

Weakly Supervised Medication Regimen Extraction from Medical Conversations

Dhruvesh Patel*

College of Information and Computer Sciences
University of Massachusetts Amherst
dhruveshpate@cs.umass.edu

Sandeep Konam

Abridge AI Inc.
san@abridge.com

Sai P. Selvaraj

Abridge AI Inc.
prabhakarsai@abridge.com

Abstract

Automated Medication Regimen (MR) extraction from medical conversations can not only improve recall and help patients follow through with their care plan, but also reduce the documentation burden for doctors. In this paper, we focus on extracting spans for *frequency*, *route* and *change*, corresponding to medications discussed in the conversation. We first describe a unique dataset of annotated doctor-patient conversations and then present a weakly supervised model architecture that can perform span extraction using noisy classification data. The model utilizes an attention bottleneck inside a classification model to perform the extraction. We experiment with several variants of attention scoring and projection functions and propose a novel transformer-based attention scoring function (TAScore). The proposed combination of TAScore and Fusedmax projection achieves a 10 point increase in Longest Common Substring F1 compared to the baseline of additive scoring plus softmax projection.

1 Introduction

Patients forget 40-80% of the medical information provided by healthcare practitioners immediately (Mcguire, 1996) and misconstrue 48% of what they think they remembered (Anderson et al., 1979), and this adversely affects patient adherence. Automatically extracting information from doctor-patient conversations can help patients correctly recall doctor’s instructions and improve compliance with the care plan (Tsulukidze et al., 2014). On the other hand, clinicians spend up to 49.2% of their overall time on EHR and desk work, and only 27.0% of their total time on direct clinical face time with

*Work done as an intern at Abridge AI Inc.

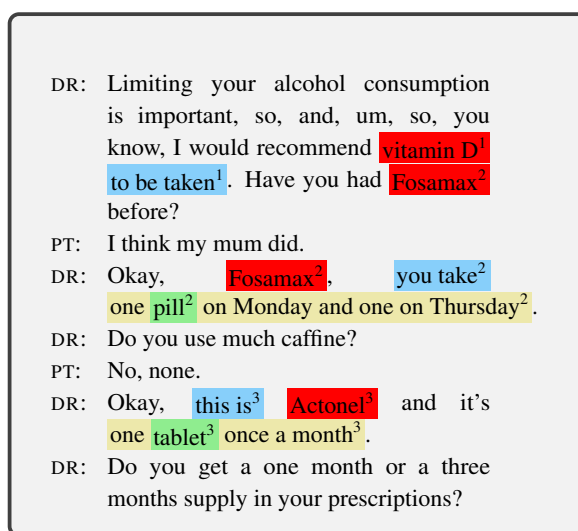


Figure 1: An example excerpt from a doctor-patient conversation transcript. Here, there are three medications mentioned indicated by the superscript. The extracted attributes, change, route and frequency, for each medications are also shown.

patients (Sinsky et al., 2016). Increased data management work is also correlated with increased doctor burnout (Kumar, 2016). Information extracted from medical conversations can also aid doctors in their documentation work (Rajkomar et al., 2019; Schloss and Konam, 2020), allow them to spend more face time with the patients, and build better relationships.

In this work, we focus on extracting Medication Regimen (MR) information (Du et al., 2019; Selvaraj and Konam, 2019) from the doctor-patient conversations. Specifically, we extract three attributes, i.e., *frequency*, *route* and *change*, corresponding to medications discussed in the conversation (Figure 1). Medication Regimen information can help doctors with medication orders cum renewals, medication reconciliation, verification of

reconciliations for errors, and other medication-centered EHR documentation tasks. It can also improve patient engagement, transparency and better compliance with the care plan (Tsulukidze et al., 2014; Grande et al., 2017).

