

Adapting Event Extractors to Medical Data: Bridging the Covariate Shift

Aakanksha Naik¹

Jill Lehman²

Carolyn Rosé¹

¹Language Technologies Institute, Carnegie Mellon University

²Human-Computer Interaction Institute, Carnegie Mellon University

{anaik, jfl, cprose}@andrew.cmu.edu

Abstract

We tackle the task of adapting event extractors to new domains without labeled data, by aligning the marginal distributions of source and target domains. As a testbed, we create two new event extraction datasets using English texts from two medical domains: (i) clinical notes, and (ii) doctor-patient conversations. We test the efficacy of three marginal alignment techniques: (i) adversarial domain adaptation (ADA), (ii) domain adaptive fine-tuning (DAFT), and (iii) a new instance weighting technique based on language model likelihood scores (LIW). LIW and DAFT improve over a no-transfer BERT baseline on both domains, but ADA only improves on notes. Deeper analysis of performance under different types of shifts (e.g., lexical shift, semantic shift) explains some of the variations among models. Our best-performing models reach F1 scores of 70.0 and 72.9 on notes and conversations respectively, using *no* labeled target data.

1 Introduction

Events are an important phenomenon in the field of computational semantics. They offer an intuitive mechanism for constructing structured representations of text, which can be used for downstream tasks such as question answering and summarization. Events also embody a crucial function of language: the ability to report happenings. Narratives from many diverse domains (e.g., news articles, literary texts, clinical notes) use events as basic building blocks. These characteristics make event extraction a key sub-task of interest for text understanding pipelines in multiple domains. Despite its importance, building high-performing and generalizable systems for event extraction has remained an elusive goal. One of the major hurdles is that the notion of what counts as an *important event* is usually task-specific or domain-specific (sometimes both). For example, to build a system that

can track a patient’s disease progression from clinical notes, event extractors only need to focus on extracting medical events relevant to that illness. This task/domain specificity has encouraged prior work to focus on specific event types (Grishman and Sundheim, 1996; Doddington et al., 2004; Kim et al., 2008) or domains (Pustejovsky et al., 2003b; Sims et al., 2019). Owing to this narrow focus, supervised event extractors often fail to adapt to new domains or event types (Keith et al., 2017). Unsupervised event extractors that use syntactic rule-based modules (Saurí et al., 2005; Chambers et al., 2014), conversely, have a tendency to over-generate by labeling most verbs and nouns as events.

In this work, we try to achieve a balance between these extremes by *adapting* event extractors using unsupervised domain adaptation techniques. We also study the behavior of these techniques under various types of linguistic shifts (e.g., lexical shift, semantic shift) to gain insight into differences among them. Exploring adaptability under no (or little) supervision is crucial, since sourcing annotated data for new domains, especially medical texts, can be expensive and time-consuming. Following prior work, we formulate event extraction as the task of labeling **triggers**, i.e., words which instantiate an event (Linguistic Data Consortium, 2005). For example, in the sentence “She was diagnosed with cancer,” *diagnosed* and *cancer* are triggers, referring to “diagnosis” and “illness” events respectively. Throughout our work, we model event trigger labeling as token-level classification.

To test adaptability, we create new event extraction test sets using English texts from two diverse medical domains: (i) clinical notes, and (ii) doctor-patient conversations. We develop comprehensive event annotation guidelines, based on TimeML (Pustejovsky et al., 2003a) and Thyme-TimeML (Styler IV et al., 2014) (§3), and use them to annotate 45 documents from each domain. As a baseline,

we train a BERT-based event extraction model on English news articles from TimeBank (Pustejovsky et al., 2003b), which is labeled using TimeML, and test its performance on our datasets. To improve this out-of-domain baseline performance, we tackle the problem of covariate shift, i.e., differences between marginal distributions of source (news) and target domains (notes or conversations). We experiment with three marginal alignment techniques: (i) adversarial domain adaptation (ADA) (Ganin and Lempitsky, 2015), (ii) domain-adaptive fine-tuning (DAFT) (Han and Eisenstein, 2019), and (iii) a new instance weighting scheme using language model likelihood scores (LIW).

Our results show that DAFT and LIW improve over BERT on both domains, whereas ADA only improves on notes. Across domains, ADA and DAFT perform best on notes and conversations respectively. To probe why some techniques are better at addressing certain source-target domain pairs, we analyze model performance on various types of covariate shifts (e.g., lexical shift, semantic shift). Our analysis uncovers interesting patterns such as varying ability of models to leverage subword morphology to generalize to technical terms, and LIW’s performance improvement on long-term state events (e.g., chronic illnesses). Our best models achieve F1 scores of 70.0 and 72.9 on notes and conversations respectively with *no* training data.¹

2 Related Work

2.1 Event Extraction

Most prior event extraction work has focused on news articles, resulting in the development of several datasets (Onyshkevych et al., 1993; Grishman and Sundheim, 1996; Pustejovsky et al., 2003b; Doddington et al., 2004; Lee et al., 2012; Cybulska and Vossen, 2014; Mitamura et al., 2016). Recently, event extraction has also been explored in other domains such as biology (Wattarujekrit et al., 2004; Kim et al., 2008, 2009; Berant et al., 2014), Wikipedia articles (Araki and Mitamura, 2018), social media data (Ritter et al., 2012; Li et al., 2014; Jain et al., 2016) and literary novels (Sims et al., 2019). Aside from data domain, event extraction paradigms (both datasets and tools) differ along three major axes: (i) event extraction granularity, (ii) event representation, and (iii) event categoriza-

¹Annotated clinical notes and all code associated with this work can be found at: <https://github.com/aakanksha19/MedicalEventExtraction>.

tion (ontology). We briefly describe these axes to contextualize our choice of event paradigm.

Event extraction granularity divides extraction paradigms into two types: (i) document-level paradigms that assume that a piece of text refers to a single event (Grishman and Sundheim, 1996), and (ii) sentence-level paradigms that assume that a single sentence describes one or more events. Event representation also divides extraction paradigms into two types: (i) span-based paradigms that represent events by marking text spans that refer to events, called **triggers** or **nuggets** (Linguistic Data Consortium, 2005; Mitamura et al., 2015; O’Gorman et al., 2016), and (ii) structured paradigms that represent events by marking text spans and adding additional arguments (e.g., participants, location etc.) to create a structured template (Grishman and Sundheim, 1996). Event categorization divides extraction paradigms into: (i) ontology-driven paradigms that are limited to specific event types (Grishman and Sundheim, 1996; Doddington et al., 2004), and (ii) ontology-free paradigms that do not place type restrictions (Pustejovsky et al., 2003b; Araki and Mitamura, 2018).

