



Yang, E., MacAvaney, S. , Lewis, D. and Frieder, O. (2022) Goldilocks: Just-Right Tuning of BERT for Technology-Assisted Review. In: 44th European Conference on Information Retrieval (ECIR 2022), Stavanger, Norway, 10-14 Apr 2022, pp. 502-517.
(doi: [10.1007/978-3-030-99736-6_33](https://doi.org/10.1007/978-3-030-99736-6_33))

The material cannot be used for any other purpose without further permission of the publisher and is for private use only. There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

https://doi.org/10.1007/978-3-030-99736-6_34

<https://eprints.gla.ac.uk/259713/>

Date deposited: 21 Jan 2022

Goldilocks: Just-Right Tuning of BERT for Technology-Assisted Review

Eugene Yang¹, Sean MacAvaney², David D. Lewis³, and Ophir Frieder⁴

¹ HLTCOE, Johns Hopkins University, USA eugene.yang@jhu.edu

² University of Glasgow, United Kingdom Sean.MacAvaney@glasgow.ac.uk

³ Reveal-Brainspace, USA ecir2022paper@davelewis.com

⁴ IRLab, Georgetown University, USA ophir@ir.cs.georgetown.edu

Abstract. Technology-assisted review (TAR) refers to iterative active learning workflows for document review in high recall retrieval (HRR) tasks. TAR research and most commercial TAR software have applied linear models such as logistic regression to lexical features. Transformer-based models with supervised tuning are known to improve effectiveness on many text classification tasks, suggesting their use in TAR. We indeed find that the pre-trained BERT model reduces review cost by 10% to 15% in TAR workflows simulated on the RCV1-v2 newswire collection. In contrast, we likewise determined that linear models outperform BERT for simulated legal discovery topics on the Jeb Bush e-mail collection. This suggests the match between transformer pre-training corpora and the task domain is of greater significance than generally appreciated. Additionally, we show that *just-right* language model fine-tuning on the task collection before starting active learning is critical. Too little or too much fine-tuning hinders performance, worse than that of linear models, even for a favorable corpus such as RCV1-v2.

1 Introduction

High recall retrieval (HRR) tasks (also called *annotation* tasks) involve identifying most or all documents of interest in a large collection. HRR tasks include electronic discovery in the law (eDiscovery) [3], systematic review in medicine [22–24, 47], document sensitivity review [34], online content moderation [55], and corpus annotation to support research and development [60].

Technology-assisted review (TAR) refers to the automated methods to reduce the number of documents reviewed in HRR projects [36]. Iterative, pool-based active learning of predictive models for review prioritization is the most commonly applied workflow [9, 10]. Linear models such as logistic regression and support vector machines (SVMs) applied to lexical and metadata features are the most common supervised learning approaches. Unlike in classification and adhoc retrieval tasks, the supervised learning model in TAR is typically discarded after use. This is because each legal case or other project has its own retrieval objective, and because of concerns of leaking confidential information. Therefore, the cost of training the supervised learning model in TAR often cannot be amortized over future data or across tasks.

Pre-trained transformers [46] such as BERT [12], GPT-3 [5], and T5 [38] are effective at a variety of natural language processing tasks. These models learn linguistic patterns from very large corpora in an unsupervised fashion (*pre-training*) and can be tuned to language characteristics of a particular task data set (*LM fine-tuning*) [12, 19]. They can then be applied to a task on that data set by zero-shot transfer learning [32, 49] or by *task fine-tuning* to labeled training data [20, 58]. Transformers have improved effectiveness at tasks related to HRR such as document classification [1], entity extraction [12], and adhoc retrieval [33]. This has inspired initial commercial use of transformers by eDiscovery providers, though not yet in an active learning context.⁵

We are not aware of published studies of transformers in TAR workflows. Several studies have evaluated task fine-tuning using active learning [30, 44], including for text classification tasks [13, 59]. These studies, however, have evaluated generalization to new data using training/test splits. HRR, like relevance feedback in adhoc search [42], is a transductive setting: evaluation is on the same task corpus from which the training data is selected by active learning.

The transductive setting makes of less importance a key advantage of transformers over traditional methods: their inclusion of language-wide linguistic regularities that might be present in unseen test data. It has already been demonstrated by Gururangan et al. [18] that BERT is more effective when the target task domain is similar to the ones on which BERT was trained (English language books [61] and English Wikipedia). Active learning also reduces transformer advantage, by reducing the labeling cost to learn corpus-specific vocabulary and regularities. Finally, the short useful life of TAR models means limited opportunity to amortize training cost, raising questions about the large computational cost of task fine-tuning for transformers.

The recent TREC-COVID evaluation provides evidence both in favor and against transformers. A SciBERT-based zero-shot reranker of BM25-based text retrieval topped several of the Round 1 evaluation measures [31, 32]. On the other hand, another transformer-based effort (which omitted language model fine tuning) struggled [29], a number of other deep learning efforts had mediocre effectiveness, and classic linear models based on lexical features and trained by active learning were highly competitive (leading on one measure) [31, 48]. Recently, Ioannidis [21] evaluated BERT and PubMedBERT [17] on CLEF eHealth Technology Assisted Reviews in Empirical Medicine Task [22, 23]. Despite the claim, Ioannidis [21] considered a simple ranking and classification setting instead of an iterative task.

Against this context, we provide the first demonstration of fine-tuned transformer-based models in the TAR transductive active learning setting. We use BERT [12] as a representative transformer. We fine-tune the language model to each of two (unlabeled) task corpora using a masked language modeling objective, kick off each prioritization task on that corpus with a single positive example, and do task fine-tuning of BERT on each TAR active learning iteration.

