

Time-Aware Language Models as Temporal Knowledge Bases

Bhuwan Dhingra^{***} Jeremy R. Cole^{*} Julian Martin Eisenschlos Daniel Gillick
Jacob Eisenstein William W. Cohen

Google Research

{bdhingra, jrcole, eisenjulian, dgillick, jeisenstein, wcohen}@google.com

Abstract

Many facts come with an expiration date, from the name of the President to the basketball team LeBron James plays for. However, most language models (LMs) are trained on snapshots of data collected at a specific moment in time. This can limit their utility, especially in the closed-book setting where the pretraining corpus must contain the facts the model should memorize. We introduce a diagnostic dataset aimed at probing LMs for factual knowledge that changes over time and highlight problems with LMs at either end of the spectrum—those trained on specific slices of temporal data, as well as those trained on a wide range of temporal data. To mitigate these problems, we propose a simple technique for jointly modeling text with its timestamp. This improves memorization of seen facts from the training time period, as well as calibration on predictions about unseen facts from future time periods. We also show that models trained with temporal context can be efficiently “refreshed” as new data arrives, without the need for retraining from scratch.

1 Introduction

Language models (LMs) have been suggested as repositories of real-world knowledge (Petroni et al., 2019) and there is much interest in using them for tasks such as closed-book question answering (QA; Roberts et al., 2020), fact verification (Lee et al., 2020), and dialogue (Adiwardana et al., 2020). Many facts, however, change with time. This raises two questions: Do pretrained LMs learn the appropriate temporal scope for the facts they encode? And what is the best way to update temporally scoped knowledge in pretrained models?

^{*}Equal contribution.

^{**}Also affiliated with Duke University, work done at Google.

Pretraining corpora for models such as BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019), and GPT (Radford et al., 2019) are typically derived from a snapshot of the web crawled at a specific moment in time (Raffel et al., 2020). While the impact on language modeling itself has been highlighted in recent work (e.g., Lazaridou et al., 2021; Röttger and Pierrehumbert, 2021; Hombaiyah et al., 2021), there are several potential problems specific to the encoding of factual knowledge:

- **Averaging:** For temporally scoped knowledge, the model may see conflicting information, for example, “Lebron James plays for the Cavaliers / Lakers.” Because LM training generally ignores temporal metadata, this can lead to an averaging effect, in which the model has low confidence in any of the correct answers.
- **Forgetting:** Corpora such as Wikipedia and web crawls are constantly growing, with documents distributed non-uniformly across time: There are more recent documents than older ones, both because old documents can be updated and because more web documents are generated recently than in the past. As a result, the model may fail to memorize facts that were valid only during underrepresented periods of time, and therefore do worse when asked questions about the more distant past.
- **Poor temporal calibration:** As language models become “stale”, they are increasingly likely to be queried about facts outside the temporal scope of their training data. While it may seem undesirable for a model to guess the answer to such questions, in many cases it is perfectly reasonable to assume that

the future will be like the present: for example, in twenty years the capital of Alaska is unlikely to change, even though the *governor* of Alaska is nearly impossible to predict. Ideally, the confidence with which the model responds to such queries should reflect this difficulty.

Temporally scoped facts are common in practice; however, QA datasets such as SQuAD (Rajpurkar et al., 2018) or Natural Questions (Kwiatkowski et al., 2019) focus on a single time period, even for questions whose answers are temporally scoped. Thus, our first contribution in this paper is a diagnostic dataset, TEMPLAMA (short for TEMPoral Language Model Analysis), of fill-in-the-blank queries for probing time-sensitive knowledge in LMs. The queries in TEMPLAMA are chosen such that the answer varies with time (§ 2.1). Using this dataset, we find empirical evidence of the problems mentioned above (§ 3).

As a first step towards addressing these problems, we propose a lightweight modification to pretraining. We parametrize the masked language modeling objective (MLM; Devlin et al., 2019) with temporal information, $P(y|x, t; \theta)$, where y is a masked token or span, x is the textual context, and t is the time (§ 2.3). The parameters θ must learn a representation of both text and time. In the T5 framework (Raffel et al., 2020), this can be accomplished by prefixing the input x with a string representation of t , for example, “year: 2018”. In addition, we pretrain from documents that are uniformly sampled from the timespan of the training corpus which, in our case, consists of news articles ranging from 2010–2018 (Lazaridou et al., 2021) (§ 2.1). These interventions accomplish two goals: the model is exposed to facts from the entire time range instead of just the most recent one, which avoids forgetting certain temporally scoped facts, and it prevents averaging because the facts are assigned to different time buckets (in our case years). This leads to improved recall of facts from the timespan of the training corpus (§ 3.1).

These interventions also improve the model’s temporal calibration. We find that jointly modeling text and time improves perplexity on future years unseen during training. On TEMPLAMA, the joint model degrades more gracefully than a model unaware of time. We also examine the model’s calibration farther into the future using

hand-crafted sets of queries whose answer is likely to change *frequently*, *rarely*, or *never*. We find qualitative evidence that the entropy of models trained uniformly across the training timespan increases most rapidly for the frequently changing facts (§ 3.2).

While calibration is desirable, models should be refreshed with new data when it becomes available. A standard practice for doing this is to combine the new and old data and retrain the model from scratch (e.g., Liu et al., 2021), but retraining can be costly for large-scale models (Strubell et al., 2019). On the other hand, finetuning only on the new data leads to catastrophic forgetting of the old data (Zhu et al., 2020), since standard LMs have no knowledge of what is “new” and what is “old”, unlike a model trained with temporal context. We show that our temporally scoped pretraining procedure makes LMs more amenable to post-hoc finetuning, as the data is implicitly bucketed into non-overlapping time slices. We observe a similar performance to models retrained from scratch with $30\times$ fewer steps, and without degradation on the knowledge encoded by the older data (§ 3.3).

Summary of Contributions: (1) We offer TEMPLAMA, a new dataset of temporally scoped knowledge probes. (2) We propose a simple modification to pretraining that facilitates the acquisition of temporal knowledge. (3) We conduct evaluations that demonstrate the impact of temporal shift on the knowledge encoded by existing LMs and the improvements offered by temporally scoped pretraining. (4) We perform a qualitative analysis of temporal calibration into the future, again demonstrating the positive impact of temporally scoped pretraining. (5) We show that temporally scoped pretraining also facilitates efficient updates to existing pretrained LMs.

2 Methods

We probe factual knowledge in masked LMs using span prediction—given an input statement x with a span y replaced by a special character, the task is to reconstruct that span. Additionally, we assume that each (x, y) pair has a timestamp t denoting the time at which it was written or a point in time at which its assertion is valid. In this paper, we discretize t into yearly buckets and leave more fine-grained groupings (e.g., at

the level of months or days) for future work. For simplicity and efficiency, all of our models are text-to-text Transformers (Vaswani et al., 2017) initialized from publicly available T5 checkpoints (Raffel et al., 2020) and then adapted to more time-dependent datasets. We first describe these datasets, followed by the approaches for jointly modeling text and time.

2.1 Datasets

We experiment with a large-scale news corpus (CUSTOMNEWS) for pretraining our models, combined with a smaller diagnostic dataset of factual queries (TEMPLAMA) for evaluation.