MR attribute information present in a conversation can be obtained as spans in text (Figure 1) or can be categorized into classification labels (Table 2). While the classification labels are easy to obtain at scale in an automated manner – for instance, by pairing conversations with billing codes or medication orders – they can be noisy and can result in a prohibitively large number of classes. Classification labels go through normalization and disambiguation, often resulting in label names which are very different from the phrases used in the conversation. This process leads to a loss of granular information present in the text (see, for example, row 2 in Table 2). Span extraction, on the other hand, alleviates this issue as the outputs are actual spans in the conversation. However, span extraction annotations are relatively hard to come by and are time-consuming to annotate manually. Hence, in this work, we look at the task of MR attribute span extraction from doctor-patient conversation using weak supervision provided by the noisy classification labels.

The main contributions of this work are as follows. We present a way of setting up an MR attribute extraction task from noisy classification data (Section 2). We propose a weakly supervised model architecture which utilizes attention bottleneck inside a classification model to perform span extraction (Section 3 & 4). In order to favor sparse and contiguous extractions, we experiment with two variants of attention projection functions (Section 3.1.2), namely, softmax and Fusedmax (Niculae and Blondel, 2017). Further, we propose a novel transformer-based attention scoring function TAScore (Section 3.1.1). The combination of TAScore and Fusedmax achieves significant improvements in extraction performance over a phrase-based (22 LCSF1 points) and additive softmax attention (10 LCSF1 points) baselines.

2 Medication Regimen (MR) using Weak Supervision

Medication Regimen (MR) consists of information about a prescribed medication akin to attributes of an entity. In this work, we specifically focus on *frequency*, *route* of the medication and any *change* in

Attribute	Normalized Classes
<i>frequency</i>	Daily Every morning At Bedtime Twice a day Three times a day Every six hours Every week Twice a week Three times a week Every month Other None
<i>route</i>	Pill Injection Topical cream Nasal spray Medicated patch Ophthalmic solution Inhaler Oral solution Other None
<i>change</i>	Take Stop Increase Decrease None Other

Table 1: The normalized labels in the classification data.

the medication’s *dosage* or *frequency* as shown in Figure 1. For example, given the conversation excerpt and the medication “Fosamax” as shown in Figure 1, the model needs to extract the spans “one pill on Monday and one on Thursday”, “pill” and “you take” for attributes *frequency*, *route* and *change*, respectively. The major challenge, however, is to perform the attribute span extraction using noisy classification labels with very few or no span-level labels. The rest of this section describes the dataset used for this task.

2.1 Data

The data used in this paper comes from a collection of human transcriptions of 63000 fully-consented and de-identified doctor-patient conversations. A total of 57000 conversations were randomly selected to construct the training (and dev) conversation pool and the remaining 6000 conversations were reserved as the test pool.

The classification dataset: All the conversations are annotated with MR tags by expert human annotators. Each set of MR tags consists of the *medication* name and its corresponding attributes *frequency*, *route* and *change*, which are normalized free-form instructions in natural language phrases corresponding to each of the three attributes (see Table 8 in A.4). Each set of MR tags is grounded to a contiguous window of utterances’ *text*,¹ around a medication mention as evidence for that set. Hence, each set of grounded MR tags can be written as $\langle medication, text, frequency, route, change \rangle$, where the last three entries correspond to the three MR attributes.

The free-form instructions for each attribute in the MR tags are normalized and categorized into manageable number of classification labels to avoid long tail and overlapping classes. This process re-

¹The *text* includes both the spoken words and the speaker information.

<i>text</i>	<i>medication</i>	Classification labels		
		<i>frequency</i>	<i>route</i>	<i>change</i>
... I would recommend vitamin D to be taken. Have you had Fosamax before?...	vitamin D	none	none	take
... I think my mum did. Okay, Fosamax, you take one pill on Monday and one on Thursday. Do you have much caffeine? No, none...	Fosamax	Twice a week	pill	take
Do you have much caffeine? No, none. Okay, this is Actonel and it's, one tablet once a month...	Actonel	Once a month	pill	take

Table 2: Classification examples resulting from the conversation shown in Figure 1.