We use a sentence-level, span-based, ontology-free event extraction paradigm. Sentence-level extraction suits our domains of interest since notes and conversations tend to discuss multiple events. Span-based and ontology-free extraction allows us to develop adaptable coding guidelines since event arguments and types are usually domain-specific or task-specific. This adaptability sets our work apart from other prior work on medical event extraction such as adverse drug event extraction (Nikfarjam et al., 2015; Sarker and Gonzalez, 2015; Cocos et al., 2017; Henry et al., 2020) and personal event extraction from online support groups (Wen et al., 2013; Naik et al., 2017), which focus on specific event types. Our guidelines draw heavily from the Thyme-TimeML guidelines (Styler IV et al., 2014) used by the Clinical TempEval challenges on event ordering in clinical notes (Bethard et al., 2015, 2016, 2017),² but also cover event extraction in a novel domain: doctor-patient conversations.

2.2 Unsupervised Domain Adaptation

Unsupervised domain adaptation is the task of transferring a model from a source domain to a target domain, using only unlabeled data from the target domain, by aligning source and target distri-

²We provide a detailed comparison with this work in §3.1.

butions. Early approaches such as structural correspondence learning (SCL) (Blitzer et al., 2006, 2007) tried to solve this by mapping source and target examples into a shared **pivot feature** space, where pivot features are selected to be features that behave the same way for discriminative learning in both domains (e.g., sentiment terms such as *amazing* and *great* show similar behavior for sentiment analysis across domains). With advances in neural representation learning, autoencoder-based methods (Glorot et al., 2011; Chen et al., 2014), neural SCL (Ziser and Reichart, 2017), adversarial domain adaptation (Ganin and Lempitsky, 2015; Ganin et al., 2016) and LM fine-tuning methods (Han and Eisenstein, 2019; Gururangan et al., 2020) have shown success in learning a shared space in which source and target domains are aligned. We propose a new method (LIW) which relies on instance weighting via language model likelihood, and contrast it with adversarial domain adaptation (ADA) and domain adaptive fine-tuning (DAFT). These two techniques have shown promise on sequence labeling tasks (Gui et al., 2017; Han and Eisenstein, 2019; Naik and Rosé, 2020), and offer an interesting contrast between approaches that jointly perform alignment and task training (ADA) and approaches that perform these steps sequentially (DAFT). Comparing all three techniques also provides us the opportunity to study which methods adapt better to different kinds of shifts between source and target domains (e.g., shifts in vocabulary, syntax, etc.).

3 Dataset Creation

To test adaptability of event extraction models, we create a testbed using data from two domains:

1. Clinical Notes: Clinical notes are records documenting physician observations from their interactions with patients. They usually detail various aspects of a patient’s care such as present illness, symptoms, medical history, treatments, and test results. They share a thematic structure, though particular specialties (e.g., cardiology) and institutions often incorporate their own modifications. We collected a set of 4999 de-identified clinical notes from 40 specialties, by scraping [mtsamples](https://www.mtsamples.com/).³ The notes are reference samples provided by various users, with names and dates edited for confidentiality. They are freely available to print, share, link and distribute, as per website policy. Average

length of a clinical note is 652 tokens.

2. Doctor-Patient Conversations: This data contains human-transcribed, de-identified conversations recorded during physician-patient visits. The conversations often follow a similar schema, with patients describing their symptoms, doctors inquiring about ongoing treatments, and then suggesting potential follow-up treatments/tests. We use a proprietary database of 63,540 conversations covering 53 specialties, collected by Abridge AI Inc. Physicians across a variety of specialties are contracted to record natural in-office conversations with their patients who agree to participate in the research by providing verbal and written consent. Recordings are made on a digital recording device or a smartphone application and are uploaded to a secure server where they are scrubbed of all identifiable information, in accordance with the Health Insurance Portability and Accountability Act of 1996 (HIPAA) privacy rule. De-identified recordings are transcribed and stored in a database, which currently contains over 100,000 recordings dating from 2006 to 2017. Average conversation transcript length is 2309 tokens.

These domains exhibit different types of linguistic shifts from the source (news). While both domains exhibit a shift in vocabulary, it is more pronounced in clinical notes since they are written by doctors (experts) who use highly technical terms. Conversely, shifts in syntax are more pronounced in conversations due to the prevalence of repetition, back-channeling, interruptions etc. Semantic shifts are more pronounced in conversations since they contain a higher proportion of hypothetical statements (e.g., when doctors ask questions, make requests or “think out loud”) than both notes and news articles which tend to serve as records of actual events. To better evaluate model performance on linguistic shifts, we control for topical variation across domains by limiting our focus to 3 specialties: Cardiovascular/Pulmonary (Cardio), Obstetrics/Gynaecology (Obgyn) and Hematology/Oncology (Onco). These specialties are well-represented in both notes and conversations, and cover events with a variety of temporalities ranging from intervals with fixed duration (e.g., pregnancy), to intervals with indeterminable endpoints (e.g., long-term cardiac failure). Table 1 gives an overview of the number of notes and conversations in each specialty.

³<https://www.mtsamples.com/>

Specialty	#Notes	#Convos
Cardio	372	4876
Obgyn	160	1784
Onco	90	7177

Table 1: Domain-wise raw data statistics for chosen medical specialties

3.1 Developing Event Annotation Guidelines

We develop a set of coding guidelines for the task of annotating event triggers in documents from these two domains. Our coding guidelines build upon TimeML (Pustejovsky et al., 2003a), a rich specification language for annotation of events and temporal expressions in text,⁴ and Thyme-TimeML (Styler IV et al., 2014), a variant of TimeML developed for clinical notes. We start with these guidelines because they use a syntax-driven domain-agnostic definition of events, allowing for an adaptable annotation scheme. In TimeML, the term *event* refers to situations that *happen* or *occur*, or circumstances in which something *obtains* or *holds true*. This is a broad definition, consistent with Bach’s definition of **eventualities** (Bach, 1986), and the idea of **fluents** (McCarthy, 2002). Events can be expressed in text by means of tensed or untensed verbs, nominalizations, adjectives, predicative clauses or prepositional phrases. TimeML describes rules to annotate events in all these syntactic categories. Styler IV et al. (2014) adapted these rules for clinical notes. They focused on the THYME corpus of 1254 de-identified notes from the Mayo Clinic, representing two fields in oncology: brain cancer and colon cancer. As a first step, we annotate one document from each of our domains following TimeML and Thyme-TimeML rules. During this phase, we identify cases where it is reasonable to deviate from these guidelines.