CUSTOMNEWS The CUSTOMNEWS dataset is a subset of web documents that are determined to be news (Lazaridou et al., 2021) and have an associated date either extracted from the article’s URL or from its html by looking for a publication date. We adapt this dataset in two main ways. First, we focus on a subset created by randomly sampling 1M news articles from each of the years 2010–2020 which had the maximum number of articles. Second, while Lazaridou et al. (2021) used this data for classic autoregressive language modeling, we instead adapt it for the MLM objective. Specifically, we split the articles into sentences x and then identify *salient spans* y in the text corresponding to named entities and dates. The salient span masking (SSM) paradigm improves question answering performance in both open-book (Guu et al., 2020) and closed-book settings (Roberts et al., 2020). SSM restricts the inputs to those which have a higher chance of requiring world knowledge and better aligns with our objective of measuring the factual knowledge captured by the LMs. Following Guu et al. (2020), we identify named entities using a BERT-based tagger trained on CoNLL-2003 data (Tjong Kim Sang and De Meulder, 2003) and a regular expression for dates.

TEMPLAMA We also construct a more targeted masked LM evaluation for probing temporally sensitive knowledge. Starting with the November 2020 Wikidata snapshot (Vrandečić and Krötzsch, 2014), we first identify all facts that have either a start or an end date after 2010 and whose subjects and objects are both entities with Wikipedia pages.¹ Among these 482K facts, we identify sub-

¹We use SLING (Ringgaard et al., 2017) for preprocessing.

Year	Input	Target
CUSTOMNEWS		
2017	The pound faces pressure from the US but the <u>_X_</u> election could hit euro	French
2020	<u>_X_</u> accused Liverpool of ‘crossing the line’ during win over his Chelsea side.	Frank Lampard
TEMPLAMA		
2012	Cristiano Ronaldo plays for <u>_X_</u> .	Real Madrid
2019	Cristiano Ronaldo plays for <u>_X_</u> .	Juventus FC

Table 1: Examples from CUSTOMNEWS, which masks named entities and dates from news articles, and TEMPLAMA, a novel synthetic dataset of temporally scoped factual statements built from Wikidata.

ject and relation pairs that have multiple objects at different times and select nine relations with the most such subjects. For these relations we manually write template cloze queries (e.g., “Subject works for _X_.”) and populate them with the 1000 most frequent subjects per relation. For each subject and each relation we gather all the objects with their associated time interval and construct a separate query for each year in that interval. When intervals for the object entities overlap, we add all of them to the list of correct answers. The query and the corresponding year form the inputs x and t , while the object entity is the target y . In total we construct 50,310 queries across 11 years.² Note that these type of cloze-style questions naturally follow the salient span masking paradigm, where the answer to the question is the span to be masked. Table 1 shows examples from both CUSTOMNEWS and TEMPLAMA. A full list of the relations in TEMPLAMA and their template queries is included in Appendix A.

2.2 Training and Evaluation

We train and evaluate each of our models on a mixture of CUSTOMNEWS and TEMPLAMA. All models are initialized from a public T5 checkpoint, and then further *adapted* for 300K steps on our data. From CUSTOMNEWS we hold out 2000 articles each for validation and testing from each of the yearly subsets. From TEMPLAMA we reserve 10% and 70% of the queries from each of the yearly subsets for validation and testing, respectively, ensuring that none of the subject entities overlap between train, validation, or test sets.

²The TEMPLAMA data is available at <https://github.com/google-research/language/tree/master/language/templama>.

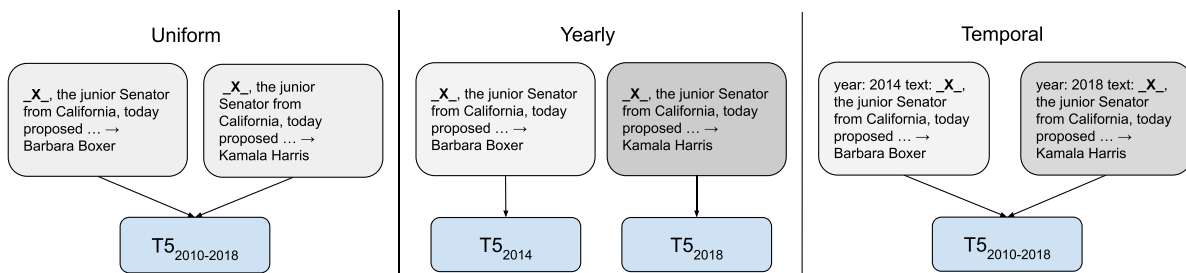


Figure 1: Three training setups to train T5 on CUSTOMNEWS: The *Uniform* model (**left**) is trained on all the data without explicit time information. The *Yearly* model (**middle**) avoids averaging over similar contexts by training separate models depending on the year, while the *Temporal* model (**right**) prepends a time prefix to each example.

Splitting along subject entities ensures that none of the facts required to answer the test queries are seen during training on TEMPLAMA (Lewis et al., 2021). Instead they must be learned in an unsupervised manner either from the T5 pretraining or when adapting to CUSTOMNEWS. We train over the combination of the two training sets such that for every 1000 inputs from CUSTOMNEWS, the model sees 1 input from TEMPLAMA. Finetuning on a small disjoint set of queries from TEMPLAMA in this manner avoids issues due to suboptimal prompts (Jiang et al., 2020b; Logan et al., 2021) by allowing the model to learn the expected format of queries and answers (e.g., “Liverpool F.C.” vs “Liverpool”).

We also partition the data into two groups based on the year: **2010–18** and **2019–20**. Models are trained only on the former, but tested on both to measure their performance for both seen and future time periods. This split was informed by the fact that the T5 checkpoints were pretrained on web text extracted in April 2019. The main metric for evaluation is a token-level F1 score between the predicted and ground truth targets, computed in the same way as for the SQuAD benchmark (Rajpurkar et al., 2018). For TEMPLAMA queries with multiple targets we take the max F1.

2.3 Jointly Modeling Text and Time

Given a dataset of (x, y, t) triples we model $P(y|x, t; \theta)$ using variants of the T5 model where, given x as the input sequence, we maximize the likelihood of the target sequence y . We compare two approaches to condition the predictions on the time t (also see Figure 1).

Yearly In the first approach we use the temporal context by training separate models specialized to different time buckets (in our case years), so

$P(y|x, t; \theta) = P(y|x; \theta_t)$. Hence, we train an ensemble of nine T5 models adapted to each year between 2010 and 2018 for an additional 300K steps. When provided with a test input, this approach routes it to the appropriate yearly expert based on its timestamp. If the timestamp falls outside 2010–18, we use the closest yearly expert (e.g., 2018 for all test inputs ≥ 2018).

Temporal Training a separate expert for each time slice reduces the averaging across conflicting contexts (§ 1), but keeping an ensemble of large-scale LMs is undesirable in practice. Moreover, there are regularities in how often facts change (e.g., the FIFA World Cup happens every 4 years, whereas NBA Championships happen every year), which a model specialized to a single time slice might not be able to learn. Hence we also train a single T5 model on the entire dataset from 2010–2018 for 300K steps. In this model, the time t is concatenated to the input, that is, $P(y|x, t; \theta) = P(y|t \oplus x; \theta)$, using a simple string representation of t as a prefix for the input x , for example, “year: 2014”.

