where each element represents a single entity reference. The second sequence u_1, \dots, u_m corresponds to the input utterance, with each element refers to a token in the utterance. The vector representation of these sequences are obtained by concatenating the embedded representation of various features as described in the following section. The input sequences are fed to a 1-layer transformer encoder, which fuses the representations of the two sequences to generate the sequence t_1, \dots, t_{n+m} . The outputs t_{n+1}, \dots, t_{n+m} are then fed to a feed forward network to generate output logits e_1, \dots, e_m of appropriate dimension, from which all the required outputs are obtained as described in the Context Tracking section below.

4.2. Input Feature Representation

The entity repository and utterance token sequences are represented by concatenating the embedded representations of the input features. Some features are only available for the entity reference repository, and hence are set to 0 for the token sequence. These features are: (i) *ID* - It is a one-hot vector encoding the ID that has been assigned to an entity. (ii) *Meta* - A Boolean vector indicating whether the reference introduces a new, previously unseen entity. (iii) *Properties* - Each categorical value of a property (e.g. gender=male) is assigned a Boolean value, and the Boolean representation for all the values for all the properties are concatenated together to represent the properties of an entity. (iv) *Membership* - A Boolean vector which denotes which

entities are members of the referenced group. The vector is zero for non-group mentions. (v) *Context* - This is a one-hot vector of size 2 indicating whether the underlying entity is the sender or recipient of the message and a binary vector indicating which of the preceding turns a reference was part of.

A few other features are available for both the entity reference repository and the utterance token sequence. These are: (i) *Type* - Boolean value indicating whether a sequence element is an entity reference or a token. (ii) *Signals* - Boolean features which indicate if the underlying element belongs to a predefined lexicon. The reported model uses a single lexicon of first names. (iii) *Word2Vec* - It contains the Word2Vec (Mikolov et al., 2013) embedding of the underlying token. For entity references, it contains the mean of the embeddings of the tokens which were marked as references to that entity in earlier turns. (iv) *BERT* - It contains the BERT (Devlin et al., 2019) embeddings for each element. They are obtained by feeding the two sequences to a BERT encoder, and then taking the average of all wordpiece tokens corresponding to each element. The Word2Vec and BERT embeddings respectively provide a context free and context dependent representation of the element, both of which are important for context tracking.

4.3. Context Tracking

The context tracking model outputs logits which encode information about the entity references in the input utterance. These references are added to the entity repository to update it for the next turn. Each entity reference contains four types of labels: (i) *Entity ID*: It clusters the references referring to the same entity together. (ii) *Meta*: This is a Boolean label indicating whether a reference corresponds to the first mention of an entity. It can be used to identify the source entity among the ones which share the same entity ID. (iii) *Properties*: This contains predictions for each of the labels such as gender, plurality, type etc. (iv) *Group memberships*: For group references, it contains the entity IDs for all the member entities. All these features are converted to Boolean vectors similar to the procedure described in the preceding section to obtain ground truth labels without any loss of information. Next, we describe how the model predicts output labels in two stages and which loss function is used for training.

Stage 1 - The first stage just identifies the spans corresponding to new entities in the input utterance. Each of the output logit vectors $e_i, 1 \leq i \leq m$, consists of two elements, which are converted to Booleans using a threshold of zero. The two Booleans respectively denote whether the underlying token is at the beginning or inside a span referring to a new entity. The tokens for which both these Booleans are zero are not part of a new entity span. Once these spans are found, the next available entity IDs are sequentially assigned to them,

and their Meta attribute is also populated.

Stage 2 - The second stage calculates the rest of the output labels. In addition to the inputs described in the Input Feature Representation section above, representations of the token level outputs obtained from the first stage (entity ID and Meta) are concatenated to the token vectors. A similar network as Stage 1 is then used to convert the input sequences to $2K + P$ output logits for each token. Here, K is the maximum number of supported entities and P is the total number of values taken by all the properties. These logits can be converted to Boolean vectors by thresholding at 0. Those logits specify the model outputs, with the first K logits giving the entity IDs, followed by P logits determining the value of each of the properties and the remaining K logits giving group memberships. Please note that no span annotations are needed in this stage because the token level entity ID predictions identify the spans. This means one can simply merge successive tokens with the same predicted entity ID into a single reference.

