

ECONET: Effective Continual Pretraining of Language Models for Event Temporal Reasoning

Rujun Han¹ Xiang Ren¹ Nanyun Peng^{1,2}

¹University of Southern California ²University of California, Los Angeles
{rujunhan, xiangren}@usc.edu; violetpeng@cs.ucla.edu

Abstract

While pre-trained language models (PTLMs) have achieved noticeable success on many NLP tasks, they still struggle for tasks that require event temporal reasoning, which is essential for event-centric applications. We present a continual pre-training approach that equips PTLMs with targeted knowledge about event temporal relations. We design self-supervised learning objectives to recover masked-out event and temporal indicators and to discriminate sentences from their corrupted counterparts (where event or temporal indicators got replaced). By further pre-training a PTLM with these objectives jointly, we reinforce its attention to event and temporal information, yielding enhanced capability on event temporal reasoning. This **E**ffective **C**ontinual pre-training framework for **E**vent **T**emporal reasoning (*ECONET*) improves the PTLMs’ fine-tuning performances across five relation extraction and question answering tasks and achieves new or on-par state-of-the-art performances in most of our downstream tasks.¹

1 Introduction

Reasoning event temporal relations is crucial for natural language understanding, and facilitates many real-world applications, such as tracking biomedical histories (Sun et al., 2013; Bethard et al., 2015, 2016, 2017), generating stories (Yao et al., 2019; Goldfarb-Tarrant et al., 2020), and forecasting social events (Li et al., 2020; Jin et al., 2020). In this work, we study two prominent event temporal reasoning tasks as shown in Figure 1: event relation extraction (ERE) (Chambers et al., 2014; Ning et al., 2018; O’Gorman et al., 2016; Mostafazadeh et al., 2016) that predicts temporal relations between a pair of events, and machine reading comprehension (MRC) (Ning et al., 2020; Zhou et al., 2019) where a passage and a question

¹Reproduction code, training data and models are available here: <https://github.com/PlusLabNLP/ECONET>.

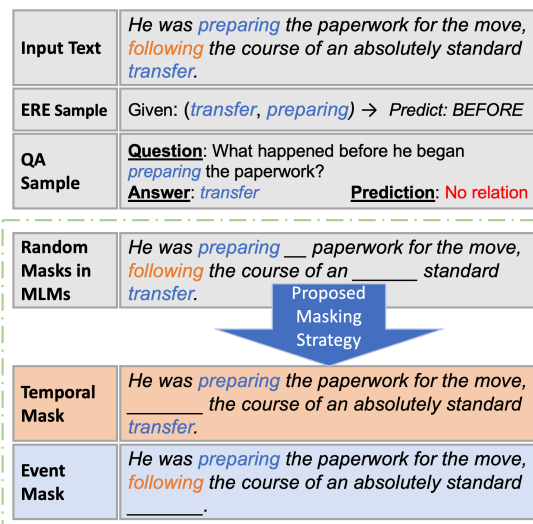


Figure 1: **Top:** an example illustrating the difference between ERE and QA / MRC samples of event temporal reasoning. **Bottom:** our targeted masking strategy for *ECONET* v.s. random masking in PTLMs.

about event temporal relations is presented, and models need to provide correct answers using the information in a given passage.

Recent approaches leveraging large pre-trained language models (PTLMs) achieved state-of-the-art results on a range of event temporal reasoning tasks (Ning et al., 2020; Pereira et al., 2020; Wang et al., 2020; Zhou et al., 2020c; Han et al., 2019b). Despite the progress, vanilla PTLMs do not focus on capturing **event temporal knowledge** that can be used to infer event relations. For example, in Figure 1, an annotator of the QA sample can easily infer from the temporal indicator “following” that “transfer” happens *BEFORE* “preparing the paperwork”, but a fine-tuned RoBERTa model predicts that “transfer” has no such relation with the event “preparing the paperwork.” Plenty of such cases exist in our error analysis on PTLMs for event temporal relation-related tasks. We hypothesize that such deficiency is caused by original PTLMs’ random masks in the pre-training where temporal indicators and event triggers are under-weighted and hence not attended well enough for our downstream tasks.

Category	Words
[before]	before, until, previous to, prior to, preceding, followed by
[after]	after, following, since, now that soon after, once**
[during]	during, while, when, at the time, at the same time, meanwhile
[past]	earlier, previously, formerly, yesterday, in the past, last time
[future]	consequently, subsequently, in turn, henceforth, later, then
[beginning]	initially, originally, at the beginning to begin, starting with, to start with
[ending]	finally, in the end, at last, lastly

Table 1: The full list of the temporal lexicon. Categories are created based on authors’ domain knowledge and best judgment. ** ‘once’ can be also placed into [past] category due to its second meaning of ‘previously’, which we exclude to keep words unique.

TacoLM (Zhou et al., 2020a) explored the idea of targeted masking and predicting textual cues of event frequency, duration and typical time, which showed improvements over vanilla PTLMs on related tasks. However, event frequency, duration and time do not directly help machines understand pairwise event temporal relations. Moreover, the mask prediction loss of TacoLM leverages a soft cross-entropy objective, which is manually calibrated with external knowledge and could inadvertently introduce noise in the continual pre-training.

We propose *ECONET*, a continual pre-training framework combining mask prediction and contrastive loss using our masked samples. Our targeted masking strategy focuses only on event triggers and temporal indicators as shown in Figure 1. This design assists models to concentrate on events and temporal cues, and potentially strengthen models’ ability to understand event temporal relations better in the downstream tasks. We further pre-train PTLMs with the following objectives jointly: the mask prediction objective trains a generator that recovers the masked temporal indicators or events, and the contrastive loss trains a discriminator that shares the representations with the generator and determines whether a predicted masked token is corrupted or original (Clark et al., 2020). Our experiments demonstrate that *ECONET* is effective at improving the original PTLMs’ performances on event temporal reasoning.

We briefly summarize our contributions. 1) We propose *ECONET*, a novel continual pre-training framework that integrates targeted masking and contrastive loss for event temporal reasoning. 2)

Our training objectives effectively learn from the targeted masked samples and inject richer event temporal knowledge in PTLMs, which leads to stronger fine-tuning performances over five widely used event temporal commonsense tasks. In most target tasks, *ECONET* achieves SOTA results in comparison with existing methods. 3) Compared with full-scale pre-training, *ECONET* requires a much smaller amount of training data and can cope with various PTLMs such as BERT and RoBERTa. 4) In-depth analysis shows that *ECONET* successfully transfers knowledge in terms of textual cues of event triggers and relations into the target tasks, particularly under low-resource settings.

2 Method

Our proposed method aims at addressing the issue in vanilla PTLMs that event triggers and temporal indicators are not adequately attended for our downstream event reasoning tasks. To achieve this goal, we propose to replace the random masking in PTLMs with a targeted masking strategy designed specifically for event triggers and temporal indicators. We also propose a continual pre-training method with mask prediction and contrastive loss that allows models to effectively learn from the targeted masked samples. The benefits of our method are manifested by stronger fine-tuning performances over downstream ERE and MRC tasks.

Our overall approach *ECONET* consists of three components. 1) Creating targeted self-supervised training data by masking out temporal indicators and event triggers in the input texts; 2) leveraging mask prediction and contrastive loss to continually train PTLMs, which produces an event temporal knowledge aware language model; 3) fine-tuning the enhanced language model on downstream ERE and MRC datasets. We will discuss each of these components in the following subsections.