Deviations from TimeML: Our guidelines⁵ differ from TimeML in their treatment of two categories:

1. Activity patterns: Activity patterns are events that are neither pure generics⁶, nor single events

⁴The complete TimeML coding manual is available here: https://catalog.ldc.upenn.edu/docs/LDC2006T08/timeml_annguide_1.2.1.pdf

⁵Our complete coding manual, including example annotations, is available at: <https://github.com/aakanksha19/MedicalEventExtraction>.

⁶Pure generics are events which discuss illnesses/treatments in general, and are not associated with a specific person and time. For example, “there is a benefit to systemic adjuvant chemotherapy.”

clearly positioned in time. For example, consider the sentence “I *take* my blood pressure regularly.” The event *take* is not grounded in time. It is also not a pure generic event as it is definitely associated with the speaker. Such events are *not* annotated in TimeML. However, in our data, these activity patterns occur frequently in crucial contexts such as taking medications, following lifestyle changes suggested by doctors, measuring vital signs, etc.

2. Long-term states: Because TimeML was geared towards the task of temporal ordering, it strictly restricted annotation of stative events to the following types: (i) states associated with a temporal expression, (ii) states undergoing a change within the document, (iii) states introduced by other events, since those can offer temporal cues, and (iv) states associated with the document creation time. However, many stative events in our data don’t fit within these strict parameters, but are nevertheless important. The most crucial category is states associated with long-term ongoing illnesses (e.g., “The patient has a long history of *COPD*”).

These event categories are not specific to medical domains only. For example, long-term state events might be salient when extracting personal events from biographies.⁷ Similarly activity patterns might be salient when extracting events from scientific procedure manuals.⁸ Considering these scenarios, we add rules to extract these two categories of events. We also expand syntactic rules to cover constructions unique to doctor-patient conversations such as repetition, especially for instructions, and hypothetical event annotation in utterances when doctors are “thinking out loud”.

Deviations from Thyme-TimeML: Our guidelines differ from Thyme-TimeML in their treatment of two categories:

1. Generic events: Thyme-TimeML annotates generic events in sections documenting discussion of risks, plans and alternative strategies. They do so because adding these events to a patient’s clinical timeline could be important from a legal perspective, as they help to establish informed consent and knowledge of risk. We do not annotate pure generics, because we do not perceive any domain-agnostic utility in annotating them. Note that we annotate verbs of discussion and comprehension which are not generics, so we do not fully ignore events associated with patient consent. For exam-

⁷e.g., “Bill Gates is currently *employed* full-time at the Bill and Melinda Gates Foundation.”

⁸“Repeat step 5 daily, over a period of 30 days.”

Domain	Entity κ	Event κ
Notes	0.9117	0.8652
Convos	0.8634	0.8327

Table 2: Inter-annotator agreement on entity and event annotation tasks in both domains, measured using chance-corrected Cohen’s κ

ple, in the sentence “She repeated the potential side effects back to me,” *repeated* is annotated, but *effects* is not. Thyme-TimeML would annotate both.

2. Entities as events: Thyme-TimeML treats some entities and non-events as events in clinical language. Two categories see this shift in semantic interpretation: (i) Medications, and (ii) Disorders. Both categories contribute significant information to a patient’s timeline, and so they are treated as events. Since we are not specifically focused on timeline construction, we do not follow the same reasoning. In particular, medications are not treated as events, while disorders may be treated as events as long as they fit the TimeML definition. To ensure that we do not discard potentially crucial information, we incorporate an additional step in which we annotate entities such as medications, body parts, abnormalities (e.g., rash), etc.

3.2 Annotation Process

After incorporating our modifications, we test our guidelines by having two expert annotators annotate one document from each domain. We see high inter-annotator agreement (measured by chance-corrected Cohen’s κ) on entity and event annotation, in both domains. Table 2 presents the agreement scores. To create our final datasets, we sample 45 documents from each domain (15 from each specialty). Each document is annotated by one expert. Annotation is carried out using the BRAT stand-off markup interface (Stenetorp et al., 2012). Figure 1 shows a sample clinical note annotated with events and entities. Table 3 gives a brief overview of statistics for our datasets, in comparison with TimeBank (news articles) (Pustejovsky et al., 2003b).

4 Methods for Marginal Alignment

To adapt event extraction models with no training data, we tackle the problem of covariate shift, which arises when the marginal distribution (or input distribution) $P(X)$ changes between train and test data. Directly applying a supervised model trained on the training set, to the test set might not

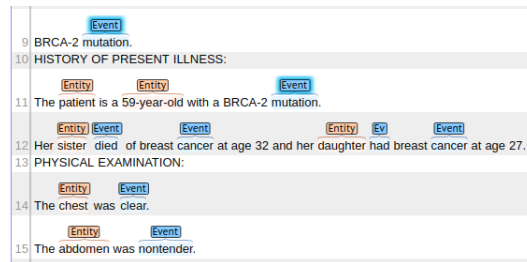


Figure 1: Sample clinical note with entity and event annotation

Statistic	News	Notes	Convos
#Files	54	45	45
#Tokens	18,263	28,935	76,711
#Events	1986	4781	7064
Event Density	10.88%	16.52%	9.21%
Vocab Size	3978	4303	3505
Event Vocab	1015	1588	1472

Table 3: Dataset statistics. Note that the statistics for TimeBank (News) are computed over the test set for fair comparison with our datasets, which are test-only.

perform well due to the gap between training and test distributions. We experiment with several techniques to align the training and test distributions, so that the supervised model transfers better to test data. The techniques can be divided into two types based on the kind of supervision used during alignment: (i) task-guided alignment techniques, and (ii) task-agnostic alignment techniques.