sults in classes shown in Table 1.² As an illustration, this annotation process when applied to the conversation piece shown in Figure 1 would result in the three data points shown in Table 2. Using this procedure on both the training and test conversation pools, we obtain 45,059 training, 11,212 validation and 5,458 test classification data points.³

The extraction dataset: Since the goal is to extract spans related to MR attributes, we would ideally need a dataset with span annotations to perform this task in a fully supervised manner. However, span annotation is laborious and expensive. Hence, we re-purpose the classification dataset (along with its classification labels) to perform the task of span extraction using weak supervision. We also manually annotate a small fraction of the train, validation and test sets (150, 150 and 500 data-points respectively) for attribute spans to see the effect of supplying a small number of strongly supervised instances on the performance of the model. In order to have a good representation of all the classes in the test set, we increase the sampling weight of data-points which have rare classes. Hence, our test set is relatively more difficult compared to a random sample of 500 data-points. All the results are reported on our test set of 500 difficult data-points annotated for attribute spans.

For annotating attribute spans, the annotators were given instructions to mark spans which provide minimally sufficient and natural evidence for the already annotated attribute class as described below.

Sufficiency: Given only the annotated span for a particular attribute, one should be able to predict the correct classification label. This aims to encourage the attribute spans to cover all distinguishing information for that attribute.

²The detailed explanation for each of the classes can be found in Table 7 in Appendix A.1.

³The dataset statistics are given in Appendix A.1.

Minimality: Peripheral words which can be replaced with other words without changing the attribute’s classification label should not be included in the extracted span. This aims to discourage marking entire utterances as attribute spans.

Naturalness: The marked span(s) if presented to a human should sound like complete English phrases (if it has multiple tokens) or a meaningful word if it has only a single token. In essence, this means that the extractions should not drop stop words from within phrases. This requirement aims to reduce the cognitive load on the human who uses the model’s extraction output.

2.2 Challenges

Using medical conversations for information extraction is more challenging compared to written doctor notes because the spontaneity of conversation gives rise to a variety of speech patterns with disfluencies and interruptions. Moreover, the vocabulary can range from colloquial to medical jargon.

In addition, we also have noise in our classification dataset with its main source being annotators’ use of information outside the grounded *text* window to produce the free-form tags. This happens in two ways. First, when the free-form MR instructions are written using evidence that was discussed elsewhere in the conversation but is not present in the grounded *text* window. Second, when the annotator uses their domain knowledge instead of using just the information in the grounded *text* window – for instance, when the *route* of a *medication* is not explicitly mentioned, the annotator might use the *medication*’s common *route* in their free-form instructions. Using manual analysis of the 800 data-points across the train, dev and test sets, we find that 22% of *frequency*, 36% of *route* and 15% of *change* classification labels, have this noise.

In this work, our approach to extraction depends

on the size of the auxiliary task’s (classification) dataset to overcome above mentioned challenges.

3 Background

There have been several successful attempts to use neural attention (Bahdanau et al., 2015) to extract information from text in an unsupervised manner (He et al., 2017; Lin et al., 2016; Yu et al., 2019). Attention scores provide a good proxy for importance of a particular token in a model. However, when there are multiple layers of attention, or if the encoder is too complex and trainable, the model no longer provides a way to produce reliable and faithful importance scores (Jain and Wallace, 2019).

We argue that, in order to bring in the faithfulness, we need to create an attention bottleneck in our classification + extraction model. The attention bottleneck is achieved by employing an attention function which generates a set of attention weights over the encoded input tokens. Attention bottleneck forces the classifier *to only see* the portions of input that pass through it, thereby enabling us to trade the classification performance for extraction performance and getting span extraction with weak supervision from classification labels.

In the rest of this section, we provide general background on neural attention and present its variants employed in this work. This is followed by the presentation of our complete model architecture in the subsequent sections.