2.1 Pre-trained Masked Language Models

The current PTLMs such as BERT (Devlin et al., 2018) and RoBERTa (Liu et al., 2019) follow a random masking strategy. Figure 1 shows such an example where random tokens / words are masked from the input sentences. More formally, let $\mathbf{x} = [x_1, \dots, x_n]$ be a sequence of input tokens and $x_t^m \in \mathbf{x}^m$ represents random masked tokens. The per-sample pre-training objective is to predict the identity (x_t) of x_t^m with a cross-entropy loss,

$$\mathcal{L}_{MLM} = - \sum_{x_t^m \in \mathbf{x}^m} \mathcal{I}[x_t^m = x_t] \log(p(x_t^m | \mathbf{x})) \quad (1)$$

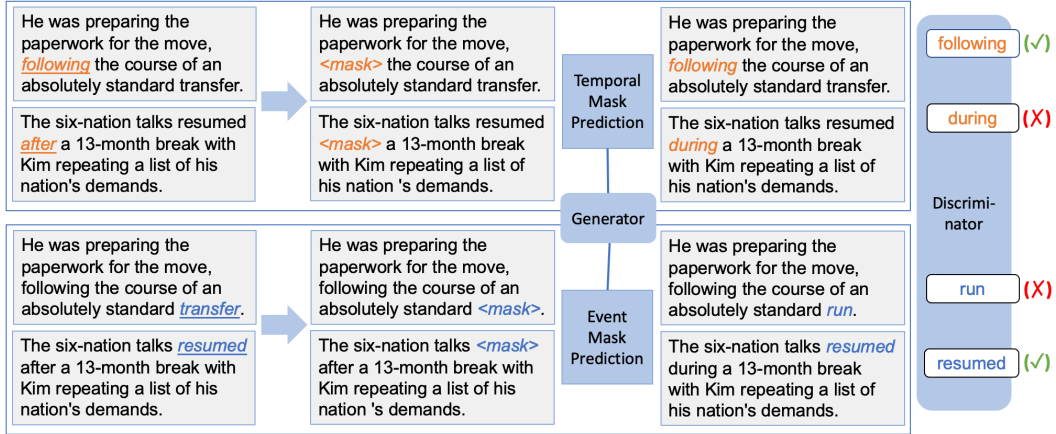


Figure 2: The proposed generator-discriminator (*ECONE*) architecture for event temporal reasoning. The upper block is the mask prediction task for temporal indicators and the bottom block is the mask prediction task for events. Both generators and the discriminator share the same representations.

Next, we will discuss the design and creation of targeted masks, training objectives and fine-tuning approaches for different tasks.

2.2 Targeted Masks Creation

Temporal Masks. We first compile a lexicon of 40 common temporal indicators listed in the Table 1 based on previous error analysis and expert knowledge in the target tasks. Those indicators in the [before], [after] and [during] categories can be used to represent the most common temporal relations between events. The associated words in each of these categories are synonyms of each other. Temporal indicators in the [past], [future], [beginning] and [ending] categories probably do not represent pairwise event relations directly, but predicting these masked tokens may still be helpful for models to understand time anchors and hence facilitates temporal reasoning.

With the temporal lexicon, we conduct string matches over the 20-year’s New York Times news articles² and obtain over 10 million 1-2 sentence passages that contain at least 1 temporal indicators. Finally, we replace each of the matched temporal indicators with a **<mask>** token. The upper block in Figure 2 shows two examples where “following” and “after” are masked from the original texts.

Event Masks. We build highly accurate event detection models (Han et al., 2019c; Zhang et al., 2021) to automatically label event trigger words in the 10 million passages mentioned above. Similarly, we replace these events with **<mask>** tokens. The bottom block in Figure 2 shows two examples where events “transfer” and “resumed” are masked from the original texts.

²NYT news articles are public from 1987-2007.

2.3 Generator for Mask Predictions

To learn effectively from the targeted samples, we train two generators with shared representations to recover temporal and event masks.

Temporal Generator. The per-sample temporal mask prediction objective is computed using cross-entropy loss,

$$\mathcal{L}_{\mathcal{T}} = - \sum_{x_t^{\mathcal{T}} \in \mathcal{x}^{\mathcal{T}}} \mathcal{I}[x_t^{\mathcal{T}} = x_t] \log(p(x_t^{\mathcal{T}} | \mathbf{x})) \quad (2)$$

where $p(x_t^{\mathcal{T}} | \mathbf{x}) = \text{Softmax}(f_{\mathcal{T}}(h_G(\mathbf{x})_t))$ and $x_t^{\mathcal{T}} \in \mathcal{x}^{\mathcal{T}}$ is a masked temporal indicator. $h_G(\mathbf{x})$ is \mathbf{x} ’s encoded representation using a transformer and $f_{\mathcal{T}}$ is a linear layer module that maps the masked token representation into label space \mathcal{T} consisting of the 40 temporal indicators.

Event Generator. The per-sample event mask prediction objective is also computed using cross-entropy loss,

$$\mathcal{L}_{\mathcal{E}} = - \sum_{x_t^{\mathcal{E}} \in \mathcal{x}^{\mathcal{E}}} \mathcal{I}[x_t^{\mathcal{E}} = x_t] \log(p(x_t^{\mathcal{E}} | \mathbf{x})) \quad (3)$$

where $p(x_t^{\mathcal{E}} | \mathbf{x}) = \text{Softmax}(f_{\mathcal{E}}(h_G(\mathbf{x})_t))$ and $x_t^{\mathcal{E}} \in \mathcal{x}^{\mathcal{E}}$ are masked events. $h_G(\mathbf{x})$ is the shared transformer encoder as in the temporal generator and $f_{\mathcal{E}}$ is a linear layer module that maps the masked token representation into label space \mathcal{E} which is a set of all event triggers in the data.

2.4 Discriminator for Contrastive Learning

We incorporate a discriminator that provides additional feedback on mask predictions, which helps correct errors made by the generators.

Contrastive Loss. For a masked token x_t , we design a discriminator to predict whether the recovered token by the mask prediction is original

or corrupted. As shown in Figure 2, “following” and “resumed” are predicted correctly, so they are labeled as **original** whereas “during” and “run” are incorrectly predicted and labeled as **corrupted**. We train the discriminator with a contrastive loss,

$$\mathcal{L}_D = - \sum_{x_t \in M} y \log(D(x_t|\mathbf{x})) + (1 - y) \log(1 - D(x_t|\mathbf{x}))$$

where $M = \mathbf{x}^E \cup \mathbf{x}^T$ and $D(x_t|\mathbf{x}) = \text{Sigmoid}(f_D(h_D(\mathbf{x})_t))$ and y is a binary indicator of whether a mask prediction is correct or not. h_D shares the same transformer encoder with h_G .

Perturbed Samples. Our mask predictions focus on temporal and event tokens, which are easier tasks than the original mask predictions in PTLMs. This could make the contrastive loss not so powerful as training a good discriminator requires relatively balanced original and corrupted samples. To deal with this issue, for $r\%$ of the generator’s output, instead of using the recovered tokens, we replace them with a token randomly sampled from either the temporal lexicon or the event vocabulary. We fix $r = 50$ to make original and corrupted samples nearly balanced.