4.1 Task-Guided Alignment Techniques

These techniques jointly optimize for two tasks: (i) aligning training and test distributions, and (ii) training an event extraction model. Since the alignment process receives supervision from task training, we refer to these techniques as task-guided alignment techniques. Under this category, we experiment with adversarial domain adaptation.

Adversarial Domain Adaptation: Adversarial domain adaptation was proposed by Ganin and Lempitsky (2015), who showed its efficacy on sentiment analysis. Recently, Naik and Rosé (2020) showed its utility in transferring event extraction models between two domains: news and literature. The adversarial domain adaptation framework for event extraction contains three components: (i) representation learner (R) which generates token-level representations for a sequence, (ii) event classifier which identifies events (E), and (iii) domain predic-

tor (D) which predicts the domain for the sequence. The key idea is to train R to generate representations which *are predictive* for event identification but *not predictive* for domain prediction, making it more domain-invariant. This aligns training and test distributions by finding a shared feature space in which training and test samples are not distinguishable, while making sure that the feature space is useful for event extraction. The technique relies on an alternating optimization procedure. The first step optimizes D on the domain prediction task, while the second step optimizes both R and E on event identification while subtracting domain prediction loss. For complete mathematical details, we refer the interested reader to Naik and Rosé (2020).

4.2 Task-Agnostic Alignment Techniques

These techniques perform training/test distribution alignment and event extraction training sequentially instead of jointly optimizing them. The alignment process does not receive supervision from task training, so these techniques are task-agnostic. We experiment with the following techniques:

Domain Adaptive Fine-tuning: Domain adaptive fine-tuning has been proposed as an effective technique for unsupervised adaption of sequence labeling models to challenging domains such as Early Modern English and social media (Han and Eisenstein, 2019). This procedure works as follows:

1. Create a large dataset containing equal proportions of sentences from source and target domains. Fine-tune contextualized embeddings using a masked language modeling objective.
2. Using fine-tuned embeddings, train an event extraction model on labeled source data.

In addition to this setup, we experiment with a variant of this procedure, which uses a syntactic objective function. This variant fine-tunes embeddings on the POS tagging task in step 1. The motivation behind this variant is two-fold. First, we observe that event annotation is heavily syntax-driven, allowing delexicalized models (i.e., models using POS tags instead of words) to achieve high performance (§5.2). This indicates that infusing more syntactic awareness into embeddings might help performance on the task. Second, syntax might offer an additional basis for generalization, since sentences that look very different lexically, might follow similar syntactic structures. Intuitively, this variant is similar to syntactic relexicalization which

has shown success in cross-lingual dependency parsing (Duong et al., 2015).

Likelihood-based Instance Weighting: We develop a new instance weighting procedure which uses likelihood scores computed by a language model. Instance selection and instance weighting strategies have frequently been used to perform domain adaptation by correcting for distributional differences (Jiang and Zhai, 2007; Foster et al., 2010; Axelrod et al., 2011; Wang et al., 2017). The main premise is that some samples from out-of-domain data and in-domain data often share some characteristics. Training only on these samples (pruning), or biasing training to focus more on these samples (weighting) can produce models that perform better on out-of-domain data. Motivated by this, our instance weighting procedure works as follows.

Let $S_t = w_1 w_2 \dots w_n$ be a sentence from the in-domain training set. Let O be a language model trained on raw text from the target domain. We first compute the likelihood of sentence S_t under O as $\mathbb{L}_t = P_O(w_1) \prod_{i=2}^n P_O(w_i | w_1 \dots w_{i-1})$, where P_O indicates probability under model O . Then we compute a weight for S_t as follows:

$$\alpha_{S_t} = \frac{\mathbb{L}_t}{\sum_{i=1}^{|N|} \mathbb{L}_i} * |N| \quad (1)$$

where $|N|$ is the size of in-domain training set. This metric gives a higher weight to in-domain sentences that are *more likely* under the target domain language model, up-weighting instances that share more characteristics with target domain sentences. The alpha values are used to weight the loss function, thus biasing the training procedure.

5 Experimental Setup

5.1 Model Details

The goal of our evaluation is to identify which alignment technique works best for each domain, as well as analyze whether there are specific kinds of source-target shifts that some techniques are better equipped to handle. We choose a strong BERT-based baseline model with no transfer, and evaluate the performance of each alignment technique when applied to this baseline.

VERB: Baseline labeling all verbs as events.

DELEX: Fully-delexicalized baseline using POS tag embeddings as features, followed by an MLP.

BERT: Single-layer BiLSTM over BERT embeddings (Devlin et al., 2019), followed by an MLP, similar to the best-performing model on LitBank

Model	In-Domain			Out-of-Domain		
	P	R	F1	P	R	F1
VERB	58.8	66.5	62.5	49.4	41.4	45.0
DELEX	75.0	66.3	70.4	74.4	42.2	53.8
BERT	80.6	86.0	83.2	85.7	55.9	67.6
CBERT	79.2	83.3	81.2	85.8	52.9	65.4
BERT-ADA	81.2	86.3	83.7	83.2	60.4	70.0
BERT-LIW	81.9	86.6	84.1	86.7	56.0	68.1
BERT-DAFT	79.1	85.9	82.3	83.9	58.6	69.0
BERT-DAFT-SYN	76.9	80.7	78.7	70.7	56.8	63.0

Table 4: Model performance on domain transfer experiments from news to clinical notes.

Model	In-Domain			Out-of-Domain		
	P	R	F1	P	R	F1
VERB	58.8	66.5	62.5	44.6	68.1	53.9
DELEX	75.0	66.3	70.4	56.9	64.5	60.4
BERT	80.6	86.0	83.2	75.0	63.6	68.9
CBERT	79.2	83.3	81.2	66.5	65.1	65.8
BERT-ADA	81.1	85.9	83.4	74.5	62.2	67.8
BERT-LIW	80.0	87.0	83.4	72.8	67.3	70.0
BERT-DAFT	78.5	84.8	81.5	72.7	73.1	72.9
BERT-DAFT-SYN	80.0	78.7	79.3	67.6	60.7	63.9

Table 5: Model performance on domain transfer experiments from news to doctor-patient conversations.