⁵ <https://www.nexlp.com/blog/nexbert-story-engine-cloud>

Surprisingly, despite the past success stories of BERT in dramatically advancing the retrieval effectiveness, in our work, we found that it only performs on par with the simple logistic regression model due to the transductivity of HRR. On the contrary, under certain scenarios, the BERT model reduces the total reviewing cost, which is the primary objective of HRR tasks. Given its data-hungry property, this cost reduction is counterintuitive but yet very favorable. We highlight our contributions in the following,

- First, we find that language model fine-tuning to the task corpus before active learning is critical, but also that too much of it can be done.
- Second, we find language model fine-tuning is not a cure-all for domain mismatch. Our fine-tuned BERT model beats linear models on a data set (RCV1-v2) similar to the text types on which BERT was trained, but falls short when operating with very different textual characteristics.
- Finally, we provide a running time analysis to demonstrate the computational overhead for applying BERT.

2 Background

HRR projects typically balance thoroughness versus cost by setting a recall target that is high, but below 100%. Targets such as 80% recall are common in eDiscovery [41] and are sometimes encoded in legal agreements [54]. Systematic review often shoots for 95% recall (on smaller and more homogeneous collections) [22, 23]. Recall is defined as the number of relevant documents found among the reviewed documents, divided by the number of relevant documents in the defined collection of interest (e.g., all emails from a set of employees relevant to a legal case, or all biomedical research papers that have passed a keyword screening).

TAR workflows reduce costs by using iterative active learning to prioritize batches of documents for review. One-phase TAR workflows continue this process until a stopping rule indicates that the reviewed documents have met the recall target [9]. Two-phase workflows have a training phase followed by a classification phase (on the same data set), with review done in both phases [34, 54]. Designing stopping rules that determine as early as possible that a recall target has been reached is an active research area [6, 10, 11, 26, 28, 43, 47, 53], but we design our evaluation to avoid the selection and the error incurred by the stopping rule based on the prior studies in TAR cost evaluation [54].

Evaluation for HRR emphasizes a recall/cost tradeoff rather than the related recall/precision tradeoff. In eDiscovery, *Depth for recall* ($DFR@x$) is the proportion of the collection reviewed to hit a recall target x .⁶ Systematic review uses *Work saved over sampling* ($WSS@x$), which subtracts $DFR@x$ from the expected cost to hit the recall target by random sampling: $WSS@x = x - DFR@x$ [8]. Some early HRR studies also use R-Precision (precision at R where R is the

⁶ <https://www.gibsondunn.com/wp-content/uploads/documents/publications/Evans-Metrics-that-Matter-Inside-Counsel-1.2015.pdf>

number of relevant documents) [16, 41] to capture the effectiveness to lower part of the rank as opposed to Precision at 5 or 10 in adhoc retrieval.

However, these evaluation metrics do not consider the cost of obtaining the labels for training documents. In this study, we adapt the cost evaluation of TAR proposed by Yang et al. [54] to jointly evaluate the effectiveness and the cost of the retrieval results. The total cost of TAR consists of the cost of reviewing (1) the training documents and (2) the minimum number of unreviewed documents ranked by the current classification model for fulfilling the recall target. This cost evaluation approach allows documents in different classes and phases to cost differently, facilitating a more practical HRR evaluation and emphasizing the cost of training the one-time classification model.

Commercial TAR technology relies on traditional text classification approaches such as logistic regression and support vector machines (SVMs) [4, 50], that have been widely studied in both active learning and transductive contexts [9, 26, 34, 53, 54]. However, the state of the art in text classification has moved to transformer-based models such as BERT [12] whose properties in these contexts are less well-understood. This gap in understanding motivates the current study.

3 Adapting BERT for TAR

In this section, we describe the adaption of the BERT model to TAR. On a high level, the BERT language model is fine-tuned on the collection of the retrieval interest. At each active learning iteration, we select a set of documents based on the predictions from the model for human review. The acquired labels are fed to the BERT model to perform classification fine-tuning for learning relevancy.

Since the entire task corpus is available before training in TAR, our first step in applying BERT to TAR was language model fine tuning to that corpus. We used the same unsupervised masked language modeling task originally used to train BERT: randomly masking 15% of the tokens in each sequence and tuning BERT’s parameters to predict the missing tokens [12]. The key question is how much to move BERT’s parameters (encoding the linguistic regularities explicit in a mammoth broad domain corpus) toward the task-specific, but less complete, explicit regularities of the task corpus. Our experiments study this by varying the number of epochs (passes through training set) in language model fine-tuning.

TAR workflows use an active learning method such as relevance feedback [40] or uncertainty sampling [25], where the model trained by supervised learning on iteration $k - 1$ is used to select the batch of documents to be labeled in iteration k . The union of labeled batches for iterations $1 \dots k - 1$ is the training set for iteration k . One random relevant document was selected at the beginning of the process as the *seed document* to initiate the active learning. All labeled documents are used for classification fine-tuning of the BERT model. Documents labeled in earlier iterations are visited more by the model based on this classification fine-tuning process. However, based on our pilot study on only fine-tuning model on the newly labeled documents at each iteration, the results were far worse than on all labeled documents. We use a cross-entropy loss on the binary

class label by adding a dense layer on the [CLS] token on top of BERT. We train for a fixed number of epochs, which previous work on active learning for BERT suggests works as well as choosing epoch number using a validation set [13].

For simple models, training can be done to convergence from scratch on each iteration (as we do for logistic regression and SVMs in our experiments). Classification fine tuning for a transformer is computationally expensive, so we instead use the model trained on iteration $k - 1$ as the starting point for optimization on iteration k . While this potentially gives more influence to examples selected on the first iteration, adaptive example selection by active learning reduces this effect.

4 Experiment Setup

4.1 Data Sets

We simulate TAR reviews on two fully labeled collections widely used in HRR studies [10, 11, 16, 35, 41, 51, 52, 56]: RCV1-v2 [27] and the Jeb Bush emails [16, 41].