Baselines The T5 checkpoints released by Raffel et al. (2020) are pretrained on long inputs with multiple masks and cannot directly be tested using our factual knowledge probes. Instead, we establish a baseline on the datasets introduced above using the pretrained models from Roberts et al. (2020), which were trained using SSM on Wikipedia for an additional 100K steps. This is referred to as *T5-CBQA* (closed-book question answering). We also experiment with additionally finetuning this model on TEMPLAMA for 5K steps (*T5-CBQA-ft*).

To isolate the effect of time-aware pretraining, we also train a *Uniform* model, which trains on

Model	#Parameters	CustomNews			TempLAMA		
		2010–18	2019–20	Overall	2010–18	2019–20	Overall
T5-CBQA	737M	20.2	19.8	20.1	5.4	4.3	5.2
T5-CBQA-ft	737M	15.2	15.7	15.3	17.8	15.3	17.3
Uniform	737M	30.6	27.8	30.1	28.1	19.8	26.6
Yearly	6.6B	33.4	26.7	32.2	28.5	21.8	27.3
Temporal	737M	32.1	29.5	31.6	29.6	22.2	28.2

Table 2: F1 scores of Large-sized model variants for salient span mask prediction on CUSTOMNEWS and TEMPLAMA. *T5-CBQA* is the pretrained model from Roberts et al. (2020), and *T5-CBQA-ft* is further finetuned on TEMPLAMA. The *Yearly* model is an ensemble of 9 models each finetuned on a yearly slice of the training data between 2010 and 2018. We use the 2018 model when testing on 2019–20. The *Uniform* and *Temporal* models are trained on the entire data from 2010–18, and the latter has additional temporal context. The F1 scores are macro-averaged across the evaluation years. The *Temporal* model performs better on TEMPLAMA, which is focused only on temporally scoped facts, as well as on the unseen years for CUSTOMNEWS.

the same uniformly sampled data as *Temporal* for the same number of steps, but without the time provided as an input. During training, examples are shuffled rather than presented in chronological order. Note that there are many ways of sampling training data across time, and the optimal choice likely depends on the relative importance of memorizing old versus recent facts. Here we assume all time slices in the training data are equally important and hence focus on uniform sampling.

Hyperparameters We primarily focus on the Large-sized T5 models with 770M parameters, but we also investigate the scaling with size by comparing to the Small (110M) and XXL (11B) versions. We use the same set of hyperparameters as Raffel et al. (2020), with a batch size of 2048, a fixed learning rate of 0.001, and a dropout rate of 0.1. All our models are trained for a fixed number of 300K steps, except when adapting to new data (§ 3.3), and then evaluated on the test set. We found the loss on held out CUSTOMNEWS was still improving at the end of 300K steps, but the overall trends were stable; to limit the experimentation time we did not explore longer training runs.

3 Experiments

We design several experiments to highlight the problems around temporally scoped knowledge in LMs and to test whether they can be addressed by joint models of text and time.

3.1 Memorizing Facts Across Time

To understand the interplay of memorization and time, we examine the TEMPLAMA and CUSTOMNEWS performance on the 2010–18 slice. This permits us to analyze the forgetting and averaging effects discussed in § 1 by comparing models trained on different slices of the data and with or without the temporal context.

Results Table 2 shows performance on the 2010–18 test sets of CUSTOMNEWS and TEMPLAMA. *T5-CBQA* and *T5-CBQA-ft* fare significantly worse on TEMPLAMA (17.8) than the more standard Natural Questions benchmark (28.5; c.f. Roberts et al., 2020). In particular, we find that training on the news domain leads to significant improvements on the temporally scoped knowledge required by TEMPLAMA (comparing *T5-CBQA-ft* and *Uniform*). The two approaches that condition the predictions on time, *Yearly* and *Temporal*, improve over *Uniform*, which trains on the same data but without temporal context. The *Yearly* ensemble, however, has linearly more parameters and requires linearly more compute to train. For 2010–18, the *Yearly* model performs better on CUSTOMNEWS, which is far more likely to describe short-lived facts, but the *Temporal* model is better on TEMPLAMA, where the facts typically span multiple years. We further investigate the relationship between fact durations and model performance below.

We show empirical evidence of averaging and forgetting effects in Figure 2, which plots the F1




Figure 2: F1 score of models trained on data from a specific year on CUSTOMNEWS (Left) and TEMPLAMA (Middle) as the gap between test and train years varies. Negative gaps indicate that the model is tested on data from *before* the slice on which it was trained. The F1-score is macro-averaged across all possible pairs of train/test years between 2010 and 2018. For comparison we also show the F1 score of Uniform and Temporal models averaged across 2010–18. Shaded area shows the 95% confidence interval around the macro-average. The performance drop on both sides shows the forgetting effect. (Right) F1 scores on TEMPLAMA grouped by the number of years for which the answer to a query persists. Shaded area shows the 95% confidence interval using bootstrap.

score of the year-specific models as we vary the gap between test and train years. The performance drops quickly on both sides, showing forgetting; however, the decline is larger for future years. The right plot compares F1-scores on TEMPLAMA for queries grouped by the number of years for which their answer is valid.³ This is computed from the duration of their corresponding facts in Wikidata. The uniformly trained model has higher performance on queries whose answers persist for a long time, but it does worse on queries whose answers persist for less than 5 years. The opposite is true for the year-specific models, which is intuitive due to the averaging effect of training on data from long periods of time. Adding temporal context strikes a trade-off between these two extremes, leading to the overall higher F1 in Table 2.

Qualitatively, examining the TEMPLAMA questions that the *Temporal* model answers correctly while the *Uniform* model answers incorrectly supports our hypothesis that the *Uniform* model is averaging over possible choices: It frequently answers with an entity that was more salient during our training period (see Table 5).

Scaling Table 3 shows the effect of increasing model size on the overall F1 scores on CUSTOMNEWS and TEMPLAMA. In general, larger model sizes lead to a bigger improvement when training with temporal context.

Longer Time Span. Table 6 compares the Large-sized *Uniform* and *Temporal* models when

³For multiple answers we pick the duration of the first one.

Size	CustomNews		TempLAMA	
	Uniform	Temporal	Uniform	Temporal
Small	21.1	21.9	20.7	20.5
Large	30.1	31.6	26.6	28.2
XXL	32.3	33.8	28.4	30.5

Table 3: Overall F1-score averaged from 2010–20 for *Uniform* and *Temporal* models for different model sizes. Larger models benefit more from the temporal context.

trained on a wider time period from 2004 to 2018.⁴ While the *Temporal* model still outperforms *Uniform*, the gap is smaller between the two compared to when training on 2010–18. In general increasing the time period entails memorizing more facts for the *Temporal* model. Hence, this result suggests that the model size should also be increased when training on longer time spans.