(Sims et al., 2019).

CBERT: Similar to BERT, but embeddings are extracted from Clinical-BERT (Alsentzer et al., 2019)

BERT-ADA: BERT trained using adversarial domain adaptation.

BERT-LIW: BERT trained on data weighted by LM likelihood. We train autoregressive language models over 3 million tokens for each domain.

BERT-DAFT: BERT with domain adaptive fine-tuning. For target domains, we use the same text as BERT-LIW, and extract 3 million tokens from CNN/ DailyMail (Hermann et al., 2015) for news.

BERT-DAFT-SYN: BERT with syntactic fine-tuning on the same text as BERT-DAFT, tagged using Stanford CoreNLP (Manning et al., 2014). Complete implementation details are provided in appendix A.

5.2 Results

Tables 4 and 5 show the performance of all models when transferring from news data to clinical notes and doctor-patient conversations respectively. From the tables, we see that DELEX is surpris-

ingly strong out-of-domain. BERT with no transfer performs well out-of-domain, improving by 8.25 F1 points on average over DELEX. C-BERT also performs well out-of-domain, but does worse than BERT. We attribute this to the fact that fine-tuning only on clinical notes does not improve alignment with the source domain (news), providing no basis for models trained on news to adapt better. BERT-ADA shows mixed results, improving over BERT by 2.4 F1 on notes, but dropping by 1.1 F1 on conversations. BERT-LIW and BERT-DAFT improve upon BERT in both domains. BERT-DAFT shows minor performance drops in-domain, due to some degree of catastrophic forgetting. BERT-DAFT-SYN shows performance drops, both in-domain and out-of-domain, in both settings. Unlike syntactic relexicalization work which used non-contextualized embeddings, we use contextualized embeddings, which possess a larger degree of syntactic information, probably reducing the need for syntax-driven training. Another source of errors is POS tagging, since off-the-shelf taggers trained

on news will be less accurate on our data. Across domains, the skew between precision and recall is higher on notes, which might stem from the specialized vocabulary dragging down recall.

6 Analysis and Discussion

Tables 4 and 5 provide an indication of model ability to handle covariate shift. However, covariate shift occurs at multiple layers in language (e.g., lexical level, syntactic level, etc.), leading to different dimensions of variation between domains (e.g., topical variation, genre variation, etc.). Looking at overall model performance does not offer insight into whether there are specific shifts that some models are better at addressing. We dig deeper into this question, focusing on two levels of shift: (i) lexical shift, and (ii) semantic (event type) shift.

Variation under lexical shift: We separate model performance on in-vocabulary (IV) and out-of-vocabulary (OOV) tokens. Note that the proportion of events that are OOV is higher in clinical notes (52%) than conversations (20.6%). Tables 6 and 7 present model performance on these token categories. Surprisingly, despite the use of specialized language, OOV performance on clinical notes is higher than conversations for all models except BERT-DAFT. Taking a closer look at the OOV event instances from clinical notes that models identify correctly, we see that a large proportion (54.8%) contain one of three morphological patterns: (i) past tense verbs ending in “-ed”, (ii) gerunds ending in “-ing”, or (iii) nouns ending in “-tion” or “-sion”. These patterns are also common among events in the news domain. For example, past tense verbs often refer to events that have already occurred and gerunds and nouns ending in “-tion” refer to processes. We hypothesize that BERT-based models might be exploiting these morphological regularities to correctly label unseen medical terms (e.g., irrigated, excision, dissected, wheezing, etc.). These patterns are more prevalent in notes (35.6%) than conversations (23.5%), explaining the surprising performance difference.

Variation under semantic shift: To determine whether model performance on OOV tokens depends on event type, we randomly sample ~ 500 OOV tokens from each domain and label them for event type. We use the same typology as TimeML (State, I-State, Occurrence, Aspectual, Reporting, Perception, I-Action, None), with additional labels for the event types we introduce (ActivityPattern,

Model	IV F1	OOV F1
BERT	73.5	61.2
BERT-ADA	75.2	65.0
BERT-LIW	73.6	62.6
BERT-DAFT	75.7	62.0
BERT-DAFT-SYN	67.7	58.4

Table 6: Model performance on in-vocabulary (IV) and out-of-vocabulary (OOV) terms from notes.

Model	IV F1	OOV F1
BERT	71.3	57.9
BERT-ADA	70.2	57.6
BERT-LIW	72.0	61.4
BERT-DAFT	74.9	63.6
BERT-DAFT-SYN	65.5	55.5

Table 7: Model performance on in-vocabulary (IV) and out-of-vocabulary (OOV) terms from conversations.

LongTermState).⁹ We run an ANOVA model with each token per model as an instance (total 5080 instances), noting Event Type, Target (notes/convo), Model (BERT/ADA/LIW/DAFT/DAFT-SYN) and Correctness (1 vs 0). Correctness is the dependent variable, while others are independent variables. We include all pairwise interaction terms and the three way interaction between Event Type, Target and Model. We see a positive main effect of Event Type on Correctness ($p < 0.0001$), indicating that some event types are more difficult. There are two significant two-way interactions, one between Target and Event type ($p < 0.0001$), indicating that difficulty of event types differs across sources, and between Model and Event type ($p < 0.0001$), indicating that which model is better depends on event type. Three way interaction between Model, Event type, and Target is also significant ($p < 0.0001$), indicating that performance differences between models per event type differs between sources.

We interpret differences in performance per event type separately for each source using a Student-t post-hoc analysis to determine which pairwise contrasts are statistically significant. This reveals that in clinical notes, LIW outperforms all models on I-State events (i.e., hypothetical, future or negated states) and LongTermState events, a category never seen in the training data. These improvements might stem from the training algo-

⁹Examples provided in appendix B

rithm used by LIW. LIW up-weights instances in news that resemble medical data, which contains a high proportion of these event categories. Therefore, despite being infrequent in news, they get up-weighted, helping LIW identify them better.