Loss function - Each logit x and its corresponding binary label $y = \pm 1$ give rise to a single loss term using hinge loss, defined as $\max(0, 1 - xy)$. The overall loss is obtained by summing all of the loss terms arising from each of the logits for each token, while masking out the summands arising from labels, which do not need to be considered at that particular position. For example, the logits for group membership predictions contribute to the loss only when the underlying reference is plural in the ground truth. Since the two stages are trained jointly, the overall loss L is defined as $\alpha L_1 + L_2$, where L_1 and L_2 are the losses of the two stages and α is a tunable parameter. Empirically, $\alpha = 6$ gave us the best results, which indicates that misidentifying new entities should have a higher penalty.

4.4. Implementation Details

All the experiments use the same hyperparameters, with the transformer encoder made up of a single self-attention layer with 9 heads followed by a feed-forward network with a hidden-dimension of 800. The model is configured to track up to twenty entities ($K = 20$) and the input sequence pair is restricted to be no longer than 100 elements altogether. We do not use dropout during training and utilize Adam optimizer (Kingma and Ba, 2015) with a learning rate of 0.0001 and batch size 20. The context tracking model is recurrent because the entity reference repository output in the current turn is used as an input for the next turn. We resolve this dependency during training by applying a teacher-forcing approach (Goyal et al., 2016) for simplicity, where the model uses ground truth labels as input instead of its output in the last step. Since the distribution of the entity IDs is skewed towards the smaller numbers in the dataset, we randomize the entity IDs in the training examples to make the model more robust. This is done by adding a single random number to all entity IDs in

Section	P	R	F1
New Entities	76.1	78.8	77.4
Existing ID	66	63.7	64.8
Properties	98.9	98.8	98.8
Membership	79.6	81.7	80.6

Table 1: Results for the four endpoints of the Contrack baseline model.

Nr of Turns	With Teacher Forcing	W/o Teacher Forcing	Difference
0	88.9	88.9	0.0
1	84.6	78.3	6.3
2	84.4	67.8	16.6
3	79.6	61.3	18.2
4	77.8	61.4	16.4
5	76.2	52.1	24.1
6	75.3	54.6	20.7
7	80.7	59.3	21.4
8	79.8	58.7	21.2
9	82.5	62.1	20.3

Table 2: Results on how error propagation affects prediction accuracy by turn. The columns give the percentage of correctly predicted entity IDs with and without Teacher Forcing.

a conversation and taking modulo K to keep the IDs within the supported range.

5. Experiments

First, we conduct several experiments to study the performance of our baseline model on the Contrack dataset. Next, we compare the baseline model to high performing task-specific models for a few conversational understanding tasks subsumed by context tracking.

5.1. Evaluation on Context Tracking

The evaluation dataset is constructed by randomly setting aside 15% of the data. This is done while ensuring that all conversations stemming from the same scenario are either assigned to the training or the evaluation set, but not split between the two. Unlike training, which uses teacher forcing, the model uses the predicted entity reference repository in a given turn as an input to the next turn. This makes the model more computationally efficient by allowing it to process each turn only once, and doesn't put any theoretical bound on the number of turns in a conversation. However, it also makes the model prone to error propagation because of a recurrent dependency.

We evaluate the performance of the two stages of model on four endpoints - (i) Identifying *new entities*, (ii) Associating entities to *existing IDs*, (iii) Predicting properties for all entities, and (iv) Calculating *membership*




Figure 6: Accuracy numbers for frequently occurring entity references. The blue bars give the percentage of entity references that were detected correctly, the red bars give the percentage with correctly predicted entity IDs.

of group entities. The first endpoint is implemented as part of the first stage, the others are computed by the second stage. We report the precision, recall and F1 score for each of these endpoints.

Table 1 lists the value of evaluation metrics. They indicate that the model is generally accurate on properties and to a lesser degree on group membership, while detecting new entities is easier than resolving existing ones. That is true in particular for location entities where the model is generally not able to detect entities as reliably as with people. Figure 6 shows the accuracy results for the most frequent occurring tokens. Not surprisingly, references to the first and second persons are the easiest to resolve while references to locations and third persons are more challenging. All in all, the baseline detected 96.8% and resolved the IDs of 64.7% of references correctly.

5.2. Effects of Error Propagation

Before we compare the baseline model’s performance on other tasks, let us try to quantify the effect of error propagation within Context Tracking. Processing the conversation turn-by-turn can lead to cases where misprediction in one turn can result in errors in all subsequent turns. This makes Context Tracking more difficult than tasks where the model can use the whole conversation. To evaluate the magnitude of this effect, we compare the percentage of correctly predicted entity IDs of the baseline Contrack model with propagat-

ing entity repositories turn by turn (i.e. without teacher forcing) with a version which uses the ground-truth input repository as an input in each turn (i.e. with teacher forcing). The results are listed in Table 2. We can see that the performance of the model drops as the index of the turn in a conversation increases. The difference is largest at five turns with a 24% accuracy difference and remains flat around 20% as the turn number increases.