3.1 Neural Attention

Given a query $\mathbf{q} \in \mathbb{R}^m$ and keys $\mathbf{K} \in \mathbb{R}^{l \times n}$, the attention function $\alpha: \mathbb{R}^m \times \mathbb{R}^{l \times n} \rightarrow \Delta^l$ is composed of two functions: a scoring function $\mathcal{S}: \mathbb{R}^m \times \mathbb{R}^{l \times n} \rightarrow \mathbb{R}^l$ which produces unnormalized importance scores, and a projection function $\Pi: \mathbb{R}^l \rightarrow \Delta^l$ which normalizes these scores by projecting them to an $(l - 1)$ -dimensional probability simplex.⁴

3.1.1 Scoring Function

The purpose of the scoring function is to produce importance scores for each entry in the key \mathbf{K} w.r.t the query \mathbf{q} for the task at hand, which in our case is classification. We experiment with two scoring functions: additive and transformer-based.

Additive: This is same as the scoring function used in Bahdanau et al. (2015), where the scores

⁴Throughout this work l represents the sequence length dimension and $\Delta^l = \{\mathbf{x} \in \mathbb{R}^l \mid \mathbf{x} > 0, \|\mathbf{x}\|_1 = 1\}$ represents a probability simplex.




Figure 2: Architecture of TAScore. \mathbf{q} and \mathbf{K} are input query and keys, respectively, and \mathbf{s} are the output scores.

are produced as follows:

$$s_j = \mathbf{v}^T \tanh(\mathbf{W}_q \mathbf{q} + \mathbf{W}_k \mathbf{k}_j),$$

where, $\mathbf{v} \in \mathbb{R}^m$, $\mathbf{W}_q \in \mathbb{R}^{m \times m}$ and $\mathbf{W}_k \in \mathbb{R}^{m \times n}$ are trainable weights.

Transformer-based Attention Score (TAScore):

While the additive scoring function is simple and easy to train, it suffers from one major drawback in our setting: since we freeze the weights of our embedder and do not use multiple layers of trainable attention (Section 4.4), the additive attention can struggle to resolve references – finding the correct attribute when there are multiple entities of interest, especially when there are multiple distinct medications (Section 6.4). For this reason, we propose a novel multi-layer transformer-based attention scoring function (TAScore) which can perform this reference resolution while also preserving the *attention bottleneck*. Figure 2 shows the architecture of TAScore. The query and key vectors are projected to the same space using two separate linear layers while also adding sinusoidal positional embeddings to the key vectors. A special trainable separator vector is added between the query and key vectors and the entire sequence is passed through a multi-layer transformer (Vaswani et al., 2017). Finally, scalar scores (one corresponding to each vector in the key) are produced from the outputs of the transformer by passing them through a feed-forward layer with dropout.

3.1.2 Projection Function

A projection function $\Pi: \mathbb{R}^l \rightarrow \Delta^l$ in the context of attention distribution, normalizes the real valued importance scores by projecting them to an $(l -$

1)-dimensional probability simplex Δ^l . Niculae and Blondel (2017) provide a unified view of the projection function as follows:

$$\Pi_{\Omega}(\mathbf{s}) = \arg \max_{\mathbf{a} \in \Delta^l} \mathbf{a}^T \mathbf{s} - \gamma \Omega(\mathbf{a}).$$

Here, $\mathbf{a} \in \Delta^l$, γ is a hyperparameter and Ω is a regularization penalty which allows us to introduce problem specific inductive bias into our attention distribution. When Ω is strongly convex, we have a closed form solution to the projection operation as well as its gradient (Niculae and Blondel, 2017; Blondel et al., 2020). Since we use the attention distribution to perform extraction, we experiment with the following instances of projection functions in this work.

Softmax: $\Omega(\mathbf{a}) = \sum_{i=1}^l a_i \log a_i$

Using the negative entropy as the regularizer, results in the usual softmax projection operator $\Pi_{\Omega}(\mathbf{s}) = \frac{\exp(\mathbf{s}/\gamma)}{\sum_{i=1}^l \exp(s_i/\gamma)}$.