RCV1-v2 consists of 804,414 news stories with coding for 658 economic news categories. We use the 45 categories subset established by previous high recall retrieval study [54] that spans across three prevalence and three difficulty bins. Text from the title and body was concatenated and tokenized using WordPiece. Documents are truncated with 512 WordPiece tokens as the leading passages of the news documents usually convey the most important aspects of the news articles [7]. The collection is also downsampled to 20% (160,833 documents) for computational efficiency.

The Jeb Bush collection consists of 274,124 unique emails between the former governor of Florida and his colleagues and constituents. The collection was annotated for 44 political topics for the 2015 and 2016 TREC Total Recall Tracks [16, 41]. Text from the subject line and body were concatenated. As with RCV1-v2, documents with more than 512 WordPiece tokens were truncated, similar to the preprocessing steps used in prior works in email classification [45]. Since the most recent replies and content are presented at the beginning of the email and the trailing parts are often duplicated from other emails, including only the leading passages are usually sufficient. A 50% random sample of the remainder (137,062 documents) was used. All 44 topics are used in the experiment. For consistency, we refer to these topics as categories in the later sections.

The RCV1-v2 news articles are professionally written texts with topics and vocabulary well covered by the book and encyclopedic text used to train BERT. We view HRR on it as an in-domain task for BERT. The Jeb Bush emails (particularly from constituents) vary wildly in style and formality from message to message, and reference many Florida personalities, places, and issues likely to be poorly covered in the BERT pre-training materials. We therefore view it as an out-of-domain task for BERT.

4.2 Software and Evaluation

We implemented the active learning workflow with `libact` [57], an open-source active learning framework. For each category, a randomly selected positive example formed the sample for the first iteration. On each subsequent iteration, 200 documents were sampled using active learning (either relevance feedback or least-confidence uncertainty sampling) by following experiment settings from prior HRR studies [51, 54].

For BERT runs we used the `BERT-base-cased` model.⁷ Masked language model fine-tuning was done with the HuggingFace script `run_mlm.py`,⁸ which uses ADAM with no weight decay and warm up period as the optimizer, and a learning rate of 5×10^{-5} . To test the importance of language model fine-tuning, we vary it from no language model fine-tuning to ten iterations over the corpus.

Then on each active learning iteration, we do classification fine-tuning using the ADAM optimizer with a linear weight decay of 0.01 with 50 warm up steps and initial learning rate of 0.001. All reviewed documents (including the ones previously reviewed) are used to fine-tune the model at each active learning iterations with 20 epochs. All hyperparameters were selected based on a pilot study on one selected category of each collection for maximizing the average R-Precision after 20 active learning iterations. The authors also experimented with fine-tuning the model with only the newly queried documents at each iteration, but the results were worse than fine-tuning on all labeled documents by a large margin.

Logistic regression is served as the baseline in our study and is implemented with `scikit-learn` [37] for comparison. It is widely used in HRR research and commercial software [2, 4, 52, 54]. We use the `scikit-learn` tokenizer and BM25 within document saturated term frequencies as feature values [39, 52]. We use L2 regularization on the logistic losses, with penalty weight 1.0 and fit to convergence with default settings from `scikit-learn`.

For comparison with prior work, we report R-Precision, which is a metric that often reports in high recall retrieval studies [16, 41, 52]. Despite being an effectiveness measure that jointly considers precision and recall, it does not reflect the actual objective of the retrieval task, which is the reviewing cost.

Therefore, our primary evaluation measure is the total optimal reviewing cost of the TAR run [54], which is the sum of reviewing the training documents and the documents ranked by the current classification model to fulfill the recall target. The latter is referred to as the optimal amount of the second phase review and can be considered as an optimal penalty for the one-phase workflow [26, 54]. We report the minimal total cost that occurs during the 20 active learning iterations. Without loss of generality, we use 80% recall target as an example, which is a widely used target in eDiscovery study. Higher targets such as 95% yield similar results.

⁷ <https://huggingface.co/bert-base-cased>

⁸ https://github.com/huggingface/transformers/blob/master/examples/pytorch/language-modeling/run_mlm.py

To emphasize the importance of the underlying classification model in the iterative process, we evaluate with both the uniform cost structure (i.e., no reviewing cost difference between documents) and expensive training cost structure. Without loss of generality, we assume the training documents cost ten times more than documents reviewed during the mass reviewing phase as an example [54]. The expensive training cost structure favors classification models that require less training data for optimizing the total cost, enabling us to distinguish the effectiveness of the classification model further.

4.3 Hardware

The active learning experiments are conducted on a cluster of 48 NVIDIA Titan RTX GPUs with 24 GB memory on each. One active learning run (one topic, one sampling strategy, one pretrained BERT model) took on average 18 hours. The entire set of experiments $((45 + 44) \times 5 \times 2 = 890$ runs) took around two weeks on our research cluster. The baseline experiments ran on a single CPU. All logistic regression runs $((45 + 44) \times 2 = 178)$ took around one hour. A detailed running time analysis is presented in the next section.

5 Results and Analysis

In this section, we aim to answer the following research questions: does language model fine-tuning improve the retrieval effectiveness? If so, what is the right amount? How much overhead are we paying for applying BERT?