CronQuestions To explore whether the improved memorization of facts translates to downstream tasks, we finetune the *Uniform* and *Temporal* models on CronQuestions, a dataset of 410K time-dependent questions based on temporal knowledge graphs (Saxena et al., 2021). It consists of questions where the answer is either an entity or a temporal expression. Similar to TEMPLAMA, the questions are based on Wikidata across time. We focus on a closed-book version of the task, similar to the setup in Roberts et al. (2020), where the model is trained to predict the

⁴CUSTOMNEWS only has a small number of articles from 2003 and before.

Size	Model	EM	F1
Small	None	3.63	9.51
	Uniform	4.01	10.27
	Temporal	4.05	10.20
Large	None	4.10	10.78
	2018	4.39	10.87
	Uniform	4.70	11.34
	Temporal	5.13	11.93
XXL	None	5.44	12.19
	Uniform	5.71	12.61
	Temporal	5.81	12.88

Table 4: Test set results for models finetuned on the CronQuestions dataset in a closed-book manner. ‘‘None’’ refers to finetuning the T5 baseline; the ‘‘2018’’ model is adapted to the 2018 slice of CUSTOMNEWS.

first answer in the list of correct answers for an input question. During evaluation, it is compared to each answer in the set of correct answers, and we take the maximum score among them. Table 4 lists the SQuAD-based EM and F1 metrics on the test set. We see an improvement in memorization for the *Uniform* and *Temporal* models, with the latter doing slightly better on the Large and XXL model sizes.

3.2 Better Calibration in the Future

We examine the model’s performance on future slices of data at two different time scales. In the first, we look at *graceful degradation*, mimicking the life-cycle of a model that has been deployed, and thus has not seen the newest slices of data yet. In the second, we ask the models to predict relations in the more distant future. While this may seem unreasonable, it is possible to articulate coherent intuitions about the future: For example, the capitals of U.S. states change far less frequently than their governors, and the probabilities emitted by language models should reflect this.

3.2.1 Graceful Degradation

Here we examine the TEMPLAMA and CUSTOMNEWS performance on the 2019–20 slices. Note that none of the models were pretrained or adapted to this slice, so these experiments allow us to measure degradation. We additionally look at the

perplexity of the masked LM, which we compute as:

$$\text{ppl} = \exp - \frac{\sum_{(x,y,t)} \log P(y|x,t;\theta)}{\sum_y \text{len}(y)}.$$

Following Lazaridou et al. (2021), we expect perplexity to increase for slices that are not covered in the training data, but we expect the temporally conditioned model to be relatively more robust.

Results Comparing the *Uniform* and *Temporal* models in Table 2, we can see that training with temporal context improves F1 scores on the 2019–20 slices. The *Yearly* ensemble, which uses the latest 2018 model when tested on 2019–20, is significantly worse on CUSTOMNEWS but comparable on TEMPLAMA; potentially because some of the answers remain the same. A closer look at the model predictions reveals that, unsurprisingly, none of the models are able to predict the TEMPLAMA facts that change after the training period. Adding temporal context simply allows the *Temporal* model to persist the unchanged facts to 2019–20. On CUSTOMNEWS it has higher performance on the SSM objective, which includes both dates and entities in articles from an unseen time period.

Table 7 shows MLM perplexity on the CUSTOMNEWS test set. The *Temporal* model has lowest perplexities on both the seen and unseen slices of evaluation data. The *Uniform* model has lower perplexity than the *Yearly* one, especially on the future slices where we use the 2018 expert for the latter. This suggests that, for language modeling, training on more data outweighs the benefit of training on the specific temporal distribution of test data.

Do the models learn how soon an answer is likely to change in the future? We do a qualitative analysis by partitioning the TEMPLAMA test queries where each model was correct in the 2018 evaluation into two sets: those with *Single* or *Multiple* answers across 2010–20. Then we measure the log-likelihood of that correct answer as we change the input year t from 2019 to 2029, and plot the change in log-likelihood relative to 2018 in Figure 3. For the *T5-CBQA-ft* and *Uniform* models, we vary the input years by prefixing queries with ‘‘In year,...’’. The confidence for all models decreases as we get into the future, which is reasonable since all relations in TEMPLAMA

Input	Year	Uniform	Temporal
X is the chair of Federal Reserve System.	2019	Janet L. Yellen	Jerome Powell
Nigel Farage is a member of the _X_.	2019	UK Independence Party	Brexit Party
Mark Sanford holds the position of _X_.	2017	Governor of South Carolina	United States representative
X is the head of the government of New York City.	2016	Michael Bloomberg	Bill de Blasio
X is the head coach of Real Madrid CF.	2015	Zinedine Zidane	Carlo Ancelotti
Theresa May holds the position of _X_.	2014	Prime Minister of Great Britain	Home Secretary
Peyton Manning plays for _X_.	2014	Indianapolis Colts	Denver Broncos
X is the head of the government of United Kingdom.	2011	Theresa May	David Cameron
Marissa Mayer works for _X_.	2011	Yahoo	Google
Rahm Emanuel holds the position of _X_.	2010	Mayor of Chicago	White House Chief of Staff

Table 5: Examples comparing the *Uniform* and *Temporal* models on TEMPLAMA. The former frequently predicts a more common or newsworthy answer from the range of the training data, without taking the year into account.

Model	2004–09	2010–18	2019–20
Uniform	34.8 (+6.3)	29.8 (−0.8)	27.4 (−0.4)
Temporal	36.3 (+5.2)	31.1 (−1.0)	28.8 (−0.7)

Table 6: F1 scores on different evaluation slices of CUSTOMNEWS for models trained on data from 2004–18. Numbers in the parentheses show the absolute difference from the same model trained on data from 2010–18.

Model	2010–18	2019–20
T5-CBQA	26.11	29.22
Uniform	11.68	14.37
Yearly	13.62	23.30
Temporal	11.33	13.58

Table 7: Masked language modeling perplexity on CUSTOMNEWS (lower is better). The Temporal model degrades less when evaluated on the future time slice.

are time-sensitive. However, the confidence of the *Temporal* model decreases more rapidly for queries with multiple answers, reflecting the intuition that facts which have changed in the past are likely to change again in the future.