5.3. Evaluation on Other Tasks

In this section, we demonstrate the applicability of the baseline model to conversational understanding tasks subsumed by context tracking. In all experiments we use a fully trained baseline model, which makes predictions for all endpoints. Its performance is compared to the target model trained on the single task under consideration. The baseline differs in that it is more parameter-efficient (because it models multiple endpoints concurrently) and easier to scale (because it makes predictions turn by turn). However, unlike the competing systems, the baseline’s errors in one turn will propagate to later turns as described in the section on error propagation above. This means the results in this section can be considered to be a measure of the accuracy costs incurred by switching to turn-by-turn prediction and by using a joint model.

Slot Tagging. The slot tagging task (Goo et al., 2018; Zhang et al., 2020) deals with identifying spans from an utterance and associating a type or slot to these spans.

Attribute Classification			
	Precision	Recall	F1
Coach	98.2	97.8	98.0
Coach-B	98.2	98.0	98.1
Contrack	99.0	99.0	99.0

Singular Coreference Resolution			
	Precision	Recall	F1
E2E Coref	97.6	87.3	91.5
Contrack	85.5	70.4	76.9

Plural Coreference Resolution			
	Precision	Recall	F1
CZC	81.2±1.0	73.3±0.4	75.9±0.5
TE	80.4±0.8	76.5±1.2	78.0±0.6
Contrack	73.2±0.6	68.5±0.7	69.7±0.8

Entity Linking			
	Precision	Recall	F1
CZC	71.8±0.4	61.4±0.4	66.2±0.4
TE	71.1±0.4	64.2±1.3	67.4±0.8
Contrack	61.2±1.0	55.5±1.6	58.2±1.2

Table 3: Experiments: Comparison with various task specific SoTA models. Contrack is the baseline described in the Contrack Baseline Model section, Coach model is from (Liu et al., 2020), Coach-B is the Coach model trained with BERT embeddings, CZC is from (Chen et al., 2017) and TE is from (Zhou and Choi, 2018). More details can be found in the Experiments section.

The baseline model subsumes slot tagging by identifying the entity reference tokens as spans and assigning a type to it by predicting a property. We benchmark the baseline model against Coach (Liu et al., 2020), which is one of the recent, best performing models for slot tagging. We train this model on the Contrack dataset, where it is only asked to identify the spans corresponding to entity references, and to assign their properties (grammatical gender for person entities, count for location entities). All other annotations are ignored. Table 3 shows the results using the default fastText (Bojanowski et al., 2017) embeddings used by the baseline and BERT embeddings used by our model. Even though both models perform very well on this task, our simple baseline outperforms the Coach model in both settings.

Singular Coreference Resolution. We compare with the state of the art E2ECoref system (Lee et al., 2018) for coreference resolution, trained using default parameters on the Contrack dataset with all the plural entity references removed. Table 3 gives the results. The E2ECoref system outperforms the baseline substantially. This can be partly explained by the fact that E2ECoref processes the entire conversation together, and hence has access to future turns while making a

decision. As mentioned above the Contrack baseline makes its decisions based only on the current utterance and outputs from the previous turn and is thus prone to error propagation. However, future work may show that one can transfer some of the techniques which make E2ECoref successful to context tracking.

Plural Coreference Resolution. We compare with (Zhou and Choi, 2018), which reports results on resolving plural mentions on conversational data. For the comparison, we train and evaluate the Contrack baseline model on the Character mining dataset¹, which contains annotated transcripts of four seasons of the TV show *Friends*. We follow the same setup as (Zhou and Choi, 2018), computing mean and standard deviation of the BLANC score (Recasens and Hovy, 2011) over five runs. Results in Table 3 show that the Contrack model’s F1 score is lower by about eight points, which indicates that the baseline is more competitive on plural than on singular coreference resolution.

Entity Linking. Zhou and Choi (2018) also discusses entity linking, which differs from coreference resolution in that the model needs to assign each mention to one of the given known entities (that is, TV show characters in the Character Mining dataset). This is conceptually close to Context Tracking, but differs in scope (only people mentions), setup (entities are known in advance) and difficulty (model uses whole conversation). We compare the baseline model with the entity linking implementation in (Zhou and Choi, 2018) as described in the previous section and report F1 scores in table 3. The F1 scores are worse by 9 points, similar to plural coreference resolution. On this dataset most of the loss is caused by plural references whose F1 score is fifteen points worse than the one reported by Zhou and Choi (2018).