5.1 Language Model Fine-Tuning

Based on our experimental results, BERT with language model (LM) fine-tuning improves the effectiveness only when the domain of the collection aligns with the domain of the pretraining corpora. In Table 1, the reported cost is the average of the proportional relative cost differences between the baseline logistic regression results and the pretrained BERT model. Since the cost varies between categories, averaging the relative differences prevent the naturally harder tasks (with higher baseline cost) from diluting the aggregated values. The paired t-tests are still conducted between the raw cost with a null hypothesis of identical cost between the BERT and the baseline model. In RCV1-v2, BERT models provide roughly the same R-Precision (0.75 to 0.77) as the baseline logistic regression model regardless of the length of LM fine-tuning, suggesting similar quality at the top of the rank list. On the other hand, BERT models reduce the cost, especially when the training documents cost more to review, compared to the baseline model when the amount of LM fine-tuning is *just right* (10% to 15% on average with expensive training cost structure). In our experiments, the *goldilock* amount is five epochs. However, this amount varies with collection size and other characteristics of the task, which is discussed later in the section. Since reducing the total cost of TAR requires improving the overall rank list [54], these results

Table 1: Averaged evaluation results on the in-domain RCV1-v2 collection and off-domain Jeb Bush collection over categories. Numbers in parentheses are the relative difference between the baseline logistic regression model (LR). Both uniform and expensive training cost (Exp. Train.) values are the relative cost difference between the BERT and the logistic regression models. Values larger than 1.0 indicate higher costs than the baseline. * indicates the statistical significance with 95% confidence between the corresponding pretrained BERT model and the baseline conducted by paired t-test with Bonferroni corrections within each evaluation metric.

Collection	LMFT Epoch	R-Precision (\uparrow)		Uni. Cost (\downarrow)		Exp. Train. (\downarrow)	
		Relevance	Uncertainty	Rel.	Unc.	Rel.	Unc.
In-domain RCV1-v2	LR	0.788 (1.00)	0.760 (1.00)	1.000	1.000	1.000	1.000
	0	0.752 (0.95)	0.756 (0.99)	1.309	1.015	1.178	*0.873
	1	0.757 (0.96)	0.768 (1.01)	1.199	1.039	1.012	0.894
	2	0.759 (0.96)	0.766 (1.01)	1.289	1.028	1.067	0.890
	5	0.756 (0.96)	0.784 (1.03)	1.173	0.893	0.980	*0.844
	10	0.764 (0.97)	0.765 (1.01)	1.192	0.950	1.051	*0.878
Off-domain Jeb Bush	LR	0.904 (1.00)	0.857 (1.00)	1.000	1.000	1.000	1.000
	0	*0.724 (0.80)	*0.719 (0.84)	6.877	5.834	*2.717	*2.194
	1	0.811 (0.90)	0.816 (0.95)	4.678	2.896	1.756	1.413
	2	0.812 (0.90)	0.808 (0.94)	3.257	3.141	1.675	1.446
	5	*0.810 (0.90)	0.813 (0.95)	3.261	2.665	1.583	1.322
	10	0.805 (0.89)	0.815 (0.95)	3.922	2.943	1.601	1.361

suggest that the BERT model with five epochs of LM fine-tuning provides a consistent improvement on the entire ranking.

If the target collection is off-domain compared to the original pre-trained corpora, BERT models cannot provide an effective classifier, even worse than simple linear logistic regression. The averaged values in the Jeb Bush collection suggest worse effectiveness (lower R-Precision and higher cost) despite that the differences are not statistically significant. However, the time overhead and computational burden of applying neural models such as BERT are massive compared to linear models. The inability to provide more effective retrieval results is already a failure. Note that the effectiveness of the BERT models could eventually improve over the baseline with more LM fine-tuning despite the decrement from five to ten epochs; the computational cost would be uneconomical. Running time analysis is presented later in this section.

Therefore, applying BERT models to TAR is not guaranteed to lead to more effective retrieval results. The alignment of the domain between the collections and the amount of LM fine-tuning constitutes a considerable variation of the effectiveness, which is counterintuitive to the common wisdom that continuing fine-tuning would result in better results [19]. If *just-right* hyperparameter is

Table 2: Cost of RCV1-v2 categories in each bin under the expensive training cost structure. Values are the relative cost difference between the corresponding BERT and baseline models averaged over the five categories in each bin.

Difficulty	Prevalence	Relevance					Uncertainty				
		0	1	2	5	10	0	1	2	5	10
Hard	Rare	0.918	0.988	1.011	0.997	1.048	1.044	0.843	0.801	0.664	0.870
	Medium	0.774	0.773	0.699	0.622	0.639	0.594	0.670	0.602	0.612	0.613
	Common	0.832	0.856	0.850	0.798	0.755	0.815	0.849	0.842	0.755	0.751
Medium	Rare	0.932	0.916	0.904	0.784	0.951	0.770	0.903	0.868	0.794	0.828
	Medium	1.275	1.311	1.293	1.175	1.229	1.065	1.199	1.203	1.088	1.211
	Common	0.951	0.778	0.830	0.743	0.820	0.946	0.945	0.933	0.845	0.915
Easy	Rare	1.688	1.225	1.362	1.430	1.540	0.587	0.638	0.702	0.632	0.621
	Medium	1.897	1.189	1.182	1.103	1.263	1.073	0.936	1.015	1.069	0.982
	Common	1.336	1.070	1.474	1.165	1.218	0.960	1.061	1.047	1.136	1.112

not available for the task, which is usually the case for real-world applications, applying BERT models could result in inferior results.

5.2 *Just-Right* Varies Across Tasks

The 45 categories selected from RCV1-v2 enable further analysis into the effect of the task characteristics. Table 2 demonstrates the averaged relative cost differences compared to the baseline model in each category bin under the expensive training cost structure. Since each bin only contains five runs (five categories), statistical tests are non-indicative; hence omitted.

For relevance feedback where training documents are selected from the top of the rank, BERT models usually perform similarly to logistic regression models with a few exceptions. BERT models are more helpful in hard categories than easy ones since the relevancy is often beyond simple token matching in the hard ones, yielding a 20% to 30% cost reduction. However, when the task is hard and the relevant documents are rare, BERT models are no better than simple linear models, even with more LM fine-tuning.