Fusedmax: $\Omega(\mathbf{a}) = \frac{1}{2} \|\mathbf{a}\|_2^2 + \sum_{i=1}^l |a_{i+1} - a_i|$
Using squared loss with fused-lasso penalty (Niculae and Blondel, 2017), results in a projection operator which produces sparse as well as contiguous attention weights⁵. The fusedmax projection operator can be written as $\Pi_{\Omega}(\mathbf{s}) = P_{\Delta^l}(P_{TV}(\mathbf{s}))$, where

$$P_{TV}(\mathbf{s}) = \arg \min_{\mathbf{y} \in \mathbb{R}^l} \|\mathbf{y} - \mathbf{s}\|_2^2 + \sum_{d=1}^{l-1} |y_{d+1} - y_d|$$

is the proximal operator for 1d Total Variation Denoising problem, and P_{Δ^l} is the euclidean projection operator. Both these operators can be computed non-iteratively as described in Condat (2013) and Duchi et al. (2008), respectively. The gradient of Fusedmax operator can be efficiently computed as described in Niculae and Blondel (2017).⁶

Fusedmax*: We observe that while softmax learns to focus on the right region of text, it tends to assign very low attention weights to some tokens of phrases resulting in multiple discontinuous spans per attribute, while Fusedmax on the other hand, almost always generates contiguous attention weights. However, Fusedmax makes more mistakes in identifying the overall region that contains

⁵Some example outputs of softmax and fusedmax on random inputs are shown in Appendix A.3

⁶The pytorch implementation to compute fusedmax used in this work is available at <https://github.com/dhruvdcoder/sparse-structured-attention>.

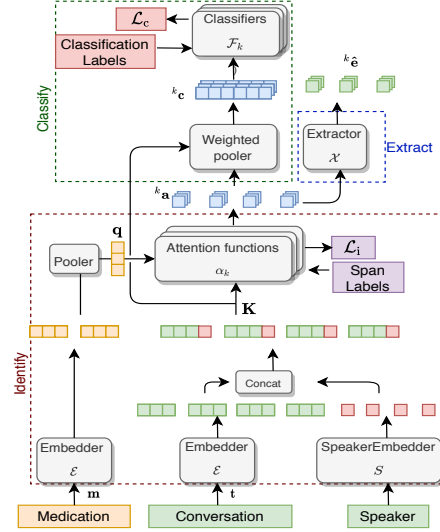


Figure 3: Complete model for weakly supervised MR attribute extraction.

the target span (Section 6.3). In order to combine the advantages of softmax and Fusedmax, we first train a model using softmax as the projector and then swap the softmax with Fusedmax in the final few epochs. We call this approach Fusedmax*.

4 Model

Our classification + extraction model uses MR attributes classification labels to extract MR attributes. The model can be divided into three phases: identify, classify and extract (Figure 3). The identify phase encodes the input text and medication name and uses the attention bottleneck to produce attention over the text. Classify phase computes the context vector using the attention from the identify phase and classifies the context vectors. Finally, the extract phase uses the attention from the identify phase to extract spans corresponding to MR attributes.

Notation: Let the dataset \mathcal{D} be $\{(\mathbf{x}^{(1)}, \mathbf{y}^{(1)}), \dots, (\mathbf{x}^{(N)}, \mathbf{y}^{(N)})\}$. Each \mathbf{x} consists of a medication m and conversation text \mathbf{t} , and each \mathbf{y} consists of classification labels for *frequency*, *route* and *change*, i.e. $\mathbf{y} = ({}^f y, {}^r y, {}^c y)$, respectively. The number of classes for each attribute is denoted by ${}^{(\cdot)} n$. As seen from Table 1, ${}^f n = 12$, ${}^r n = 10$ and ${}^c n = 8$. The length of a text excerpt is denoted by l . The extracted span for attribute $k \in \{f, r, c\}$ is denoted by a binary vector ${}^k \mathbf{e}$ of length l , such that ${}^k e_j = 1$, if j^{th} token is in the extracted span for attribute k .