2.5 Joint Training

To optimize the combining impact of all components in our model, the final training loss calculates the weighted sum of each individual loss, $\mathcal{L} = \mathcal{L}_T + \alpha \mathcal{L}_E + \beta \mathcal{L}_D$, where α and β are hyperparameters that balance different training objectives. The temporal and event masked samples are assigned a unique identifier (1 for temporal, 0 for event) so that the model knows which linear layers to feed the output of transformer into. Our overall generator-discriminator architecture resembles ELECTRA (Clark et al., 2020). However, our proposed method differs from this work in 1) we use targeted masking strategy as opposed to random masks; 2) both temporal and event generators and the discriminator, i.e. h_G and h_D share the hidden representations, but we allow task-specific final linear layers f_T , f_E and f_D ; 3) we do not train from scratch and instead continuing to train transformer parameters provided by PTLMs.

2.6 Fine-tuning on Target Tasks

After training with *ECONET*, we fine-tune the updated MLM on the downstream tasks. ERE samples can be denoted as $[P, e_i, e_j, r_{i,j}]$, where P is the passage and (e_i, e_j) is a pair of event trigger tokens in P . As Figure 3a shows, we feed (P, e_i, e_j)

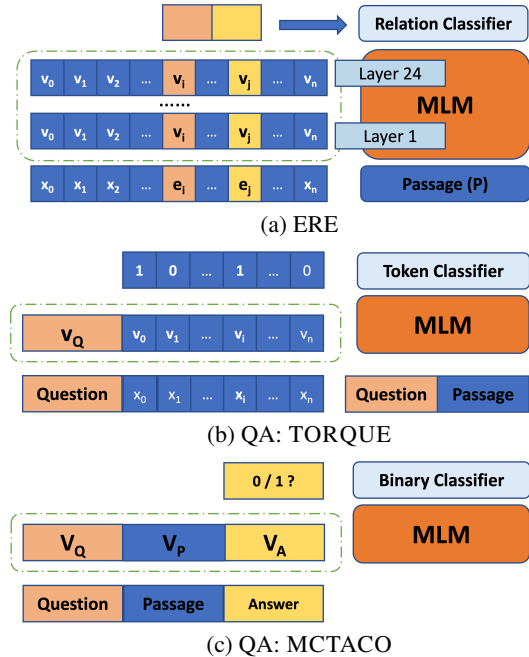


Figure 3: Target ERE and QA task illustrations.

into an MLM (trained with *ECONET*). Following the setup of Han et al. (2019a) and Zhang et al. (2021), we concatenate the final event representations v_i, v_j associated with (e_i, e_j) to predict temporal relation $r_{i,j}$. The relation classifier is implemented by a multi-layer perceptron (MLP).

MRC/QA samples can be denoted as $[P, Q, A]$, where Q represents a question and A denotes answers. Figure 3b illustrates an extractive QA task where we feed the concatenated $[P, Q]$ into an MLM. Each token $x_i \in P$ has a label with 1 indicating $x_i \in A$ and 0 otherwise. The token classifier implemented by MLP predicts labels for all x_i . Figure 3c illustrates another QA task where A is a candidate answer for the question. We feed the concatenated $[P, Q, A]$ into an MLM and the binary classifier predicts a 0/1 label of whether A is a true statement for a given question.

3 Experimental Setup

In this section, we describe details of implementing *ECONET*, datasets and evaluation metrics, and discuss compared methods reported in Section 4.

3.1 Implementation Details

Event Detection Model. As mentioned briefly in Section 2, we train a highly accurate event prediction model to mask event (triggers). We experimented with two models using event annotations in TORQUE (Ning et al., 2020) and TB-Dense (Chambers et al., 2014). These two event annotations both follow previous event-centric reasoning

research by using a trigger word (often a verb or an noun that most clearly describes the event’s occurrence) to represent an event (UzZaman et al., 2013; Glavaš et al., 2014; O’Gorman et al., 2016). In both cases, we fine-tune RoBERTa_{LARGE} on the **train set** and select models based on the performance on the **dev set**. The primary results shown in Table 2 uses TORQUE’s annotations, but we conduct additional analysis in Section 4 to show both models produce comparable results.

Continual Pretraining. We randomly selected only 200K out of 10 million samples to speed up our experiments and found the results can be as good as using a lot more data. We used half of these 200K samples for temporal masked samples and the other half for the event masked samples. We ensure none of these sample passages overlap with the target **test data**. To keep the mask tokens balanced in the two training samples, we masked only 1 temporal indicator or 1 event (closest to the temporal indicator). We continued to train BERT and RoBERTa up to 250K steps with a batch size of 8. The training process takes 25 hours on a single GeForce RTX 2080 GPU with 11G memory. Note that our method requires much fewer samples and is more computation efficient than the full-scale pre-training of language models, which typically requires multiple days of training on multiple large GPUs / TPUs.

For the generator only models reported in Table 2, we excluded the contrastive loss, trained models with a batch size of 16 to fully utilize GPU memories. We leveraged the **dev set** of TORQUE to find the best hyper-parameters.

Fine-tuning. Dev set performances were used for early-stop and average **dev performances** over three random seeds were used to pick the best hyper-parameters. Note that **test set** for the target tasks were never observed in any of the training process and their performances are reported in Table 2. All hyper-parameter search ranges can be found in Appendix C.

3.2 Datasets

We evaluate our approach on five datasets concerning temporal ERE and MRC/QA. We briefly describe these data below and list detailed statistics in Appendix A.

ERE Datasets. TB-Dense (Chambers et al., 2014), MATRES (Ning et al., 2018) and RED (O’Gorman et al., 2016) are all ERE datasets. Their

samples follow the input format described in Section 2.6 where a pair of event (triggers) together with their context are provided. The task is to predict pairwise event temporal relations. The differences are how temporal relation labels are defined. Both TB-Dense and MATRES leverage a *VAGUE* label to capture relations that are hard to determine even by humans, which results in denser annotations than RED. RED contains the most fine-grained temporal relations and thus the lowest sample/relation ratio. MATRES only considers start time of events to determine their temporal order, whereas TB-Dense and RED consider start and end time, resulting in lower inter-annotator agreement.

TORQUE (Ning et al., 2020) is an MRC/QA dataset where annotators first identify event triggers in given passages and then ask questions regarding event temporal relations (ordering). Correct answers are event trigger words in passages. TORQUE can be considered as reformulating temporal ERE tasks as an MRC/QA task. Therefore, both ERE datasets and TORQUE are highly correlated with our continual pre-training objectives where targeted masks of both events and temporal relation indicators are incorporated.

MCTACO (Zhou et al., 2019) is another MRC/QA dataset, but it differs from TORQUE in 1) events are not explicitly identified; 2) answers are statements with true or false labels; 3) questions contain broader temporal commonsense regarding not only temporal ordering, but also event frequency, during and typical time that may not be directly helpful for reasoning temporal relations. For example, knowing how often a pair of events happen doesn’t help us figure out which event happens earlier. Since our continual pre-training focuses on temporal relations, MCTACO could be the least compatible dataset in our experiments.

3.3 Evaluation Metrics

Three metrics are used to evaluate the fine-tuning performances.

F₁: for TORQUE and MCTACO, we follow the data papers (Ning et al., 2020) and (Zhou et al., 2019) to report macro average of each question’s F_1 score. For TB-Dense, MATRES and RED, we report standard micro-average F_1 scores to be consistent with the baselines.

Exact-match (EM): for both MRC datasets, EM = 1 if answer predictions match perfectly with gold annotations; otherwise, EM = 0.