For uncertainty sampling, where the training documents are ones that the model is the least certain about (with predicted probability around 0.5), BERT models provide a substantial improvement of 20% to 40% cost reduction in both hard and rare categories. These results indicate that BERT models are still more effective in challenging situations – either extremely unbalanced training set or relevancy requires subtle semantic understanding. These are cases where linear models tend to fail if no specific treatments to the collection are made.

However, even in these cases where BERT models demonstrate a clear advantage over the linear models, the amount of LM fine-tuning is still critical. The optimal length of LM fine-tuning varies across difficulty and prevalence bins, which were developed by Yang et al. [54]. For example, the best performing pre-trained model for the *hard-medium* bin is no LM fine-tuning (0.5935, i.e., 41%

Table 3: Running time in minutes. Running time for LM fine-tuning (LMFT) is agnostic to the categories. Time reported for TAR is the average running time for each category to complete a 20-iteration TAR process, which consists of 20 classification fine-tuning (or training for logistic regression) and scoring the entire collection. Values in parentheses are the standard deviations of the averaged time.

Collection	LMFT	Relevance			Uncertainty		
	Epoch	LMFT	TAR	Total	LMFT	TAR	Total
In-domain RCV1-v2	0	–	1095	1095 (31.49)	–	1098	1098 (28.97)
	1	98	1094	1192 (17.55)	98	1102	1200 (20.57)
	2	196	1096	1292 (20.53)	196	1103	1296 (28.33)
	5	490	1103	1593 (23.57)	490	1103	1593 (19.93)
	10	980	1101	2081 (20.26)	980	1105	2085 (20.90)
	LR	–	0.32	0.34 (0.04)	–	0.38	0.38 (0.05)
Off-domain Jeb Bush	0	–	999	999 (19.02)	–	1008	1008 (19.48)
	1	98	1002	1100 (16.56)	98	1003	1101 (24.78)
	2	196	1002	1198 (15.03)	196	1002	1198 (21.80)
	5	490	1007	1497 (19.34)	490	1004	1494 (27.37)
	10	981	996	1977 (22.48)	981	1006	1987 (26.67)
	LR	–	0.33	0.33 (0.04)	–	0.41	0.41 (0.06)

cost reduction). However, LM fine-tuning for five epochs gives us the lowest cost (0.6637) and seems to be the minimum for *hard-rare*. For *hard-common*, more fine-tuning tends to be consistently improving the model with the lowest cost (0.7512) occurred at ten epochs in our experiment. The trend is different for medium and easy difficulty bins.

Beyond minimum cost during the run, the trajectory of cost over the iterations also varies among different numbers of LM fine-tuning epochs. For the *hard-rare* category (I65100) in Figure 1(a), the transition from the trajectory of 1 epoch of LM fine-tuning to 2 is not smooth and the shape is nowhere similar. The *hard-common* category (I81501 in Figure 1(c)) also convey no clear relationship between different number of LM fine-tuning epochs.

While BERT models provide significant improvement over the failure cases such as the *medium-rare* category (I42600, Figure 1(d)) and *hard-medium* category (C182, Figure 1(b)), the trajectory is nearly identical for the easy categories regardless of the LM fine-tuning epochs, especially with relevance feedback.

Despite making no clear conclusion on the optimal amount of LM fine-tuning, we observe that this hyperparameter is critical and independent of the collection. All TAR runs in Table 2 are based on the same 20% subset of RCV1-v2 collection but with different categories. This poses a challenge for TAR practitioners when applying BERT or potentially other transformer-based classification models to projects: the joint effect of this hyperparameter and the characteristics of the task is so large that it ranges from extremely helpful (50% cost reduction in *hard-medium* categories using uncertainty sampling without LM fine-tuning) to large