3.2.2 Future Relations

To further probe the models’ understanding of expected versus unexpected changes in the future, we curate a small diagnostic dataset of queries about future relations. We restrict the queries such that the answer is always either one of the 200 largest US cities or one of the 249 countries in the world. This allows us to compute the entropy of the predictions over a fixed set. To relate model

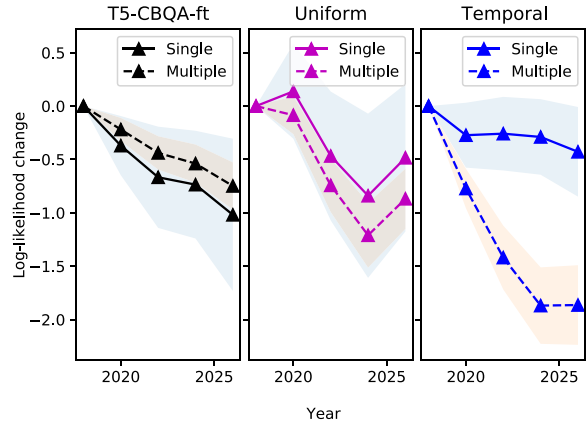


Figure 3: Change in log-likelihood over time of the most recent answer (from 2018) for TEMPLAMA queries with *Single* or *Multiple* answers. The difference is taken from the value for the 2018 answer. The *Temporal* model exhibits a more pronounced confidence gap for facts that changed in the past.

predictions to commonsense intuitions, we construct three sets of queries based on how frequently they are expected to change: *frequent*, *rare*, and *never*. For example, the location of an awards show might change every year, while the city an athlete plays in changes every few years, and the location of a landmark almost never changes. Then, given queries like “In 2022, the Space Needle will be in _X_” and “In 2022, the NBA All-Star Game will be in _X_”, a model with a reasonable representation of time should have lower entropy for the former rather than the latter. Moreover, the entropy should increase with time as the queries address the more distant future, and the *rate* of increase should be greatest for frequently-changing relations. Note that we do not expect models to provide the *correct* answers

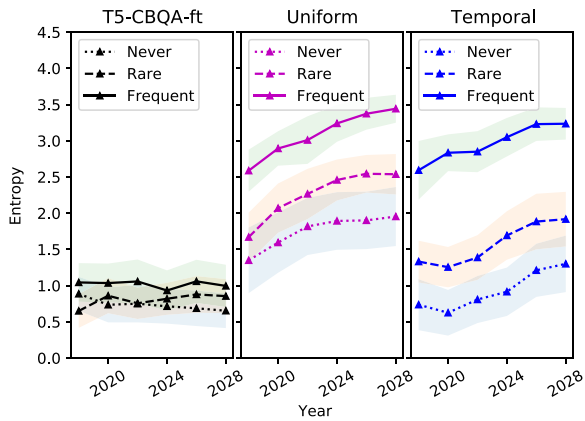


Figure 4: Entropy over time for *frequent*, *rare*, and *never*-changing queries. The *Temporal* model is more uncertain about frequently changing queries as time passes, and has a flatter entropy for constant facts.

for these queries (which we do not know anyway), but only assign confidence in a manner consistent with human intuitions. In total, we constructed 86 queries across the three sets, which are included in Appendix B.

Results Figure 4 shows the entropy of different model variants averaged across the three sets of queries and plotted over time. The baseline *T5-CBQA-ft* model has a low constant entropy throughout, irrespective of the query type. Combined with its low accuracy on future slices from Table 2, this suggests it remains confidently incorrect and has poor calibration about which facts are likely to change. Both the *Uniform* and *Temporal* models have increasing uncertainty in the future, which is ordered correctly according to intuition: highest for the queries of frequently changing facts, and lowest for queries whose answers are expected not to change. Interestingly, the *Temporal* model has a largely constant entropy for rare- and never-changing queries until 2022, after which it begins to increase. While this agrees with intuition, ideally a model should have low entropy on the never-changing set further into the future.

Overall, these results suggests that: (1) models trained uniformly over a wide range of time-sensitive data show improved calibration about expected changes in the future; and (2) training with temporal context further improves this calibration for the first few years beyond the training period, in our case from 2019 to 2022. We also

note the limitations with this evaluation, however: (1) due to manual curation by the authors there are only 86 queries in these sets, and are likely to be biased in the facts they probe; and (2) entropy mixes different kinds of uncertainty: that which is inherent in the query (e.g., there are more distinct countries than cities with NFL teams), as well as that due to the lack of confidence in the model. We are interested in the latter, but our evaluation does not disentangle the two effects.

3.3 Cheaper Adaptation to New Data

Improved calibration about the future can help minimize mistakes after the training time period (e.g., by abstaining), but eventually models need to be refreshed as the world changes and new data arrives. In this section, we consider the setting where we have an already trained model on the 2010–18 slices, as well as new data from the 2019 slice. We attempt to update the model on this new data (as measured by the combined performance on 2019–20 held out data) without forgetting the 2010–18 slices. These experiments are similar to the task posed by Lazaridou et al. (2020), but we compare the impact of adapting versus retraining from scratch. Finetuning only on the newest data (2019) is suboptimal as the model forgets facts about the past (Figure 5), which was also observed by Zhu et al. (2020). Here we explore a simple alternative—training on a mixture which samples a data point from the new slice (2019) with probability α and a data point from the old slices (2010–18) with probability $1-\alpha$. We finetune both the *Temporal* and *Uniform* models on this mixture for an additional 50K steps and compare the resulting performance to models retrained from scratch for 300K steps on data sampled uniformly from all slices (2010–19). Note that the latter strategy can be costly for large-scale LMs (Strubell et al., 2019).

Results Figure 5 shows the F1-score on CUSTOM-NEWS and TEMPLAMA as we vary α . Across all values of α , the *Uniform* model improves significantly on the 2019 slice, but this comes at the cost of degrading on the 2010–18 slices. The *Temporal* model also adapts to 2019, but shows minimal degradation on the 2010–18 slice up to $\alpha = 0.6$. For $\alpha = 0.5$ we found that its performance with 10K additional steps matches that of the *Temporal* model trained from scratch for 300K steps, suggesting that models

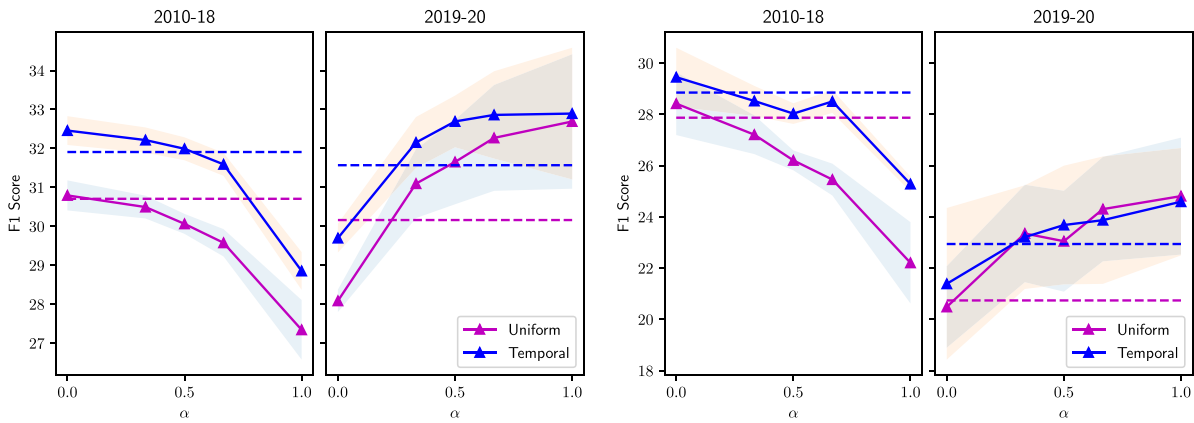


Figure 5: CUSTOMNEWS (**left**) and TEMPLAMA (**right**) F1 score as models are *adapted* to new data from 2019 for 50K steps. α denotes the fraction of training examples which come from the 2019 slice (remaining examples come from the 2010–18 slices). Dotted lines indicate models *retrained* from scratch for 300K steps on equal proportions of all data from 2010–19. The *Temporal* model degrades less than *Uniform* on the 2010–18 slice when adapted.

trained with temporal context can be efficiently adapted to new data without forgetting facts from the old data.