Methods	TORQUE			MCTACO		TB-Dense	MATRES	RED
	F ₁	EM	C	F ₁	EM	F ₁	F ₁	F ₁
TacoLM	65.4(±0.8)	37.1(±1.0)	21.0(±0.8)	69.3(±0.6)	40.5(±0.5)	64.8(±0.7)	70.9(±0.3)	40.3(±1.7)
BERT_{LARGE} + ECUNET	70.6(±1.2) 71.4(±0.7)	43.7(±1.6) 44.8(±0.4)	27.5(±1.2) 28.5(±0.5)	70.3(±0.9) 69.2(±0.9)	43.2(±0.6) 42.3(±0.5)	62.8(±1.4) 63.0(±0.6)	70.5(±0.9) 70.4(±0.9)	39.4(±0.6) 40.2(±0.8)
RoBERTa_{LARGE} + Generator + ECUNET	75.1(±0.4) 75.8(±0.4) 76.1(±0.2)	49.6(±0.5) 51.2(±1.1) 51.6(±0.4)	35.3(±0.8) 35.8(±0.9) 36.8(±0.2)	75.5(±1.0) 75.1(±1.4) 76.3(±0.3)	50.4(±0.9) 50.2(±1.2) 52.8(±1.9)	62.8(±0.3) 65.2(±0.6) 64.8(±1.4)	78.3(±0.5) 77.0(±0.9) 78.8(±0.6)	39.4(±0.4) 41.0(±0.6) 42.8(±0.7)
<i>ECUNET</i> (best)	76.3	52.0	37.0	76.8	54.7	66.8	79.3	43.8
Current SOTA	75.2*	51.1	34.5	79.5 [†]	56.5	66.7 ^{††}	80.3 [‡]	34.0 ^{‡‡}

Table 2: Overall experimental results. Refer to Section 3.4 for naming conventions. The SOTA performances for TORQUE* are provided by Ning et al. (2020) and the numbers are average over 3 random seeds. The SOTA performances for MCTACO[†] are provided by Pereira et al. (2020); TB-Dense^{††} and MATRES[‡] by Zhang et al. (2021) and RED^{‡‡} by Han et al. (2019b). [†], ^{††}, [‡] and ^{‡‡} only report the best single model results, and to make fair comparisons with these baselines, we report both average and best single model performances. TacoLM baseline uses the provided and recommended checkpoint for extrinsic evaluations.

EM-consistency (C): in TORQUE, some questions can be clustered into the same group due to the data collection process. This metric reports the average EM score for a group as opposed to a question in the original EM metrics.

3.4 Compared Methods

We compare several pre-training methods with *ECUNET*: 1) **RoBERTa_{LARGE}** is the original PTLM and we fine-tune it directly on target tasks; 2) **RoBERTa_{LARGE} + ECUNET** is our proposed continual pre-training method; 3) **RoBERTa_{LARGE} + Generator** only uses the generator component in continual pre-training; 4) **RoBERTa_{LARGE} + random mask** keeps the original PTLMs’ objectives and replaces the targeted masks in *ECUNET* with randomly masked tokens. The methods’ names for continual pre-training BERT_{LARGE} can be derived by replacing RoBERTa_{LARGE} with BERT_{LARGE}.

We also fine-tune pre-trained TacoLM on target datasets. The current SOTA systems we compare with are provided by Ning et al. (2020), Pereira et al. (2020), Zhang et al. (2021) and Han et al. (2019b). More details are presented in Section 4.1.

4 Results and Analysis

As shown in Table 2, we report two baselines. The first one, TacoLM is a related work that focuses on event duration, frequency and typical time. The second one is the current SOTA results reported to the best of the authors’ knowledge. We also report our own implementations of fine-tuning BERT_{LARGE} and RoBERTa_{LARGE} to compare fairly with *ECUNET*. Unless pointing out specifically, all gains mentioned in the following sections are in the unit of **absolute percentage**.

4.1 Comparisons with Existing Systems

TORQUE. The current SOTA system reported in Ning et al. (2020) fine-tunes RoBERTa_{LARGE} and our own fine-tuned RoBERTa_{LARGE} achieves on-par F₁, EM and C scores. The gains of RoBERTa_{LARGE} + *ECUNET* against the current SOTA performances are 0.9%, 0.5% and 2.3% per F₁, EM and C metrics.

MCTACO. The current SOTA system ALICE (Pereira et al., 2020) also uses RoBERTa_{LARGE} as the text encoder, but leverages adversarial attacks on input samples. ALICE achieves 79.5% and 56.5% per F₁ and EM metrics on the test set for the best single model, and the best performances for RoBERTa_{LARGE} + *ECUNET* are 76.8% and 54.7% per F₁ and EM scores, which do not outperform ALICE. This gap can be caused by the fact that the majority of samples in MCTACO reason about event frequency, duration and time, which are not directly related to event temporal relations.

TB-Dense + MATRES. The most recent SOTA system reported in Zhang et al. (2021) uses both BERT_{LARGE} and RoBERTa_{LARGE} as text encoders, but leverages syntactic parsers to build large graphical attention networks on top of PTLMs. RoBERTa_{LARGE} + *ECUNET*’s fine-tuning performances are essentially on-par with this work without additional parameters. For TB-Dense, our best model outperforms Zhang et al. (2021) by 0.1% while for MATRES, our best model underperforms by 1.0% per F₁ scores.

RED. The current SOTA system reported in Han et al. (2019b) uses BERT_{BASE} as word representations (no finetuning) and BiLSTM as feature extractor. The single best model achieves 34.0% F₁ score and RoBERTa_{LARGE} + *ECUNET* is 9.8% higher than the baseline.

4.2 The Impact of *ECONET*

Overall Impact. *ECONET* in general works better than the original RoBERTa_{LARGE} across 5 different datasets, and the improvements are more salient in TORQUE with 1.0%, 2.0% and 1.5% gains per F_1 , EM and C scores, in MCTACO with 2.4% lift over the EM score, and in TB-Dense and RED with 2.0% and 3.4% improvements respectively over F_1 scores. We observe that the improvements of *ECONET* over BERT_{LARGE} is smaller and sometimes hurts the fine-tuning performances. We speculate this could be related to the property that BERT is less capable of handling temporal reasoning tasks, but we leave more rigorous investigations to future research.

Impact of Contrastive Loss. Comparing the average performances of continual pre-training with generator only and with *ECONET* (generator + discriminator), we observe that generator alone can improve performances of RoBERTa_{LARGE} in 3 out of 5 datasets. However, except for TB-Dense, *ECONET* is able to improve fine-tuning performances further, which shows the effectiveness of using the contrastive loss.

Significance Tests. As current SOTA models are either not publicly available or under-perform RoBERTa_{LARGE}, we resort to testing the statistical significance of the best single model between *ECONET* and RoBERTa_{LARGE}. Table 8 in the appendix lists all improvements’ p-values per McNemar’s test (McNemar, 1947). MATRES appears to be the only one that is not statistically significant.

4.3 Impact of Event Models

Event trigger definitions have been consistent in previous event temporal datasets (O’Gorman et al., 2016; Chambers et al., 2014; Ning et al., 2020). Trigger detection models built on TORQUE and TB-Dense both achieve $> 92\%$ F_1 scores and $> 95\%$ precision scores. For the 100K pre-training data selected for event masks, we found an 84.5% overlap of triggers identified by both models. We further apply *ECONET* trained on both event mask data to the target tasks and achieve comparable performances shown in Table 10 of the appendix. These results suggest that the impact of different event annotations is minimal and triggers detected in either model can generalize to different tasks.