6. Conclusion

In this paper we introduce a new machine learning framework which tracks the entities in social open-domain conversations turn by turn, thereby building a repository of rich semantic data on the entities, their properties, and entity relationships for plural mentions. To enable research on the merits of this framework, we release the Contrack dataset and describe the implementation of a baseline model. While the experimental results are encouraging, it is clear that more work is necessary to turn context tracking into a mature practically useful service.

There are a few directions for future research: For one, we plan to extend the dataset with more conversations and more annotations for multiple domains. This would allow us to track a more complete set of entities more reliably. On the other hand, more research is necessary to investigate which neural network architectures are best suited for accurate context tracking.

¹The publicly available code for the system described in (Zhou and Choi, 2018) is missing some parts, so we were not able to train it on Contrack data.

7. Bibliographical References

- Adiwardana, D., Luong, M.-T., So, D. R., Hall, J., Fiedel, N., Thoppilan, R., Yang, Z., Kulshreshtha, A., Nemade, G., Lu, Y., et al. (2020). Towards a human-like open-domain chatbot. *arXiv preprint arXiv:2001.09977*.
- Bojanowski, P., Grave, E., Joulin, A., and Mikolov, T. (2017). Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5:135–146.
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., and Amodei, D. (2020). Language models are few-shot learners. In H. Larochelle, et al., editors, *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Byrne, B., Krishnamoorthi, K., Sankar, C., Neelakantan, A., Goodrich, B., Duckworth, D., Yavuz, S., Dubey, A., Kim, K.-Y., and Cedilnik, A. (2019). Taskmaster-1: Toward a realistic and diverse dialog dataset. 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 4516–4525, Hong Kong, China, November. Association for Computational Linguistics.
- Chen, Y.-H. and Choi, J. D. (2016). Character identification on multiparty conversation: Identifying mentions of characters in TV shows. In *Proceedings of the 17th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 90–100, Los Angeles, September. Association for Computational Linguistics.
- Chen, H. Y., Zhou, E., and Choi, J. D. (2017). Robust coreference resolution and entity linking on dialogues: Character identification on TV show transcripts. In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*, pages 216–225, Vancouver, Canada, August. Association for Computational Linguistics.
- Chen, X., Jia, S., and Xiang, Y. (2020). A review: Knowledge reasoning over knowledge graph. *Expert Systems with Applications*, 141:112948.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (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.
- Doddington, G., Mitchell, A., Przybocki, M., Ramshaw, L., Strassel, S., and Weischedel, R. (2004). The automatic content extraction (ACE) program – tasks, data, and evaluation. In *Proceedings of the Fourth International Conference on Language Resources and Evaluation (LREC’04)*, Lisbon, Portugal, May. European Language Resources Association (ELRA).
- Dunbar, R. I., Marriott, A., and Duncan, N. D. (1997). Human conversational behavior. *Human nature*, 8(3):231–246.
- El-Assady, M., Sevastjanova, R., Gipp, B., Keim, D., and Collins, C. (2017). Nerex: Named-entity relationship exploration in multi-party conversations. In *Computer Graphics Forum*, volume 36, pages 213–225. Wiley Online Library.
- Goo, C.-W., Gao, G., Hsu, Y.-K., Huo, C.-L., Chen, T.-C., Hsu, K.-W., and Chen, Y.-N. (2018). Slot-gated modeling for joint slot filling and intent prediction. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 753–757, New Orleans, Louisiana, June. Association for Computational Linguistics.
- Gopalakrishnan, K., Hedayatnia, B., Chen, Q., Gottardi, A., Kwatra, S., Venkatesh, A., Gabriel, R., Hakkani-Tür, D., and AI, A. A. (2019). Topicalchat: Towards knowledge-grounded open-domain conversations. In *INTERSPEECH*, pages 1891–1895.
- Goyal, A., Lamb, A., Zhang, Y., Zhang, S., Courville, A., and Bengio, Y. (2016). Professor forcing: A new algorithm for training recurrent networks. In *Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS’16*, page 4608–4616, Red Hook, NY, USA. Curran Associates Inc.
- Huang, X., Zhang, J., Li, D., and Li, P. (2019). Knowledge graph embedding based question answering. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining, WSDM ’19*, page 105–113, New York, NY, USA. Association for Computing Machinery.
- Huang, L., Wu, L., and Wang, L. (2020). Knowledge graph-augmented abstractive summarization with semantic-driven cloze reward. In Dan Jurafsky, et al., editors, *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 5094–5107. Association for Computational Linguistics.
- Kingma, D. P. and Ba, J. (2015). Adam: A method for stochastic optimization. In Yoshua Bengio et al., editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Lee, K., He, L., Lewis, M., and Zettlemoyer, L. (2017).