4 Discussion and Limitations

Our experiments have shown that current models have practical limitations in their ability to memorize the past and reasonably estimate the future. These limitations can be mitigated by providing the model the date at which a text was created. While our results show consistent advantages, they also represent a narrow understanding of time. In particular, the publication date of a news articles does not necessarily correspond to the temporal scope of *all* events described in the article. For example, articles may talk about historical events or discuss events scheduled to happen in the future. In CUSTOMNEWS around 3.9% sentences explicitly mention a year between 2010 and 2018, and 2.1% mention the same year as the publication date of the article. This fraction is likely responsible for the improvement of the *Uniform* model. The *Temporal* model further assigns an approximate scope to the remaining 96% sentences and it is encouraging to see improvements from that. One avenue for future work is to explore better strategies for assigning dates to these sentences.

We have focused on closed-book question answering, but temporal staleness of language models may have impacts in other applications as well. For example, in *open*-book question answering, it

is still necessary to align the question with relevant text in the retrieved passage, and this could be challenging when the question cannot be properly encoded by a stale LM: For example, the query “which countries were affected by the 2020 hurricane season?” would not match the passage “Iota caused damages of \$564 million in Nicaragua” in an LM that did not have access to training data mentioning “Iota” as a hurricane.

Another limitation of our work is that TEMPLAMA is constructed in a synthetic manner from WikiData. Incomplete or incorrect facts in the KB can result in incorrect queries in TEMPLAMA; for instance, we assume a missing start date implies the fact is valid from the beginning of our time period of interest. We partition the TEMPLAMA and CUSTOMNEWS dataset on the same yearly slices despite the nature of the datasets being quite different. Moreover, we did not investigate using longer or shorter temporal partitions. Additionally, we did not test the ability to model temporal expressions such as “before” or “during”, and we did not investigate temporal commonsense (e.g., Zhou et al. 2019), temporal ordering (e.g., Ning et al. 2020), or events (e.g., Zhou et al. 2021).

Lastly, it is worth noting that like all closed-book models the models presented in this paper are also likely to only memorize common facts about popular entities. This has the danger of reinforcing stereotypes and leading to unfair outcomes. Additionally, training the multitude of large-scale language models presented in this paper required the use of 32 Cloud TPU v3 cores for several

hundred hours, which has a significant environmental impact (Strubell et al., 2019). However, our hope is that efficient schemes for updating temporally sensitive knowledge in LMs will eventually save energy costs in the long run.

5 Related Work

There is extensive prior work on learning diachronic embeddings of individual words (e.g., Wijaya and Yeniterzi, 2011; Hamilton et al., 2016; Bamler and Mandt, 2017). Particularly related is the approach of Dubossarsky et al. (2019), who learn time-sensitive embeddings by concatenating each word token with the decade in which it appears. As contextualized embedding models have largely replaced non-contextual word embeddings (Peters et al., 2018; Devlin et al., 2019), the main application of diachronic word embeddings is to detect and model lexical semantic changes (e.g., Frermann and Lapata, 2016), rather than to improve temporal awareness on downstream tasks. Our work fills this gap by adding a temporal component to T5, a pretrained language model that can complete multi-token spans. While Giulianelli et al. (2020) use contextualized embeddings from BERT to model lexical semantic changes post hoc, they do not add a time-sensitive component to the language model itself. Thus, their approach cannot support time-aware fact completion.

Several studies have focused on degradation of models on test data from a different time period than their training data (Huang and Paul, 2018, 2019; Jaidka et al., 2018; Lukes and Søgaard, 2018; Florio et al., 2020). Delasalles et al. (2019) introduced an LSTM language model that conditions on dynamic author representations computed separately, and showed that it improves perplexity on both seen and unseen (future) time periods. Most recently, Röttger and Pierrehumbert (2021) analyzed the interplay between temporal adaptation during pretraining and finetuning, and concluded that while both stages benefit from adaptation separately, adaptation during pretraining does not help the downstream task. Here we show that the benefits of adaptation can be achieved using a single model that conditions on time. We further show that the benefits of adaptation come, at least in part, from better memorization of time-sensitive facts.

In production contexts, an important form of temporal generalization is the deployment of models trained on data up to a certain time T but applied on data after T : that is, the present. Lazaridou et al. (2021) show that language models gradually degrade in performance under such a time-stratified setting, and propose *dynamic evaluation* (Krause et al., 2018) as a potential mitigation. However, LMs are frequently applied to past data as well, for example, for extracting representations, and here we show that updating on only the new data degrades performance on old data. Our approach of conditioning on the temporal context alleviates this issue.

A related line of work has explored editing neural predictions after training given a dataset of revised input and output pairs (Sinitsin et al., 2020; Zhu et al., 2020; De Cao et al., 2021). Here we introduce a different setting where we have access to new unlabeled text after model training, which must be used implicitly to update the factual predictions of the model. In this case the update procedure also needs to figure out *which* facts must be updated and which ones remain the same.

Petroni et al. (2019) introduced the LAMA benchmark for probing the factual knowledge memorized by LMs, which consists of cloze queries about facts, for example, ‘‘Dante was born in $_X__$ ’’. Follow up studies have introduced improved prompts for eliciting such knowledge (Jiang et al., 2020b) as well as multilingual versions (Jiang et al., 2020a; Kassner et al., 2021). However, all these benchmarks assume a static view of the knowledge inside an LM, and consider all answers across time to be correct for a given query. The TEMPLAMA dataset instead focuses on relations where the answers change with time and uses temporal scopes to determine the correct answer.

TEMPLAMA is similar in spirit to KB-QA benchmarks which focus on temporal reasoning such as TempQuestions (Jia et al., 2018) and CronQuestions (Saxena et al., 2021). Its format, however, mimics the masked LM task typically used in pretraining, since it is intended as a zero/few-shot probe. Unlike those datasets, we further restrict the queries to subject and relation pairs for which multiple objects exist at different points in time, and ensure a balanced distribution over the entire time period of interest from 2010–2020.

WikiData ID	Relation	# Queries	Template
P54	member of sports team	9033	<subject> plays for <object>.
P39	position held	7343	<subject> holds the position of <object>.
P108	employer	9049	<subject> works for <object>.
P102	political party	7324	<subject> is a member of the <object>.
P286	head coach	4886	<object> is the head coach of <subject>.
P69	educated at	1672	<subject> attended <object>.
P488	chairperson	4190	<object> is the chair of <subject>.
P6	head of government	4125	<object> is the head of the government of <subject>.
P127	owned by	2688	<subject> is owned by <object>.

Table 8: Templates used for converting WikiData facts into natural language queries.

6 Conclusion

Though temporally scoped facts are common in practice, there has been little prior work exploring how these are encoded in pretrained LMs. We show that T5 does poorly on such facts and training on the news domain improves it significantly. However, simply training on more data is sub-optimal; conditioning on the temporal context of the data improves memorization of facts further. Hence, we propose a time-aware language model which conditions on string prefixes of time. Other benefits of time-aware LMs include a better calibration of expected changes in the future, and a cheaper adaptation to new slices of time-stamped data.