4.4 Additional Ablation Studies

To better understand our proposed model, we experiment with additional continual training methods and compare their fine-tuning performances.

Methods	TORQUE			TB-D	RED
	F_1	EM	C	F_1	F_1
RoBERTa _{LARGE}	75.1	49.6	<u>35.3</u>	62.8	39.4
+ random mask	74.9	49.5	<u>35.1</u>	58.7	<u>38.3</u>
+ <i>ECONET</i>	76.1	51.6	36.8	64.8	42.8

Table 3: Fine-tuning performances with different pre-training methods. All numbers are average over 3 random seeds. Std. Dev. $\geq 1\%$ is underlined.

Random Masks. As most target datasets we use are in the news domain, to study the impact of potential domain-adaption, we continue to train PTLMs with the original objective on the same data using random masks. To compare fairly with the generator and *ECONET*, we only mask 1 token per training sample. The search range of hyperparameters is the same as in Section 3. As Table 3 and 11 (appendix) show, continual pre-training with random masks, in general, does not improve and sometimes hurt fine-tuning performances compared with fine-tuning with original PTLMs. We hypothesize that this is caused by masking a smaller fraction (1 out of ≈ 50 average) tokens than the original 15%. RoBERTa_{LARGE} + *ECONET* achieves the best fine-tuning results across the board.

	Full Train Data			10% Train Data		
	RoBERTa	Δ	$\Delta\%$	RoBERTa	Δ	$\Delta\%$
TORQUE	75.1	+1.0	+1.3%	59.7	+7.2	+12.1%
MCTACO	75.5	+0.8	+1.1%	44.0	+5.6	+12.7%
TB-Dense	62.8	+2.0	+3.2%	48.8	+2.8	+5.7%
MATRES	78.3	+0.5	+1.3%	71.0	+2.4	+3.4%
RED	39.4	+2.4	+6.0%	27.2	+1.8	+6.6%

Table 4: RoBERTa_{LARGE} + *ECONET*’s improvements over RoBERTa_{LARGE} using full train data v.s. 10% of train data. Δ indicates absolute points improvements while $\Delta\%$ indicates relative gains per F_1 scores.

4.5 Fine-tuning under Low-resource Settings

In Table 4, we compare the improvements of fine-tuning RoBERTa_{LARGE} + *ECONET* over RoBERTa_{LARGE} using full and 10% of the training data. Measured by both absolute and relative percentage gains, the majority of the improvements are much more significant under low-resource settings. This suggests that the transfer of event temporal knowledge is more salient when data is scarce. We further show fine-tuning performance comparisons using different ratios of the training data in Figure 6a-6b in the appendix. The results demonstrate that *ECONET* can outperform RoBERTa_{LARGE} consistently when fine-tuning TORQUE and RED.

4.6 Attention Scores on Temporal Indicators

In this section, we attempt to show explicitly how *ECONET* enhances MLMs’ attentions on temporal indicators for downstream tasks. As mentioned in





Figure 4: Cumulative attention score comparisons between RoBERTa_{LARGE} and *ECONET* on TB-Dense test data. All numbers are multiplied by 100 and averaged over 3 random seeds for illustration clarity.

Sec. 2.6, for a particular ERE task (e.g. TB-Dense), we need to predict the temporal relations between a pair of event triggers $e_i, e_j \in P_{i,j}$ with associated vector representations $v_i^{l,h}, v_j^{l,h}, l \in L, h \in H$ in an MLM. L and H are the number of layers and attention heads respectively. We further use $T_m \in T$ to denote a temporal indicator category listed in Table 1, and $t_{m,n} \in T_m$ denote a particular temporal indicator. If we let $attn(v_i^{l,h}, v_x^{l,h})$ represents the attention score between an event vector and any other hidden vectors, we can aggregate the per-layer attention score between e_i and $t_{m,n}$ as, $a_{i,t_{m,n}}^l = \frac{1}{H} \sum_h attn(v_i^{l,h}, v_{t_{m,n}}^{l,h})$. Similarly, we can compute $a_{j,t_{m,n}}^l$. The final per-layer attention score for (e_i, e_j) is $a_{t_{m,n}}^l = \frac{1}{2} (a_{i,t_{m,n}}^l + a_{j,t_{m,n}}^l)$. To compute the attention score for the T_m category, we take the average of $\{a_{t_{m,n}}^l \mid \forall t_{m,n} \in T_m \text{ and } \forall t_{m,n} \in P_{i,j}\}$. Note we assume a temporal indicator is a single token to simplify notations above; for multiple-token indicators, we take the average of $attn(v_i^{l,h}, v_{x \in t_{m,n}}^{l,h})$.

Figure 4 shows the cumulative attention scores for temporal indicator categories, **[before]**, **[after]** and **[during]** in ascending order of model layers. We observe that the attention scores for RoBERTa_{LARGE} and *ECONET* align well on the bottom layers, but *ECONET* outweighs RoBERTa_{LARGE} in middle to top layers. Previous research report that upper layers of pre-trained language models focus more on complex semantics as opposed to shallow surface forms or syntax on the lower layers (Tenney et al., 2019; Jawahar et al., 2019). Thus, our findings here show another piece of evidence that targeted masking is effective at capturing temporal indicators, which could facilitate semantics tasks including temporal reasoning.

4.7 Temporal Knowledge Injection

We hypothesize in the introduction that vanilla PTLMs lack special attention to temporal indica-



(a) Random Mask - RoBERTa_{LARGE}



(b) *ECONET* - RoBERTa_{LARGE}

Figure 5: Performance (F₁ score) differences by temporal indicator categories and label classes in TB-Dense. Fine-tuning on 10% TB-Dense training data.

tors and events, and our proposed method addresses this issue by a particular design of mask prediction strategy and a discriminator that is able to distinguish reasonable events and temporal indicators from noises. In this section, we show more details of how such a mechanism works.

The heat maps in Figure 5 calculate the fine-tuning performance differences between 1) RoBERTa_{LARGE} and continual pre-training with random masks (Figure 5a); and 2) between RoBERTa_{LARGE} and *ECONET* (Figure 5b). Each cell shows the difference for each label class in TB-Dense conditional on samples' input passage containing a temporal indicator in the categories specified in Table 1. Categories with less than 50 sample matches are excluded from the analysis.

In Figure 5a, the only gains come from *VAGUE*, which is an undetermined class in TB-Dense to handle unclear pairwise event relations. This shows that continual pre-training with random masks works no better than original PTLMs to leverage existing temporal indicators in the input passage to distinguish positive temporal relations from unclear ones. On the other hand, in Figure 5b, having temporal indicators in general benefits much more for *BEFORE*, *AFTER*, *IS_INCLUDED* labels. The only exception is *INCLUDES*, but it is a small class with only 4% of the data.

More interestingly, notice the diagonal cells, i.e. ([before], *BEFORE*), ([after], *AFTER*) and ([during], *INCLUDES*) have the largest values in the respective columns. These results are intuitive as temporal indicators should be most beneficial for temporal relations associated with their categories.

Combining these two sets of results, we provide additional evidence that *ECONET* helps PTLMs better capture temporal indicators and thus results in stronger fine-tuning performances.