- End-to-end neural coreference resolution. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 188–197, Copenhagen, Denmark, September. Association for Computational Linguistics.
- Lee, K., He, L., and Zettlemoyer, L. (2018). Higher-order coreference resolution with coarse-to-fine inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 687–692, New Orleans, Louisiana, June. Association for Computational Linguistics.
- Li, X., Tur, G., Hakkani-Tür, D., and Li, Q. (2014). Personal knowledge graph population from user utterances in conversational understanding. In *2014 IEEE Spoken Language Technology Workshop (SLT)*, pages 224–229. IEEE.
- Li, Y., Su, H., Shen, X., Li, W., Cao, Z., and Niu, S. (2017). DailyDialog: A manually labelled multi-turn dialogue dataset. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 986–995, Taipei, Taiwan, November. Asian Federation of Natural Language Processing.
- Liu, Z., Winata, G. I., Xu, P., and Fung, P. (2020). Coach: A coarse-to-fine approach for cross-domain slot filling. In Dan Jurafsky, et al., editors, *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 19–25. Association for Computational Linguistics.
- Martin, L., Muller, B., Ortiz Suárez, P. J., Dupont, Y., Romary, L., de la Clergerie, É., Seddah, D., and Sagot, B. (2020a). CamemBERT: a tasty French language model. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7203–7219, Online, July. Association for Computational Linguistics.
- Martin, S., Poddar, S., and Upasani, K. (2020b). MuDoCo: Corpus for multidomain coreference resolution and referring expression generation. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 104–111, Marseille, France, May. European Language Resources Association.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., and Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In C. J. C. Burges, et al., editors, *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc.
- Mrkšić, N., Séaghdha, D. Ó., Wen, T.-H., Thomson, B., and Young, S. (2017). Neural belief tracker: Data-driven dialogue state tracking. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1777–1788.
- Pradhan, S., Moschitti, A., Xue, N., Uryupina, O., and Zhang, Y. (2012). CoNLL-2012 shared task: Modeling multilingual unrestricted coreference in OntoNotes. In *Joint Conference on EMNLP and CoNLL - Shared Task*, pages 1–40, Jeju Island, Korea, July. Association for Computational Linguistics.
- Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., and Liu, P. J. (2020). Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21:1–67.
- Rastogi, A., Hakkani-Tür, D., and Heck, L. (2017). Scalable multi-domain dialogue state tracking. In *2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 561–568. IEEE.
- Recasens, M. and Hovy, E. (2011). Blanc: Implementing the rand index for coreference evaluation. *Nat. Lang. Eng.*, 17(4):485–510, October.
- Reddy, S., Chen, D., and Manning, C. D. (2019). Coqa: A conversational question answering challenge. *Transactions of the Association for Computational Linguistics*, 7:249–266.
- Sevgili, O., Shelmanov, A., Arkhipov, M., Panchenko, A., and Biemann, C. (2020). Neural entity linking: A survey of models based on deep learning. *arXiv preprint arXiv:2006.00575*.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L. u., and Polosukhin, I. (2017). Attention is all you need. In I. Guyon, et al., editors, *Advances in Neural Information Processing Systems*, volume 30, pages 5998–6008. Curran Associates, Inc.
- Yadav, V. and Bethard, S. (2018). A survey on recent advances in named entity recognition from deep learning models. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2145–2158, Santa Fe, New Mexico, USA, August. Association for Computational Linguistics.
- Zhang, S., Dinan, E., Urbanek, J., Szlam, A., Kiela, D., and Weston, J. (2018). Personalizing dialogue agents: I have a dog, do you have pets too? In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2204–2213, Melbourne, Australia, July. Association for Computational Linguistics.
- Zhang, Y., Sun, S., Galley, M., Chen, Y.-C., Brockett, C., Gao, X., Gao, J., Liu, J., and Dolan, B. (2020). DIALOGPT: Large-scale generative pre-training for conversational response generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 270–278, Online, July. Association for Computational Linguistics.
- Zhou, E. and Choi, J. D. (2018). They exist! introducing plural mentions to coreference resolution and entity linking. In *Proceedings of the 27th Inter-*

national Conference on Computational Linguistics,
pages 24–34.

8. Language Resource References

Ralph Weischedel et al. (2013). *OntoNotes Release 5.0*. Linguistic Data Consortium, ISLRN 151-738-649-048-2.