Acknowledgments

We would like to thank the Action Editor and Reviewers for comments on an earlier draft of this

work, and the T5X team at Google for their T5 implementation.

Supplementary Material

A TEMPLAMA Templates

Table 8 lists the 9 WikiData relations used for constructing TEMPLAMA. We instantiate the template for the relation in each fact by replacing ‘<subject>’ with the name of the subject entity, and ‘<object>’ with ‘_X_’. The answer to the query is the name of the corresponding object entity. We construct a separate query for each year that the fact is valid.

B Future Relations

Table 9 shows the queries used as part of the Future Relations experiment in § 3.2. These queries were constructed by searching for lists of events, popular athletes, and issuing targeted queries to the WikiData Query Service.

Frequent	Rare	Never
Cities		
The Super Bowl will take place in .X..	Visa Inc.'s headquarters are located in .X..	South by Southwest will take place in .X..
The NCAA Men's Final Four will take place in .X..	SEGA of America's headquarters are located in .X..	Lollapalooza will take place in .X..
The first game of the World Series will take place in .X..	Barack Obama lives in .X..	Summerfest will take place in .X..
The US PGA Championship will take place in .X..	Hillary Clinton lives in .X..	Outside Lands will take place in .X..
The golf US Open will take place in .X..	Donald Trump works in .X..	Spoletto Festival USA will take place in .X..
The NBA all-star game will take place in .X..	The Chargers play their home games in .X..	CMA Music Festival will take place in .X..
The NFL Draft will take place in .X..	The Raiders play their home games in .X..	Made in America Festival will take place in .X..
The Netroots Nation conference will take place in .X..	The Rams play their home games in .X..	The US Open Tennis Championships will take place in .X..
The MLB all-star game will take place in .X..	General Electric's headquarters are located in .X..	The Masters tournament will take place in .X..
The team from .X.. won the NBA championship.	Toyota's US headquarters are located in .X..	The Kentucky Derby will take place in .X..
The team from .X.. won the Stanley Cup.	Nestle's headquarters are located in .X..	The capital of Washington state is .X..
The team from .X.. won the World Series.	Tesla's headquarters are located in .X..	The capital of California state is .X..
The team from .X.. won the Super Bowl.	Lebron James plays in .X..	The capital of Texas is .X..
The golf US Women's Open will take place in .X..	Tom Brady plays in .X..	The capital of Florida is .X..
Wrestlemania will take place in .X..	Kevin Durant plays in .X..	The Space Needle is located in .X..
	Stephen Curry plays in .X..	The Statue of Liberty is located in .X..
	Sidney Crosby plays in .X..	Golden Gate Bridge is located in .X..
	Mike Trout plays in .X..	The White House is located in .X..
	The Democratic National Convention will next take place in .X..	The Liberty Bell is located in .X..
	The Republican National Convention will next take place in .X..	
Countries		
The Six Nations Championship will be held in .X..	The UN Secretary general is from .X..	The Oxford Literary Festival will take place in .X..
The Association for Computational Linguistics will meet in .X..	The Pope hails from .X..	Wimbledon will take place in .X..
The Neural Information Processing Systems conference will be held in .X..	The FIFA world cup was last held in .X..	Tomorrowland will take place in .X..
The Palme d'Or winner is from .X..	The Cricket world cup was last held in .X..	Hajj will take place in .X..
The Tour De France winner is from .X..	The UEFA European Football Championship was last held in .X..	The Eiffel Tower is located in .X..
The Wimbledon Men's Singles winner is from .X..	The Olympics were last held in .X..	The Taj Mahal is located in .X..
The UEFA Champions League final will take place in .X..	The Winter Olympics were last held in .X..	Burj Khalifa is located in .X..
The G20 summit will be held in .X..	The FIFA world cup was last won by .X..	Machu Picchu is located in .X..
The G7 summit will be held in .X..	The Cricket world cup was last won by .X..	Stonehenge is located in .X..
The United Nations Climate Change conference will take place in .X..	.X.. won the most gold medals in the last Olympics.	The world's largest country by land area is .X..
		The world's longest river is in .X..
		The world's tallest mountain is in .X..

Table 9: The Future Relations dataset used to test model calibration over future years. The three columns represent queries whose answers, intuitively, change *frequently* or every year, *rarely* or once every few years, and *never*. The top section includes queries whose answer is a US city, while the bottom section includes queries whose answer is a country.

References

- Daniel Adiwardana, Minh-Thang Luong, David R. So, Jamie Hall, Noah Fiedel, Romal Thoppilan, Zi Yang, Apoorv Kulshreshtha, Gaurav Nemade, Yifeng Lu, and Quoc V. Le. 2020. Towards a human-like open-domain chatbot. *CoRR*, abs/2001.09977.
- Robert Bamler and Stephan Mandt. 2017. Dynamic word embeddings. In *International Conference on Machine Learning*, pages 380–389. PMLR.
- Nicola De Cao, Wilker Aziz, and Ivan Titov. 2021. Editing factual knowledge in language models. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6491–6506, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.emnlp-main.522>
- Edouard Delasalles, Sylvain Lamprier, and Ludovic Denoyer. 2019. Learning dynamic author representations with temporal language models. In *2019 IEEE International Conference on Data Mining (ICDM)*, pages 120–129. <https://doi.org/10.1109/ICDM.2019.00022>
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Haim Dubossarsky, Simon Hengchen, Nina Tahmasebi, and Dominik Schlechtweg. 2019. Time-out: Temporal referencing for robust modeling of lexical semantic change. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 457–470, Florence, Italy. Association for Computational Linguistics. <https://doi.org/10.18653/v1/P19-1044>
- Komal Florio, Valerio Basile, Marco Polignano, Pierpaolo Basile, and Viviana Patti. 2020. Time of your hate: The challenge of time in hate speech detection on social media. *Applied Sciences*, 10(12). <https://doi.org/10.3390/app10124180>
- Lea Frermann and Mirella Lapata. 2016. A Bayesian model of diachronic meaning change. *Transactions of the Association for Computational Linguistics*, 4:31–45. https://doi.org/10.1162/tacl_a_00081
- Mario Giulianelli, Marco Del Tredici, and Raquel Fernández. 2020. Analysing lexical semantic change with contextualised word representations. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3960–3973, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.acl-main.365>
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Mingwei Chang. 2020. Retrieval augmented language model pre-training. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 3929–3938. PMLR.
- William L. Hamilton, Jure Leskovec, and Dan Jurafsky. 2016. Diachronic word embeddings reveal statistical laws of semantic change. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1489–1501, Berlin, Germany. Association for Computational Linguistics. <https://doi.org/10.18653/v1/P16-1141>
- Spurthi Amba Hombaiah, Tao Chen, Mingyang Zhang, Mike Bendersky, and Marc Najork. 2021. Dynamic language models for continuously evolving content. In *Knowledge Discovery and Data Mining (KDD)*. <https://doi.org/10.1145/3447548.3467162>
- Xiaolei Huang and Michael J. Paul. 2018. Examining temporality in document classification. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 694–699, Melbourne, Australia. Association for Computational Linguistics. <https://doi.org/10.18653/v1/P18-2110>
- Xiaolei Huang and Michael J. Paul. 2019. Neural temporality adaptation for document classification: Diachronic word embeddings and domain adaptation models. In *Proceedings of*