Our final analysis attempts to show why discriminator helps. We feed 1K unused masked samples into the generator of the best *ECONET* in Table 2 to predict either the masked temporal indicators or masked events. We then examine the accuracy of the discriminator for correctly and incorrectly predicted masked tokens. As shown in Table 12 of the appendix, the discriminator aligns well with the event generator’s predictions. For the temporal generator, the discriminator disagrees substantially (82.2%) with the “incorrect” predictions, i.e. the generator predicts a supposedly wrong indicator, but the discriminator thinks it looks original.

To understand why, we randomly selected 50 disagreed samples and found that 12 of these “incorrect” predictions fall into the same temporal indicator group of the original ones and 8 of them belong to the related groups in Table 1. More details and examples can be found in Table 13 in the appendix. This suggests that despite being nearly perfect replacements of the original masked indicators, these 40% samples are penalized as wrong predictions when training the generator. The discriminator, by disagreeing with the generator, provides opposing feedback that trains the overall model to better capture indicators with similar temporal signals.

5 Related Work

Language Model Pretraining. Since the breakthrough of BERT (Devlin et al., 2018), PTLMs have become SOTA models for a variety of NLP applications. There have also been several modifications/improvements built on the original BERT model. RoBERTa (Liu et al., 2019) removes the next sentence prediction in BERT and trains with longer text inputs and more steps. ELECTRA (Clark et al., 2020) proposes a generator-discriminator architecture, and addresses the sample-inefficiency issue in previous PTLMs.

Recent research explored methods to continue to train PTLMs so that they can adapt better to downstream tasks. For example, TANDA (Garg et al., 2019) adopts an intermediate training on modified Natural Questions dataset (Kwiatkowski et al., 2019) so that it performs better for the Answer Sentence Selection task. Zhou et al. (2020b) proposed continual training objectives that require a model to distinguish natural sentences from those

with concepts randomly shuffled or generated by models, which enables language models to capture large-scale commonsense knowledge.

Event Temporal Reasoning. There has been a surge of attention to event temporal reasoning research recently. Some noticeable datasets include ERE samples: TB-Dense (Chambers et al., 2014), MATRES (Ning et al., 2018) and RED (O’Gorman et al., 2016). Previous SOTA systems on these data leveraged PTLMs and structured learning (Han et al., 2019c; Wang et al., 2020; Zhou et al., 2020c; Han et al., 2020) and have substantially improved model performances, though none of them tackled the issue of lacking event temporal knowledge in PTLMs. TORQUE (Ning et al., 2020) and MCTACO (Zhou et al., 2019) are recent MRC datasets that attempt to reason about event temporal relations using natural language rather than ERE formalism.

Zhou et al. (2020a) and Zhao et al. (2020) are two recent works that attempt to incorporate event temporal knowledge in PTLMs. The formal one focuses on injecting temporal commonsense with targeted event time, frequency and duration masks while the latter one leverages distantly labeled pairwise event temporal relations, masks before/after indicators, and focuses on ERE application only. Our work differs from them by designing a targeted masking strategy for event triggers and comprehensive temporal indicators, proposing a continual training method with mask prediction and contrastive loss, and applying our framework on a broader range of event temporal reasoning tasks.

6 Conclusion and Future Work

In summary, we propose a continual training framework with targeted mask prediction and contrastive loss to enable PTLMs to capture event temporal knowledge. Extensive experimental results show that both the generator and discriminator components can be helpful to improve fine-tuning performances over 5 commonly used data in event temporal reasoning. The improvements of our methods are much more pronounced in low-resource settings, which points out a promising direction for few-shot learning in this research area.

Acknowledgments

This work is supported by the Intelligence Advanced Research Projects Activity (IARPA), via Contract No. 2019-19051600007 and DARPA under agreement number FA8750-19-2-0500.

References

- Steven Bethard, Leon Derczynski, Guergana Savova, James Pustejovsky, and Marc Verhagen. 2015. [SemEval-2015 task 6: Clinical TempEval](#). In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pages 806–814, Denver, Colorado. Association for Computational Linguistics.
- Steven Bethard, Guergana Savova, Wei-Te Chen, Leon Derczynski, James Pustejovsky, and Marc Verhagen. 2016. [SemEval-2016 task 12: Clinical TempEval](#). In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 1052–1062, San Diego, California. Association for Computational Linguistics.
- Steven Bethard, Guergana Savova, Martha Palmer, and James Pustejovsky. 2017. [SemEval-2017 task 12: Clinical TempEval](#). In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 565–572, Vancouver, Canada. Association for Computational Linguistics.
- Nathanael Chambers, Taylor Cassidy, Bill McDowell, and Steven Bethard. 2014. [Dense event ordering with a multi-pass architecture](#). In *ACL*.
- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. [Electra: Pre-training text encoders as discriminators rather than generators](#). In *International Conference on Learning Representations*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*.
- Siddhant Garg, Thuy Vu, and Alessandro Moschitti. 2019. [Tanda: Transfer and adapt pre-trained transformer models for answer sentence selection](#).
- Goran Glavaš, Jan Šnajder, Marie-Francine Moens, and Parisa Kordjamshidi. 2014. [HiEve: A corpus for extracting event hierarchies from news stories](#). In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 3678–3683, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Seraphina Goldfarb-Tarrant, Tuhin Chakrabarty, Ralph Weischedel, and Nanyun Peng. 2020. Content planning for neural story generation with aristotelian rescoring. In *the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4319–4338.
- Rujun Han, I-Hung Hsu, Mu Yang, Aram Galstyan, Ralph Weischedel, and Nanyun Peng. 2019a. [Deep structured neural network for event temporal relation extraction](#). In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 666–106, Hong Kong, China. Association for Computational Linguistics.
- Rujun Han, Mengyue Liang, Bashar Alhafni, and Nanyun Peng. 2019b. Contextualized word embeddings enhanced event temporal relation extraction for story understanding. In *2019 Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL-HLT 2019), Workshop on Narrative Understanding*.
- Rujun Han, Qiang Ning, and Nanyun Peng. 2019c. [Joint event and temporal relation extraction with shared representations and structured prediction](#). 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 434–444, Hong Kong, China. Association for Computational Linguistics.
- Rujun Han, Yichao Zhou, and Nanyun Peng. 2020. [Domain knowledge empowered structured neural net for end-to-end event temporal relation extraction](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5717–5729, Online. Association for Computational Linguistics.
- Ganesh Jawahar, Benoît Sagot, and Djamel Seddah. 2019. [What does BERT learn about the structure of language?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3651–3657, Florence, Italy. Association for Computational Linguistics.
- Woojeong Jin, Suji Kim, and Xiang Ren. 2020. Forecastqa: A question answering challenge for event forecasting. *arXiv preprint arXiv:2005.00792*.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Matthew Kelcey, Jacob Devlin, Kenton Lee, Kristina N. Toutanova, Llion Jones, Ming-Wei Chang, Andrew Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. Natural questions: a benchmark for question answering research. *Transactions of the Association of Computational Linguistics*.
- Manling Li, Qi Zeng, Ying Lin, Kyunghyun Cho, Heng Ji, Jonathan May, Nathanael Chambers, and Clare Voss. 2020. [Connecting the dots: Event graph schema induction with path language modeling](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 684–695, Online. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint*, arXiv:1907.11692.
- Quinn McNemar. 1947. Note on the sampling error of the difference between correlated proportions or percentages. *Psychometrika*, 12(2):153–157.