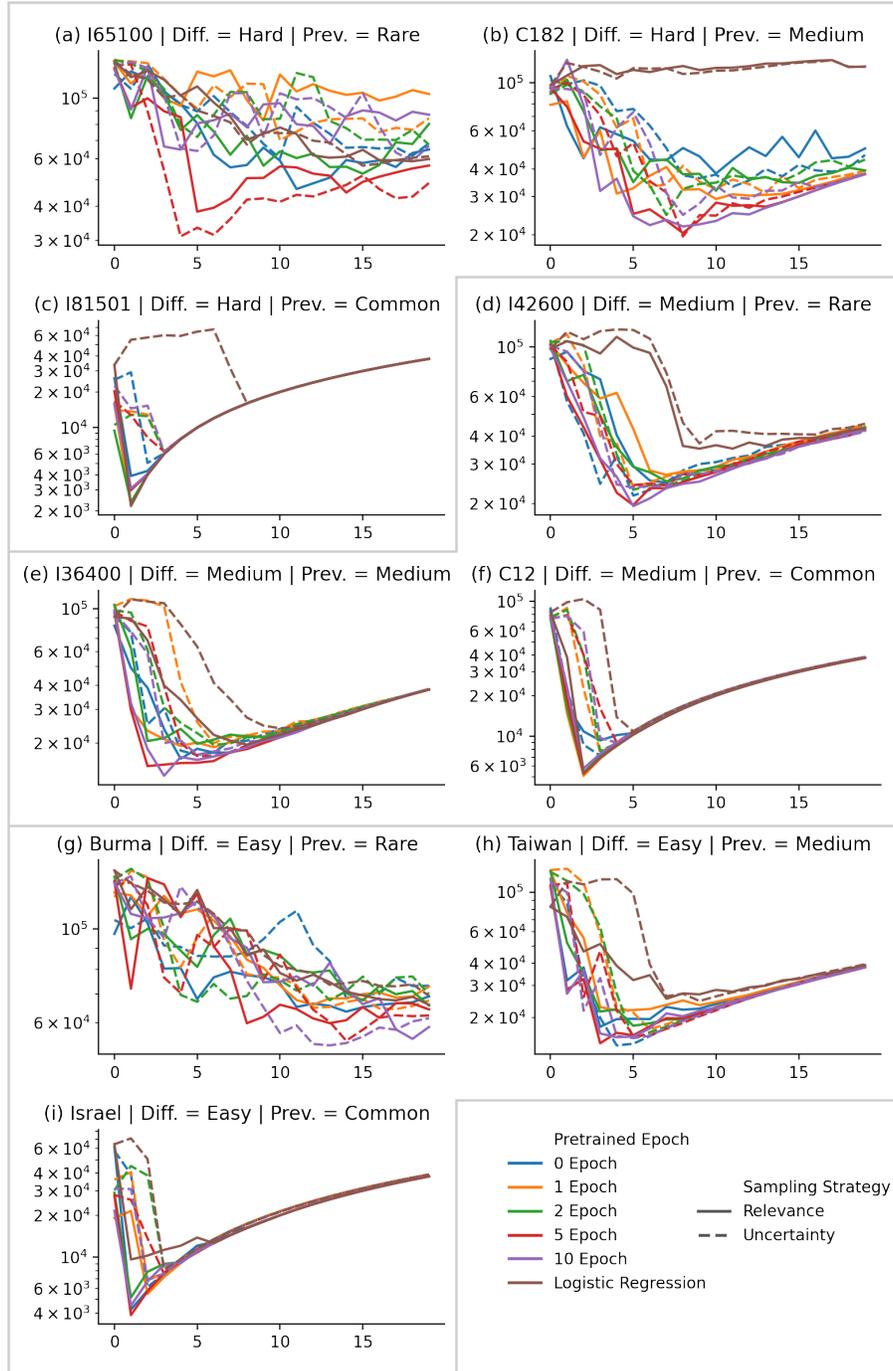


Fig. 1: Example total cost of TAR runs on RCV1-v2 collection over the rounds with expensive training cost structure. The y-axis is the total cost in log-scaled to demonstrate the differences and the x-axis is the number of TAR rounds.

cost overhead (89% cost overhead in *easy-medium* categories using relevance feedback without LM fine-tuning). Understanding the characteristics of the task remains crucial but challenging without sufficient annotations, which is one of the purposes for applying TAR.

5.3 Running Time

Finally, we analyze the running time of the TAR runs. In Table 3, the computational overhead of applying BERT is massive. While the training and scoring of the collection during TAR using logistic regression takes on average 20 to 25 seconds (0.32 to 0.41 minutes), the BERT model takes around 18 hours (1100 minutes). The majority of the time was spent on scoring the collection, which takes around 40 minutes at each iteration. The LM fine-tuning is done before the TAR iterative process, taking around 100 minutes per epoch for the collections we experimented with.

In real high recall retrieval applications where the iterative process spans weeks or even months, each round of reviewing documents takes around half a day. Adding one hour overhead to each iteration is potentially acceptable. However, for smaller projects, this significant time overhead could directly prevent BERT from applying. The computational cost for applying BERT is also not amortized to millions of queries after deployment. Spending 18 hours training a single-usage model in exchange for a mild effectiveness improvement could be unnecessary overhead for many HRR projects.

6 Summary and Future Works

We evaluated the effectiveness of TAR with pre-trained BERT as the underlying predictive model. Before entering active learning, the pre-trained BERT model is fine-tuned by the masked language modeling objective with several epochs. Through experiments, we show that the amount of LM fine-tuning is critical even on an in-domain task. For tasks with out-of-domain text, as compared to the BERT model pre-training corpora, LM fine-tuning requires more training, potentially with other similar corpora. Without proper LM fine-tuning, BERT models underperform typical linear models used with TAR. However, our experiments also show that category characteristics also impact how beneficial the BERT models are and the large computational overhead might discourage the application of BERT in real-world HRR projects.

As the first study of applying transformer models to TAR, there is still much to explore in this area. In the future, we will investigate a wider variety of HRR tasks and sampling strategies that are designed for neural models such as Monte Carlo Dropout [14] and Discriminative Active Learning [15]. A comprehensive approach for handling documents with more than 512 tokens should also be studied. Pre-training a transformer model with large email corpora would benefit the community as many eDiscovery tasks are working on emails. Whether the pre-training corpora would carry biases into the final retrieval results in each TAR project is also demanding for future research.