7 Conclusion

In this work, we focused on unsupervised adaptation of event extractors to new domains by aligning the marginal distributions of source and target domains. We created two event extraction test sets using English texts from two medical domains: (i) clinical notes, and (ii) doctor-patient conversations, and tested the efficacy of three alignment techniques: (i) adversarial domain adaptation (ADA), (ii) domain adaptive fine-tuning (DAFT), and (iii) a new instance weighting technique based on language model likelihood scores (LIW). None of these models consistently outperformed the others, but a deeper analysis of model performance under different types of shifts (e.g., lexical shift, semantic shift) uncovered interesting variations among models. Our best-performing models attained F1 scores of 70.0 and 72.9 on notes and conversations respectively, using *no* labeled target data. We believe these models define a good starting point and can be further improved using few-shot learning.

Acknowledgements

This work was supported by the University of Pittsburgh Medical Center (UPMC) and Abridge AI Inc through the Center for Machine Learning and Health at Carnegie Mellon University. The authors would like to thank the anonymous reviewers for their helpful feedback on this work.

References

Emily Alsentzer, John Murphy, William Boag, Wei-Hung Weng, Di Jindi, Tristan Naumann, and Matthew McDermott. 2019. [Publicly available clinical BERT embeddings](#). In *Proceedings of the 2nd Clinical Natural Language Processing Workshop*, pages 72–78, Minneapolis, Minnesota, USA. Association for Computational Linguistics.

Jun Araki and Teruko Mitamura. 2018. [Open-domain event detection using distant supervision](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 878–891, Santa Fe, New Mexico, USA. Association for Computational Linguistics.

Amittai Axelrod, Xiaodong He, and Jianfeng Gao. 2011. [Domain adaptation via pseudo in-domain data](#)

[selection](#). In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 355–362, Edinburgh, Scotland, UK. Association for Computational Linguistics.

Emmon Bach. 1986. The algebra of events. *Linguistics and philosophy*, pages 5–16.

Jonathan Berant, Vivek Srikumar, Pei-Chun Chen, Abby Vander Linden, Brittany Harding, Brad Huang, Peter Clark, and Christopher D. Manning. 2014. [Modeling biological processes for reading comprehension](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1499–1510, Doha, Qatar. Association for Computational Linguistics.

Steven Bethard, Leon Derczynski, Guergana Savova, James Pustejovsky, and Marc Verhagen. 2015. [SemEval-2015 task 6: Clinical TempEval](#). In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pages 806–814, Denver, Colorado. Association for Computational Linguistics.

Steven Bethard, Guergana Savova, Wei-Te Chen, Leon Derczynski, James Pustejovsky, and Marc Verhagen. 2016. [SemEval-2016 task 12: Clinical TempEval](#). In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 1052–1062, San Diego, California. Association for Computational Linguistics.

Steven Bethard, Guergana Savova, Martha Palmer, and James Pustejovsky. 2017. [SemEval-2017 task 12: Clinical TempEval](#). In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 565–572, Vancouver, Canada. Association for Computational Linguistics.

John Blitzer, Mark Dredze, and Fernando Pereira. 2007. [Biographies, Bollywood, boom-boxes and blenders: Domain adaptation for sentiment classification](#). In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*, pages 440–447, Prague, Czech Republic. Association for Computational Linguistics.

John Blitzer, Ryan McDonald, and Fernando Pereira. 2006. [Domain adaptation with structural correspondence learning](#). In *Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing*, pages 120–128, Sydney, Australia. Association for Computational Linguistics.

Nathanael Chambers, Taylor Cassidy, Bill McDowell, and Steven Bethard. 2014. [Dense event ordering with a multi-pass architecture](#). *Transactions of the Association for Computational Linguistics*, 2:273–284.

Minmin Chen, Kilian Q. Weinberger, Fei Sha, and Yoshua Bengio. 2014. [Marginalized denoising auto-encoders for nonlinear representations](#). In *Proceedings of the 31th International Conference on Machine Learning, ICML 2014, Beijing, China, 21-26*

- June 2014, volume 32 of *JMLR Workshop and Conference Proceedings*, pages 1476–1484. JMLR.org.
- Anne Cocos, Alexander G. Fiks, and Aaron J. Masino. 2017. [Deep learning for pharmacovigilance: recurrent neural network architectures for labeling adverse drug reactions in twitter posts](#). *J. Am. Medical Informatics Assoc.*, 24(4):813–821.
- Agata Cybulska and Piek Vossen. 2014. [Using a sledgehammer to crack a nut? lexical diversity and event coreference resolution](#). In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 4545–4552, Reykjavik, Iceland. European Language Resources Association (ELRA).
- 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.
- George Doddington, Alexis Mitchell, Mark Przybocki, Lance Ramshaw, Stephanie Strassel, and Ralph Weischedel. 2004. [The automatic content extraction \(ACE\) program – tasks, data, and evaluation](#). In *Proceedings of the Fourth International Conference on Language Resources and Evaluation (LREC'04)*, Lisbon, Portugal. European Language Resources Association (ELRA).
- Long Duong, Trevor Cohn, Steven Bird, and Paul Cook. 2015. [Cross-lingual transfer for unsupervised dependency parsing without parallel data](#). In *Proceedings of the Nineteenth Conference on Computational Natural Language Learning*, pages 113–122, Beijing, China. Association for Computational Linguistics.
- George Foster, Cyril Goutte, and Roland Kuhn. 2010. [Discriminative instance weighting for domain adaptation in statistical machine translation](#). In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, pages 451–459, Cambridge, MA. Association for Computational Linguistics.
- Yaroslav Ganin and Victor S. Lempitsky. 2015. [Unsupervised domain adaptation by backpropagation](#). In *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, volume 37 of *JMLR Workshop and Conference Proceedings*, pages 1180–1189. JMLR.org.
- Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor S. Lempitsky. 2016. [Domain-adversarial training of neural networks](#). *J. Mach. Learn. Res.*, 17:59:1–59:35.
- Xavier Glorot, Antoine Bordes, and Yoshua Bengio. 2011. [Domain adaptation for large-scale sentiment classification: A deep learning approach](#). In *Proceedings of the 28th International Conference on Machine Learning, ICML 2011, Bellevue, Washington, USA, June 28 - July 2, 2011*, pages 513–520. Omnipress.
- Ralph Grishman and Beth Sundheim. 1996. [Message understanding conference- 6: A brief history](#). In *COLING 1996 Volume 1: The 16th International Conference on Computational Linguistics*.
- Tao Gui, Qi Zhang, Haoran Huang, Minlong Peng, and Xuanjing Huang. 2017. [Part-of-speech tagging for twitter with adversarial neural networks](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2411–2420, Copenhagen, Denmark. Association for Computational Linguistics.
- Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. 2020. [Don't stop pretraining: Adapt language models to domains and tasks](#).
- Xiaochuang Han and Jacob Eisenstein. 2019. [Unsupervised domain adaptation of contextualized embeddings for sequence labeling](#). 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 4238–4248, Hong Kong, China. Association for Computational Linguistics.
- Sam Henry, Kevin Buchan, Michele Filannino, Amber Stubbs, and Özlem Uzuner. 2020. [2018 n2c2 shared task on adverse drug events and medication extraction in electronic health records](#). *J. Am. Medical Informatics Assoc.*, 27(1):3–12.
- Karl Moritz Hermann, Tomáš Kociský, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. [Teaching machines to read and comprehend](#). In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 1693–1701.
- Ajit Jain, Girish Kasiviswanathan, and Ruihong Huang. 2016. [Towards accurate event detection in social media: A weakly supervised approach for learning implicit event indicators](#). In *Proceedings of the 2nd Workshop on Noisy User-generated Text (WNUT)*, pages 70–77, Osaka, Japan. The COLING 2016 Organizing Committee.
- Jing Jiang and ChengXiang Zhai. 2007. [Instance weighting for domain adaptation in NLP](#). In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*, pages 264–271, Prague, Czech Republic. Association for Computational Linguistics.