- Nasrin Mostafazadeh, Alyson Grealish, Nathanael Chambers, James Allen, and Lucy Vanderwende. 2016. [Caters: Causal and temporal relation scheme for semantic annotation of event structures](#). In *NAACL*, San Diego, USA.
- Qiang Ning, Hao Wu, Rujun Han, Nanyun Peng, Matt Gardner, and Dan Roth. 2020. [TORQUE: A reading comprehension dataset of temporal ordering questions](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1158–1172, Online. Association for Computational Linguistics.
- Qiang Ning, Hao Wu, and Dan Roth. 2018. [A multi-axis annotation scheme for event temporal relations](#). In *ACL*. Association for Computational Linguistics.
- Tim O’Gorman, Kristin Wright-Bettner, and Martha Palmer. 2016. [Richer event description: Integrating event coreference with temporal, causal and bridging annotation](#). In *Proceedings of the 2nd Workshop on Computing News Storylines (CNS 2016)*, pages 47–56, Austin, Texas. Association for Computational Linguistics.
- Lis Pereira, Xiaodong Liu, Fei Cheng, Masayuki Asahara, and Ichiro Kobayashi. 2020. [Adversarial training for commonsense inference](#). In *Proceedings of the 5th Workshop on Representation Learning for NLP*, pages 55–60, Online. Association for Computational Linguistics.
- Weiyi Sun, Anna Rumshisky, and Ozlem Uzuner. 2013. [Evaluating temporal relations in clinical text: 2012 i2b2 challenge](#).
- Ian Tenney, Dipanjan Das, and Ellie Pavlick. 2019. [BERT rediscovers the classical NLP pipeline](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4593–4601, Florence, Italy. Association for Computational Linguistics.
- Naushad UzZaman, Hector Llorens, Leon Derczynski, James Allen, Marc Verhagen, and James Pustejovsky. 2013. [SemEval-2013 task 1: TempEval-3: Evaluating time expressions, events, and temporal relations](#). In *Second Joint Conference on Lexical and Computational Semantics (*SEM), Volume 2: Proceedings of the Seventh International Workshop on Semantic Evaluation (SemEval 2013)*, pages 1–9, Atlanta, Georgia, USA. Association for Computational Linguistics.
- Haoyu Wang, Muhao Chen, Hongming Zhang, and Dan Roth. 2020. [Joint constrained learning for event-event relation extraction](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 696–706, Online. Association for Computational Linguistics.
- Lili Yao, Nanyun Peng, Weischedel Ralph, Kevin Knight, Dongyan Zhao, and Rui Yan. 2019. [Plan-and-write: Towards better automatic storytelling](#). In *The Thirty-Third AAAI Conference on Artificial Intelligence (AAAI-19)*.
- Shuaicheng Zhang, Lifu Huang, and Qiang Ning. 2021. [Extracting temporal event relation with syntactic-guided temporal graph transformer](#). *arXiv preprint arXiv:arXiv:2104.09570*.
- Xinyu Zhao, Shih-ting Lin, and Greg Durrett. 2020. [Effective distant supervision for temporal relation extraction](#). *arXiv preprint arXiv:2010.12755*.
- Ben Zhou, Daniel Khashabi, Qiang Ning, and Dan Roth. 2019. [“going on a vacation” takes longer than “going for a walk”: A study of temporal commonsense understanding](#). 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 3363–3369, Hong Kong, China. Association for Computational Linguistics.
- Ben Zhou, Qiang Ning, Daniel Khashabi, and Dan Roth. 2020a. [Temporal Common Sense Acquisition with Minimal Supervision](#). In *Proc. of the Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Wangchunshu Zhou, Dong-Ho Lee, Ravi Kiran Selvam, Seyeon Lee, Bill Yuchen Lin, and Xiang Ren. 2020b. [Pre-training text-to-text transformers for concept-centric common sense](#).
- Yichao Zhou, Yu Yan, Rujun Han, J. Harry Caulfield, Kai-Wei Chang, Yizhou Sun, Peipei Ping, and Wei Wang. 2020c. [Clinical temporal relation extraction with probabilistic soft logic regularization and global inference](#). *arXiv preprint, arXiv:2012.08790*.

A Data Summary

Table 5 describes basic statistics for target datasets used in this work. The numbers of train/dev/test samples for TORQUE and MCTACO are question based. There is no training set provided in MCTACO. So we train on the dev set and report the evaluation results on the test set following Pereira et al. (2020). The numbers of train/dev/test samples for TB-Dense, MATRES and RED refer to (event, event, relation) triplets. The standard dev set is not provide by MATRES and RED, so we follow the split used in Zhang et al. (2021) and Han et al. (2019b).

Data	#Train	#Dev	#Test	#Label
TORQUE	24,523	1,483	4,668	2
MCTACO	-	3,783	9,442	2
TB-Dense	4,032	629	1,427	6
MATRES	5,412	920	827	4
RED	3,336	400	473	11

Table 5: The numbers of samples for TORQUE refers to number of questions; the numbers for MCTACO are valid question and answer pairs; and the numbers of samples for TB-Dense, MATRES and RED are all (event, event, relation) triplets.

Downloading link for the (processed) continual pretraining data is provided in the README file of the code package.

B Event Detection Model

As mentioned briefly in Secion 2, we train an event prediction model using event annotations in TORQUE. We finetune RoBERTa_{LARGE} on the training set and select models based on the performance on the dev set. The best model achieves > 92% event prediction F_1 score with > 95% precision score after just 1 epoch of training, which indicates that this is a highly accurate model.

C Reproduction Checklist

Number of parameters. We continue to train BERT_{LARGE} and RoBERTa_{LARGE} and so the number of parameters are the same as the original PTLMs, which is 336M.

Hyper-parameter Search Due to computation constraints, we had to limit the search range of hyper-parameters for *ECONET*. For learning rates, we tried ($1e^{-6}$, $2e^{-6}$); for weights on the contrastive loss (β), we tried (1.0, 2.0).

Best Hyper-parameters. In Table 6 and Table 7, we provide hyper-parameters for our best performing language model using RoBERTa_{LARGE} + *ECONET* and BERT_{LARGE} + *ECONET* and best hyper-parameters for fine-tuning them on downstream tasks. For fine-tuning on the target datasets. We conducted grid search for learning rates in the range of ($5e^{-6}$, $1e^{-5}$) and for batch size in the range of (2, 4, 6, 12). We fine-tuned all models for 10 epochs with three random seeds (5, 7, 23).

Method	learning rate	batch size	β
<i>ECONET</i>	$1e^{-6}$	8	1.0
TORQUE	$1e^{-5}$	12	-
MCTACO	$5e^{-6}$	4	-
TB-Dense	$5e^{-6}$	4	-
MATRES	$5e^{-6}$	2	-
RED	$5e^{-6}$	2	-

Table 6: Hyper-parameters of our best performing LM with RoBERTa_{LARGE} + *ECONET* as well as best hyper-parameters for fine-tuning on downstream tasks.

Method	learning rate	batch size	β
<i>ECONET</i>	$2e^{-6}$	8	1.0
TORQUE	$1e^{-5}$	12	-
MCTACO	$1e^{-5}$	2	-
TB-Dense	$1e^{-5}$	2	-
MATRES	$5e^{-6}$	4	-
RED	$1e^{-5}$	6	-

Table 7: Hyper-parameters of our best performing LM with BERT_{LARGE} + *ECONET* as well as best hyper-parameters for fine-tuning on downstream tasks.

Dev Set Performances We show average dev set performances in Table 9 corresponding to our main results in Table 2.