References

1. Adhikari, A., Ram, A., Tang, R., Lin, J.: Docbert: Bert for document classification. arXiv preprint arXiv:1904.08398 (2019)
2. Bannach-Brown, A., Przybyła, P., Thomas, J., Rice, A.S., Ananiadou, S., Liao, J., Macleod, M.R.: Machine learning algorithms for systematic review: reducing workload in a preclinical review of animal studies and reducing human screening error. *Systematic reviews* **8**(1), 1–12 (2019)
3. Baron, J., Losey, R., Berman, M.: *Perspectives on Predictive Coding: And Other Advanced Search Methods for the Legal Practitioner*. American Bar Association, Section of Litigation (2016), ISBN 9781634256582, URL <https://books.google.com/books?id=TdJ2AQAACAAJ>
4. Brown, S.: Peeking inside the black box: A preliminary survey of technology assisted review (tar) and predictive coding algorithms for ediscovery. *Suffolk J. Trial & App. Advoc.* **21**, 221 (2015)
5. Brown, T.B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D.M., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., Amodei, D.: Language models are few-shot learners (2020)
6. Callaghan, M.W., Müller-Hansen, F.: Statistical stopping criteria for automated screening in systematic reviews. *Systematic Reviews* **9**(1), 1–14 (2020)
7. Catena, M., Frieder, O., Muntean, C.I., Nardini, F.M., Perego, R., Tonelotto, N.: Enhanced news retrieval: Passages lead the way! In: *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 1269–1272 (2019)
8. Cohen, A.M., Hersh, W.R., Peterson, K., Yen, P.Y.: Reducing workload in systematic review preparation using automated citation classification. *Journal of the American Medical Informatics Association* **13**(2), 206–219 (2006)
9. Cormack, G.F., Grossman, M.F.: Evaluation of machine-learning protocols for technology-assisted review in electronic discovery. *SIGIR 2014* pp. 153–162 (2014), <https://doi.org/10.1145/2600428.2609601>.
10. Cormack, G.V., Grossman, M.R.: Engineering Quality and Reliability in Technology-Assisted Review. In: *SIGIR*, pp. 75–84, ACM Press, Pisa, Italy (2016), ISBN 978-1-4503-4069-4, <https://doi.org/10.1145/2911451.2911510>, URL <http://dl.acm.org/citation.cfm?doid=2911451.2911510>, 00024
11. Cormack, G.V., Grossman, M.R.: Scalability of continuous active learning for reliable high-recall text classification. In: *Proceedings of the 25th ACM international on conference on information and knowledge management*, pp. 1039–1048 (2016)
12. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805 (2018)
13. Ein-Dor, L., Halfon, A., Gera, A., Shnarch, E., Dankin, L., Choshen, L., Danilevsky, M., Aharonov, R., Katz, Y., Slonim, N.: Active learning for bert:

- An empirical study. In: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), pp. 7949–7962 (2020)
14. Gal, Y., Ghahramani, Z.: Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In: international conference on machine learning, pp. 1050–1059, PMLR (2016)
 15. Gissin, D., Shalev-Shwartz, S.: Discriminative active learning. arXiv preprint arXiv:1907.06347 (2019)
 16. Grossman, M.R., Cormack, G.V., Roegiest, A.: Trec 2016 total recall track overview. (2016)
 17. Gu, Y., Tinn, R., Cheng, H., Lucas, M., Usuyama, N., Liu, X., Naumann, T., Gao, J., Poon, H.: Domain-specific language model pretraining for biomedical natural language processing. arXiv preprint arXiv:2007.15779 (2020)
 18. Gururangan, S., Dang, T., Card, D., Smith, N.A.: Variational pretraining for semi-supervised text classification. arXiv preprint arXiv:1906.02242 (2019)
 19. Gururangan, S., Marasović, A., Swayamdipta, S., Lo, K., Beltagy, I., Downey, D., Smith, N.A.: Don’t stop pretraining: Adapt language models to domains and tasks. In: Proceedings of ACL (2020)
 20. Hou, Y., Che, W., Lai, Y., Zhou, Z., Liu, Y., Liu, H., Liu, T.: Few-shot slot tagging with collapsed dependency transfer and label-enhanced task-adaptive projection network. In: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, pp. 1381–1393 (2020)
 21. Ioannidis, A.: An analysis of a bert deep learning strategy on a technology assisted review task. arXiv preprint arXiv:2104.08340 (2021)
 22. Kanoulas, E., Li, D., Azzopardi, L., Spijker, R.: Clef 2017 technologically assisted reviews in empirical medicine overview. In: CEUR workshop proceedings, vol. 1866, pp. 1–29 (2017)
 23. Kanoulas, E., Li, D., Azzopardi, L., Spijker, R.: Clef 2018 technologically assisted reviews in empirical medicine overview. CEUR Workshop Proceedings **2125** (July 2018), URL <https://strathprints.strath.ac.uk/66446/>
 24. Kanoulas, E., Li, D., Azzopardi, L., Spijker, R.: Clef 2019 technology assisted reviews in empirical medicine overview. In: CEUR workshop proceedings, vol. 2380 (2019)
 25. Lewis, D.D., Gale, W.A.: A sequential algorithm for training text classifiers. In: SIGIR 1994, pp. 3–12 (1994)
 26. Lewis, D.D., Yang, E., Frieder, O.: Certifying one-phase technology-assisted reviews. In: Proceedings of 30th ACM International Conference on Information and Knowledge Management (2021)
 27. Lewis, D.D., Yang, Y., Rose, T.G., Li, F.: RCV1: A New Benchmark Collection for Text Categorization Research. JMLR **5**, 361–397 (2004)
 28. Li, D., Kanoulas, E.: When to stop reviewing in technology-assisted reviews: Sampling from an adaptive distribution to estimate residual relevant documents. ACM Transactions on Information Systems (TOIS) **38**(4), 1–36 (2020)
 29. Lima, L.C., Hansen, C., Hansen, C., Wang, D., Maistro, M., Larsen, B., Simonsen, J.G., Lioma, C.: Denmark’s participation in the search engine trec covid-19 challenge: Lessons learned about searching for precise biomedical scientific information on covid-19. arXiv preprint arXiv:2011.12684 (2020)