- the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4113–4123, Florence, Italy. Association for Computational Linguistics. <https://doi.org/10.18653/v1/P19-1403>
- Robert L. Logan IV, Ivana Balazevic, Eric Wallace, Fabio Petroni, Sameer Singh, and Sebastian Riedel. 2021. Cutting down on prompts and parameters: Simple few-shot learning with language models. *CoRR*, abs/2106.13353.
- Kokil Jaidka, Niyati Chhaya, and Lyle Ungar. 2018. Diachronic degradation of language models: Insights from social media. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 195–200, Melbourne, Australia. Association for Computational Linguistics. <https://doi.org/10.18653/v1/P18-2032>
- Zhen Jia, Abdalghani Abujabal, Rishiraj Saha Roy, Jannik Strötgen, and Gerhard Weikum. 2018. Tempquestions: A benchmark for temporal question answering. In *Companion Proceedings of the The Web Conference 2018*, pages 1057–1062. <https://doi.org/10.1145/3184558.3191536>
- Zhengbao Jiang, Antonios Anastasopoulos, Jun Araki, Haibo Ding, and Graham Neubig. 2020a. X-FACTR: Multilingual factual knowledge retrieval from pretrained language models. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5943–5959, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.emnlp-main.479>
- Zhengbao Jiang, Frank F. Xu, Jun Araki, and Graham Neubig. 2020b. How can we know what language models know? *Transactions of the Association for Computational Linguistics*, 8:423–438. https://doi.org/10.1162/tacl_a_00324
- Nora Kassner, Philipp Dufter, and Hinrich Schütze. 2021. Multilingual LAMA: Investigating knowledge in multilingual pretrained language models. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 3250–3258, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.eacl-main.284>
- Ben Krause, Emmanuel Kahembwe, Iain Murray, and Steve Renals. 2018. Dynamic evaluation of neural sequence models. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 2766–2775. PMLR.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. Natural questions: A benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:452–466. https://doi.org/10.1162/tacl_a_00276
- Angeliki Lazaridou, Adhi Kuncoro, Elena Gribovskaya, Devang Agrawal, Adam Liska, Tayfun Terzi, Mai Gimenez, Cyprien de Masson d’Autume, Tomas Kocisky, Sebastian Ruder, et al. 2021. Mind the gap: Assessing temporal generalization in neural language models. *Advances in Neural Information Processing Systems*, 34.
- Konstantina Lazaridou, Alexander Löser, Maria Mestre, and Felix Naumann. 2020. Discovering biased news articles leveraging multiple human annotations. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 1268–1277, Marseille, France. European Language Resources Association.
- Nayeon Lee, Belinda Z. Li, Sinong Wang, Wen-tau Yih, Hao Ma, and Madian Khabsa. 2020. Language models as fact checkers? In *Proceedings of the Third Workshop on Fact Extraction and VERification (FEVER)*, pages 36–41, Online. Association for Computational Linguistics.
- Patrick Lewis, Pontus Stenetorp, and Sebastian Riedel. 2021. Question and answer test-train overlap in open-domain question answering datasets. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1000–1008, Online. Association for

- Computational Linguistics. <https://doi.org/10.18653/v1/2021.eacl-main.86>
- Jialu Liu, Tianqi Liu, and Cong Yu. 2021. Newsembed: Modeling news through pre-trained document representations. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining, KDD '21*, pages 1076–1086, New York, NY, USA. Association for Computing Machinery. <https://doi.org/10.1145/3447548.3467392>
- 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. *CoRR*, abs/1907.11692.
- Jan Lukes and Anders Søgaard. 2018. Sentiment analysis under temporal shift. In *Proceedings of the 9th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 65–71, Brussels, Belgium. Association for Computational Linguistics. <https://doi.org/10.18653/v1/W18-6210>
- 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. <https://doi.org/10.18653/v1/2020.emnlp-main.88>
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics. <https://doi.org/10.18653/v1/N18-1202>
- Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. Language models as knowledge bases? 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 2463–2473.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(140):1–67.
- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. Know what you don’t know: Unanswerable questions for SQuAD. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 784–789, Melbourne, Australia. Association for Computational Linguistics. <https://doi.org/10.18653/v1/P18-2124>
- Michael Ringgaard, Rahul Gupta, and Fernando C. N. Pereira. 2017. SLING: A framework for frame semantic parsing. *CoRR*, abs/1710.07032.
- Adam Roberts, Colin Raffel, and Noam Shazeer. 2020. How much knowledge can you pack into the parameters of a language model? In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5418–5426, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.emnlp-main.437>
- Paul Röttger and Janet Pierrehumbert. 2021. Temporal adaptation of BERT and performance on downstream document classification: Insights from social media. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2400–2412, Punta Cana, Dominican Republic. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.findings-emnlp.206>
- Apoorv Saxena, Soumen Chakrabarti, and Partha Talukdar. 2021. Question answering over temporal knowledge graphs. In *Proceedings*

- of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers), pages 6663–6676, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.findings-emnlp.206>
- Anton Sinitsin, Vsevolod Plokhotnyuk, Dmitry Pyrkin, Sergei Popov, and Artem Babenko. 2020. Editable neural networks. In *International Conference on Learning Representations*.
- Emma Strubell, Ananya Ganesh, and Andrew McCallum. 2019. Energy and policy considerations for deep learning in NLP. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3645–3650, Florence, Italy. Association for Computational Linguistics. <https://doi.org/10.18653/v1/P19-1355>
- Erik F. Tjong Kim Sang and Fien De Meulder. 2003. Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition. In *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003*, pages 142–147. <https://doi.org/10.3115/1119176.1119195>
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Denny Vrandečić and Markus Krötzsch. 2014. Wikidata: A free collaborative knowledgebase. *Commun. ACM*, 57(10):78–85. <https://doi.org/10.1145/2629489>
- Derry Tanti Wijaya and Reyyan Yeniterzi. 2011. Understanding semantic change of words over centuries. In *Proceedings of the 2011 International Workshop on DETecting and Exploiting Cultural Diversity on the Social Web, DETECT ’11*, pages 35–40, New York, NY, USA. Association for Computing Machinery. <https://doi.org/10.1145/2064448.2064475>
- 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. <https://doi.org/10.18653/v1/D19-1332>
- Ben Zhou, Kyle Richardson, Qiang Ning, Tushar Khot, Ashish Sabharwal, and Dan Roth. 2021. Temporal reasoning on implicit events from distant supervision. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1361–1371, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.naacl-main.107>
- Chen Zhu, Ankit Singh Rawat, Manzil Zaheer, Srinadh Bhojanapalli, Daliang Li, Felix X. Yu, and Sanjiv Kumar. 2020. Modifying memories in transformer models. *CoRR*, abs/2012.00363.