- Katherine Keith, Abram Handler, Michael Pinkham, Cara Magliozzi, Joshua McDuffie, and Brendan O'Connor. 2017. [Identifying civilians killed by police with distantly supervised entity-event extraction](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1547–1557, Copenhagen, Denmark. Association for Computational Linguistics.
- Jin-Dong Kim, Tomoko Ohta, Sampo Pyysalo, Yoshinobu Kano, and Jun'ichi Tsujii. 2009. [Overview of BioNLP'09 shared task on event extraction](#). In *Proceedings of the BioNLP 2009 Workshop Companion Volume for Shared Task*, pages 1–9, Boulder, Colorado. Association for Computational Linguistics.
- Jin-Dong Kim, Tomoko Ohta, and Jun'ichi Tsujii. 2008. [Corpus annotation for mining biomedical events from literature](#). *BMC Bioinform.*, 9.
- Heeyoung Lee, Marta Recasens, Angel Chang, Mihai Surdeanu, and Dan Jurafsky. 2012. [Joint entity and event coreference resolution across documents](#). In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 489–500, Jeju Island, Korea. Association for Computational Linguistics.
- Jiwei Li, Alan Ritter, Claire Cardie, and Eduard Hovy. 2014. [Major life event extraction from twitter based on congratulations/condolences speech acts](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1997–2007, Doha, Qatar. Association for Computational Linguistics.
- Linguistic Data Consortium. 2005. [Ace \(automatic content extraction\) english annotation guidelines for events](#).
- Christopher Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven Bethard, and David McClosky. 2014. [The Stanford CoreNLP natural language processing toolkit](#). In *Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 55–60, Baltimore, Maryland. Association for Computational Linguistics.
- John McCarthy. 2002. [Actions and other events in situation calculus](#). In *Proceedings of the Eighth International Conference on Principles and Knowledge Representation and Reasoning (KR-02), Toulouse, France, April 22-25, 2002*, pages 615–628. Morgan Kaufmann.
- Teruko Mitamura, Zhengzhong Liu, and Eduard H. Hovy. 2016. [Overview of TAC-KBP 2016 event nugget track](#). In *Proceedings of the 2016 Text Analysis Conference, TAC 2016, Gaithersburg, Maryland, USA, November 14-15, 2016*. NIST.
- Teruko Mitamura, Yukari Yamakawa, Susan Holm, Zhiyi Song, Ann Bies, Seth Kulick, and Stephanie Strassel. 2015. [Event nugget annotation: Processes and issues](#). In *Proceedings of the The 3rd Workshop on EVENTS: Definition, Detection, Coreference, and Representation*, pages 66–76, Denver, Colorado. Association for Computational Linguistics.
- Aakanksha Naik, Chris Bogart, and Carolyn Rose. 2017. [Extracting personal medical events for user timeline construction using minimal supervision](#). In *BioNLP 2017*, pages 356–364, Vancouver, Canada. Association for Computational Linguistics.
- Aakanksha Naik and Carolyn Rosé. 2020. [Towards open domain event trigger identification using adversarial domain adaptation](#).
- Azadeh Nikfarjam, Abeed Sarker, Karen O'Connor, Rachel E. Ginn, and Graciela Gonzalez-Hernandez. 2015. [Pharmacovigilance from social media: mining adverse drug reaction mentions using sequence labeling with word embedding cluster features](#). *J. Am. Medical Informatics Assoc.*, 22(3):671–681.
- Tim O'Gorman, Kristin Wright-Bettner, and Martha Palmer. 2016. [Richer event description: Integrating event coreference with temporal, causal and bridging annotation](#). In *Proceedings of the 2nd Workshop on Computing News Storylines (CNS 2016)*, pages 47–56, Austin, Texas. Association for Computational Linguistics.
- Boyan Onyshkevych, Mary Ellen Okurowski, and Lynn Carlson. 1993. [Tasks, domains, and languages for information extraction](#). In *TIPSTER TEXT PROGRAM: PHASE I: Proceedings of a Workshop held at Fredericksburg, Virginia, September 19-23, 1993*, pages 123–133, Fredericksburg, Virginia, USA. Association for Computational Linguistics.
- James Pustejovsky, José M. Castaño, Robert Ingria, Roser Saurí, Robert J. Gaizauskas, Andrea Setzer, Graham Katz, and Dragomir R. Radev. 2003a. [Timeml: Robust specification of event and temporal expressions in text](#). In *New Directions in Question Answering, Papers from 2003 AAAI Spring Symposium, Stanford University, Stanford, CA, USA*, pages 28–34. AAAI Press.
- James Pustejovsky, Patrick Hanks, Roser Sauri, Andrew See, Robert Gaizauskas, Andrea Setzer, Dragomir Radev, Beth Sundheim, David Day, Lisa Ferro, et al. 2003b. [The timebank corpus](#). In *Corpus linguistics*, volume 2003, page 40. Lancaster, UK.
- Alan Ritter, Mausam, Oren Etzioni, and Sam Clark. 2012. [Open domain event extraction from twitter](#). In *The 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '12, Beijing, China, August 12-16, 2012*, pages 1104–1112. ACM.
- Abeed Sarker and Graciela Gonzalez. 2015. [Portable automatic text classification for adverse drug reaction detection via multi-corpus training](#). *J. Biomed. Informatics*, 53:196–207.