30. Liu, M., Tu, Z., Zhang, T., Su, T., Wang, Z.: Ltp: A new active learning strategy for crf-based named entity recognition. arXiv preprint arXiv:2001.02524 (2020)
31. MacAvaney, S., Cohan, A., Goharian, N.: Sledge: a simple yet effective baseline for covid-19 scientific knowledge search. arXiv e-prints pp. arXiv-2005 (2020)
32. MacAvaney, S., Cohan, A., Goharian, N.: Sledge-z: A zero-shot baseline for covid-19 literature search. In: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (2020), <https://doi.org/10.18653/v1/2020.emnlp-main.341>, URL <https://arxiv.org/abs/2010.05987>
33. MacAvaney, S., Yates, A., Cohan, A., Goharian, N.: Cedr: Contextualized embeddings for document ranking. In: Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 1101–1104 (2019), <https://doi.org/10.1145/3331184.3331317>, URL <https://arxiv.org/abs/1904.07094>
34. McDonald, G., Macdonald, C., Ounis, I.: Active learning strategies for technology assisted sensitivity review. In: European Conference on Information Retrieval, pp. 439–453, Springer (2018)
35. Oard, D.W., Sebastiani, F., Vinjumar, J.K.: Jointly Minimizing the Expected Costs of Review for Responsiveness and Privilege in E-Discovery. *ACM Transactions on Information Systems* **37**(1), 1–35 (Nov 2018), ISSN 10468188, <https://doi.org/10.1145/3268928>, URL <http://dl.acm.org/citation.cfm?doid=3289475.3268928>, 00000
36. Oard, D.W., Webber, W.: Information retrieval for e-discovery. *Information Retrieval* **7**(2-3), 99–237 (2013)
37. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E.: Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* **12**, 2825–2830 (2011)
38. Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., Liu, P.J.: Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research* **21**, 1–67 (2020)
39. Robertson, S., Zaragoza, H., et al.: The probabilistic relevance framework: Bm25 and beyond. *Foundations and Trends® in Information Retrieval* **3**(4), 333–389 (2009)
40. Rocchio, J.J.: Relevance feedback in information retrieval (1971)
41. Roegiest, A., Cormack, G.V.: Trec 2015 total recall track overview (2015)
42. Ruthven, I., Lalmas, M.: A survey on the use of relevance feedback for information access systems. *Knowledge engineering review* **18**(2), 95–145 (2003)
43. Saha, T.K., Hasan, M.A., Burgess, C., Habib, M.A., Johnson, J.: Batch-mode active learning for technology-assisted review. In: 2015 IEEE International Conference on Big Data (Big Data), pp. 1134–1143 (Oct 2015), <https://doi.org/10.1109/BigData.2015.7363867>, 00003

44. Shelmanov, A., Liventsev, V., Kireev, D., Khromov, N., Panchenko, A., Fedulova, I., Dylov, D.V.: Active learning with deep pre-trained models for sequence tagging of clinical and biomedical texts. In: 2019 IEEE International Conference on Bioinformatics and Biomedicine (BIBM), pp. 482–489, IEEE (2019)
45. Shu, K., Mukherjee, S., Zheng, G., Awadallah, A.H., Shokouhi, M., Dumais, S.: Learning with weak supervision for email intent detection. In: Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 1051–1060 (2020)
46. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. ArXiv [abs/1706.03762](https://arxiv.org/abs/1706.03762) (2017)
47. Wallace, B.C., Trikalinos, T.A., Lau, J., Brodley, C., Schmid, C.H.: Semi-automated screening of biomedical citations for systematic reviews. *BMC bioinformatics* **11**(1), 55 (2010)
48. Wang, X.J., Grossman, M.R., Hyun, S.G.: Participation in trec 2020 covid track using continuous active learning. arXiv preprint arXiv:2011.01453 (2020)
49. Wang, Y., Che, W., Guo, J., Liu, Y., Liu, T.: Cross-lingual bert transformation for zero-shot dependency parsing. 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), pp. 5725–5731 (2019)
50. Yang, E., Grossman, D., Frieder, O., Yurchak, R.: Effectiveness results for popular e-discovery algorithms. In: Proceedings of the 16th edition of the International Conference on Artificial Intelligence and Law, pp. 261–264 (2017)
51. Yang, E., Lewis, D.D., Frieder, O.: A regularization approach to combining keywords and training data in technology-assisted review. In: Proceedings of the Seventeenth International Conference on Artificial Intelligence and Law, pp. 153–162 (2019)
52. Yang, E., Lewis, D.D., Frieder, O.: Text retrieval priors for bayesian logistic regression. In: Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 1045–1048 (2019)
53. Yang, E., Lewis, D.D., Frieder, O.: Heuristic stopping rules for technology-assisted review. In: Proceedings of the 21st ACM Symposium on Document Engineering (2021)
54. Yang, E., Lewis, D.D., Frieder, O.: On minimizing cost in legal document review workflows. In: Proceedings of the 21st ACM Symposium on Document Engineering (2021)
55. Yang, E., Lewis, D.D., Frieder, O.: Tar on social media: A framework for online content moderation. In: 2nd International Conference on Design of Experimental Search & Information REtrieval Systems (2021)
56. Yang, E., Lewis, D.D., Frieder, O., Grossman, D., Yurchak, R.: Retrieval and richness when querying by document. In: International Conference on Design of Experimental Search & Information REtrieval Systems (2018)

57. Yang, Y.Y., Lee, S.C., Chung, Y.A., Wu, T.E., Chen, S.A., Lin, H.T.: libact: Pool-based active learning in python. Tech. rep., National Taiwan University (Oct 2017), URL <https://github.com/ntucllab/libact>, available as arXiv preprint <https://arxiv.org/abs/1710.00379>
58. Yang, Z., Wang, Y., Chen, X., Liu, J., Qiao, Y.: Context-transformer: tackling object confusion for few-shot detection. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 34, pp. 12653–12660 (2020)
59. Zhang, L., Zhang, L.: An ensemble deep active learning method for intent classification. In: Proceedings of the 2019 3rd International Conference on Computer Science and Artificial Intelligence, pp. 107–111 (2019)
60. Zhu, J., Wang, H., Hovy, E., Ma, M.: Confidence-based stopping criteria for active learning for data annotation. *ACM Transactions on Speech and Language Processing (TSLP)* **6**(3), 1–24 (2010)
61. Zhu, Y., Kiros, R., Zemel, R., Salakhutdinov, R., Urtasun, R., Torralba, A., Fidler, S.: Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. In: Proceedings of the IEEE international conference on computer vision, pp. 19–27 (2015)