D Significance Tests.

We leverage McNemar’s tests (McNemar, 1947) to show *ECONET*’s improvements against RoBERTa_{LARGE}. McNemar’s tests compute statistics by aggregating all samples’ prediction correctness. For ERE tasks, this value is simply classification correctness; for QA tasks (TORQUE and MCTACO), we use EM per question-answer pairs.

E Fine-tuning under Low-resource Settings

Table 4 shows improvements of *ECONET* over RoBERTa_{LARGE} are much more salient under low-

Datasets	p-values
TORQUE	0.002**
MCTACO	0.007**
TB-Dense	0.004**
MATRES	0.292
RED	0.059*

Table 8: McNemar’s tests for improvement significance between best single models of RoBERTa_{LARGE} and *ECONET* on the test data. Tests with p-values < 0.05 (**) indicate strong statistical significance; tests with p-values < 0.1 (*) indicate weak statistical significance.

resource setting.




Figure 6: Performances (F₁ scores) comparison between fine-tuning RoBERTa_{LARGE} vs. RoBERTa_{LARGE} + *ECONET* over different ratios of the training data.

F Variants of *ECONET*

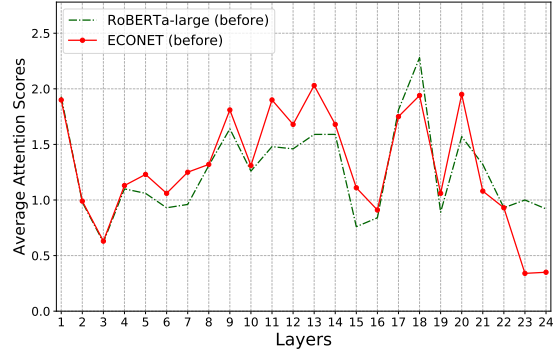
We also experimented with a variant of *ECONET* by first pretraining RoBERTa_{LARGE} + Generator for a few thousands steps and then continue to pre-train with *ECONET*. However, this method leads worse finetuning results, which seems to contradict the suggestions in Zhou et al. (2020b) and Clark et al. (2020) that the generator needs to be first trained to obtain a good prediction distribution for the discriminator. We speculate that this is due to our temporal and event mask predictions being easier tasks than those in the previous work, which makes the “warm-up steps” for the generator not necessary.

G Impact of Event Models

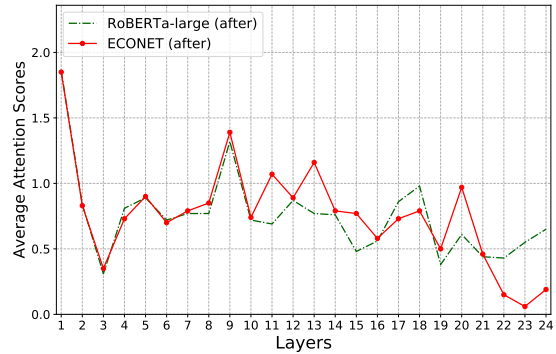
Table 10 compares results based on two event annotations.

H Ablation Studies for BERT


I Attention Scores



(a) Attentions scores for [before] indicators.



(b) Attentions scores for [after] indicators.



(c) Attentions scores for [during] indicators.

Figure 7: Attentions score comparisons between RoBERTa_{LARGE} and *ECONET* for all model layers. All numbers are multiplied by 100 and average over 3 random seeds for illustration purpose

J Analysis of the Discriminator

Table 12 shows the alignment of between the generator and the discriminator, and Table 13 shows the examples of “disagreed” samples between the

Methods	TORQUE			TB-Dense	MATRES	RED
	F ₁	EM	C	F ₁	F ₁	F ₁
TacoLM	65.5(±0.8)	36.2(±1.6)	22.7(±1.4)	56.9(±0.7)	75.1(±0.8)	40.7(±0.3)
BERT _{LARGE}	70.9(±1.0)	42.8(±1.7)	29.0(±0.9)	56.7(±0.2)	73.9(±1.0)	40.6(±0.0)
+ <i>ECONET</i>	71.8(±0.3)	44.8(±0.4)	31.1(±1.6)	56.1(±0.8)	74.4(±0.4)	41.7(±1.5)
RoBERTa _{LARGE}	76.7(±0.2)	50.5(±0.5)	36.2(±1.1)	59.8(±0.3)	79.9(±1.0)	43.7(±1.0)
+ Generator	76.6(±0.1)	51.0(±1.0)	36.3(±0.8)	61.5(±0.9)	79.8(±0.9)	43.2(±1.2)
+ <i>ECONET</i>	76.9(±0.4)	52.2(±0.9)	37.7(±0.4)	60.8(±0.6)	79.5(±0.1)	44.1(±1.7)

Table 9: Average Dev Performances Corresponding to Table 2. Note that for MCTACO, we train on dev set and evaluate on the test set as mentioned in Section 3, so we do not report test performance again here.

Event Annotations	TORQUE			TB-D	RED
	F ₁	EM	C	F ₁	F ₁
TORQUE	76.1	51.6	36.8	<u>64.8</u>	42.8
TB-Dense	76.1	51.3	36.4	65.1	<u>42.6</u>

Table 10: Fine-tuning performance comparisons using event detection models trained on TORQUE v.s. TB-Dense event annotations. All numbers are average over 3 random seeds. Std. Dev. $\geq 1\%$ is underlined.

Methods	TORQUE			TB-D	RED
	F ₁	EM	C	F ₁	F ₁
BERT _{LARGE}	<u>70.6</u>	<u>43.7</u>	<u>27.5</u>	<u>62.8</u>	39.4
+ random mask	70.6	44.1	27.2	63.4	<u>35.3</u>
+ <i>ECONET</i>	71.4	44.8	28.5	63.0	40.2

Table 11: Fine-tuning performances. All numbers are average over 3 random seeds. Std. Dev. $\geq 1\%$ is underlined.

generator and the discriminator. Detailed analysis can be found in Section 4 in the main text.

	Temporal Generator		Event Generator	
	Corr.	Incorr.	Corr.	Incorr.
Total #	837	163	26	974
Discr. Corr. #	816	29	25	964
Accuracy	97.5%	17.8%	96.2%	99.0%

Table 12: Discriminator’s alignment with generator’s mask predictions in *ECONET*. Second column shows that discriminator strongly disagree with the “errors” made by the temporal generator.

Type I. Same Group: 12/50 (24%)
>> Ex 1. original: when; predicted: while
Text: A letter also went home a week ago in Pelham, in Westchester County, New York, (mask) a threat made by a student in a neighboring town circulated in several communities within hours...
>> Ex 2. original: prior to; predicted: before
Text: ... An investigation revealed that rock gauges were picking up swifter rates of salt movement in the ceiling of the room, but at Wipp no one had read the computer printouts for at least one month (mask) the collapse.
Type II. Related Group: 8/50 (16%)
>> Ex 3. original: in the past; predicted: before
text: Mr. Douglen confessed that Lautenberg, which had won (mask) , was “a seasoned roach and was ready for this race...”
>> Ex 4. original: previously; predicted: once
text: Under the new legislation enacted by Parliament, divers who (mask) had access to only 620 miles of the 10,000 miles of Greek coast line will be able to explore ships and “archaeological parks” freely...

Table 13: Categories and examples of highly related “incorrect” temporal indicator predictions by the generator, but labeled as “correct” by the discriminator.