- Roser Saurí, Robert Knippen, Marc Verhagen, and James Pustejovsky. 2005. *Evita: A robust event recognizer for QA systems*. In *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing*, pages 700–707, Vancouver, British Columbia, Canada. Association for Computational Linguistics.
- Matthew Sims, Jong Ho Park, and David Bamman. 2019. *Literary event detection*. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3623–3634, Florence, Italy. Association for Computational Linguistics.
- Pontus Stenetorp, Sampo Pyysalo, Goran Topić, Tomoko Ohta, Sophia Ananiadou, and Jun’ichi Tsujii. 2012. *brat: a web-based tool for NLP-assisted text annotation*. In *Proceedings of the Demonstration of the 13th Conference of the European Chapter of the Association for Computational Linguistics*, pages 102–107, Avignon, France. Association for Computational Linguistics.
- William F. Styler IV, Steven Bethard, Sean Finan, Martha Palmer, Sameer Pradhan, Piet C de Groen, Brad Erickson, Timothy Miller, Chen Lin, Guergana Savova, and James Pustejovsky. 2014. *Temporal annotation in the clinical domain*. *Transactions of the Association for Computational Linguistics*, 2:143–154.
- Rui Wang, Masao Utiyama, Lemao Liu, Kehai Chen, and Eiichiro Sumita. 2017. *Instance weighting for neural machine translation domain adaptation*. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1482–1488, Copenhagen, Denmark. Association for Computational Linguistics.
- Tuangthong Wattarujekrit, Parantu K. Shah, and Nigel Collier. 2004. *Pasbio: predicate-argument structures for event extraction in molecular biology*. *BMC Bioinform.*, 5:155.
- Miaomiao Wen, Zeyu Zheng, Hyeju Jang, Guang Xiang, and Carolyn Penstein Rosé. 2013. *Extracting events with informal temporal references in personal histories in online communities*. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 836–842, Sofia, Bulgaria. Association for Computational Linguistics.
- Yftah Ziser and Roi Reichart. 2017. *Neural structural correspondence learning for domain adaptation*. In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*, pages 400–410, Vancouver, Canada. Association for Computational Linguistics.

Appendix

A Implementation Details

BERT: The BERT baseline model uses the uncased variant of BERT-Base (with no additional

fine-tuning) for feature extraction. We generate token representations by running BERT-Base and concatenating the outputs of the model’s last 4 hidden layers. The BiLSTM layer has a hidden size of 100, with an input dropout of 0.5. The MLP layer is 100-dimensional. These values are consistent with the setup in [Naik and Rosé \(2020\)](#).

BERT-ADA: The domain predictor (adversary) is a 3-layer MLP with each layer having a dimensionality of 100 and ReLU activations between layers. For the hyperparameter λ , which is the constant used to weight domain prediction loss, we experiment with values from [0.5, 1.0, 2.0, 5.0], and choose the best model based on F1 scores on the source domain validation set. We run one search trial with a fixed random seed (0) for all settings. The best performing model on clinical notes uses $\lambda = 1.0$ and on conversations uses $\lambda = 0.5$.

BERT-LIW: The autoregressive word-level language models used for weighting are 3-layer LSTMs, with a hidden size of 300 and layer dropout of 0.2 at each layer. Input embeddings are initialized using 300-dimensional GloVe embeddings, with parameter tying between input and output embedding matrices. The models are trained using SGD with gradient clipping at 0.25 and a batch size of 16 for 25 epochs. Training starts with a learning rate of 20, which is divided by 4 whenever validation loss plateaus.

BERT-DAFT/BERT-DAFT-SYN: BERT-Base is fine-tuned for 3 epochs, using a batch size of 4 and default parameter settings in the Huggingface transformers library.

All event extraction models are trained with a batch size of 16 and use Adam optimizer with a learning rate of 0.001. Models are trained for 1000 epochs with early stopping. All experiments are run on an NVIDIA GeForce GTX 1080 Ti machine.

B Event Typology Examples

1. **Occurrence:** Occurrence refers to all events describing something that happens or occurs in the world. This is the broadest class of events. For example, “I *took* Midol yesterday.”
2. **Aspectual:** Aspectual events refer to events which focus on various aspects of a different event’s history, such as initiation, termination, continuation etc. For example, “I *started* taking this medicine last Friday.” Here *started* is an aspectual event describing the initiation

of the event *taking*.

3. **Reporting:** Reporting events describe the action of an entity (person/group/organization) declaring something, narrating an event, providing information about an event etc. For example, "So you *said* you have been experiencing symptoms since yesterday?"
4. **Perception:** Perception events refer to events involving the physical perception of a different event. For example, "I *watched* my weight gain throughout the pregnancy."
5. **State:** States describe circumstances in which something obtains or holds true. For example, "My blood pressure is *higher* today". Note that annotation of state events in TimeML is subject to certain rules.
6. **Intensional Action (I-Action):** Intensional actions introduce an explicit event argument describing an action or situation, from which we can infer something given its relation with the intensional action. For example, "We will *investigate* your symptoms further via this test." Here *investigate* is an intensional action associated with the *symptoms* event.
7. **Intensional State (I-State):** Intensional states contain stative events that refer to alternative or possible worlds. For example, "You might observe *higher* blood pressure for a few days when you start taking this medicine."
8. **Activity Pattern:** Activity patterns refer to events that are not clearly grounded to a single occurrence in time, but still considered events since the presence of a participant stops them from being purely generic. For example, "You should *take* your blood pressure regularly."
9. **Long Term State:** Long-term states expand the annotation of states beyond TimeML restrictions, allowing the inclusion of long-term chronic conditions. For example, "You have a history of *COPD*."