4.1 Identify

As shown in the Figure 3, the *identify* phase finds the most relevant parts of the text w.r.t each of the three attributes. For this, we first encode the text as well as the given medication using a contextualized token embedder \mathcal{E} . In our case, this is 1024 dimensional BERT (Devlin et al., 2019)⁷. Since BERT uses WordPiece representations (Wu et al., 2016), we average these wordpiece representations to form the word embeddings. In order to supply the speaker information, we concatenate a 2-dimensional fixed vocabulary speaker embedding to every token embedding in the text to obtain speaker-aware word representations.

We then perform average pooling of the medication representations to get a single vector representation for the medication⁸. Finally, with the given medication representation as the query and the speaker-aware token representations as the key, we use three separate attention functions (attention bottleneck), one for each attribute (no weight sharing), to produce three sets of normalized attention distributions $^f\hat{\mathbf{a}}$, $^r\hat{\mathbf{a}}$ and $^c\hat{\mathbf{a}}$ over the tokens of the text. The *identify* phase can be succinctly described as follows:

$$^k\mathbf{a} = {}^k\alpha(\mathcal{E}(m), \mathcal{E}(t)), \text{ where } k \in \{f, r, c\}$$

Here, each $^k\hat{\mathbf{a}}$ is an element of the probability simplex Δ^l and is used to perform attribute extraction (Section 4.3).

4.2 Classify

We obtain the attribute-wise context vectors $^k\mathbf{c}$, as the weighted sum of the encoded tokens (\mathbf{K} in Figure 3) where the weights are given by the attribute-wise attention distributions $^k\mathbf{a}$. To perform the classification for each attribute, the attribute-wise context vectors are used as input to feed-forward neural networks \mathcal{F}_k (one per attribute), as shown below:⁹

$$^k\mathbf{p} = \text{softmax} \left(\mathcal{F}_k(^k\mathbf{c}) \right)$$

$$^k\hat{y} = \arg \max_{j \in \{1, 2, \dots, k_n\}} {}^k p_j, \text{ where } k \in \{f, r, c\}.$$

⁷The pre-trained weight for BERT is from the Hugging-Face library (Wolf et al., 2019)

⁸Most medication names are single word, however a few medicines have names which are upto 4-5 words.

⁹Complete set of hyperparameters used is given in Appendix A.2

4.3 Extract

The spans are extracted from the attention distribution using a fixed extraction function $\mathcal{X}: \Delta^l \rightarrow \{0, 1\}^l$, defined as:

$$^k\hat{e}_j = \mathcal{X}_k(^k\mathbf{a})_j = \begin{cases} 1 & \text{if } ^k a_j > ^k\gamma \\ 0 & \text{if } ^k a_j \leq ^k\gamma, \end{cases}$$

where $^k\gamma$ is the extraction threshold for attribute k . For softmax projection function, it is important to tune the attribute-wise extraction thresholds γ . We tune these using extraction performance on the extraction validation set. For fusedmax projection function which produces sparse weights, the thresholds need not be tuned, and hence are set to 0.

4.4 Training

We train the model end-to-end using gradient descent, except the *extract* module (Figure 3), which does not have any trainable weights, and the embedder \mathcal{E} . Freezing the embedder is vital for the performance, since not doing so results in excessive dispersion of token information to other nearby tokens, resulting in poor extractions.

The total loss for the training is divided into two parts as described below.

(1) Classification Loss \mathcal{L}_c : In order to perform classification with highly class imbalanced data (see Table 1), we use weighted cross-entropy:

$$\mathcal{L}_c = \sum_{k \in \{f, r, c\}} - {}^k w_{k_y} \log \left({}^k p_{k_y} \right),$$

where the class weights ${}^k w_{k_y}$ are obtained by inverting each class' relative proportion.

(2) Identification Loss \mathcal{L}_i : If span labels \mathbf{e} are present for some subset \mathcal{A} of training examples, we first normalize these into ground truth attention probabilities \mathbf{a} :

$$^k a_j = \frac{{}^k e_j}{\sum_{j=1}^l {}^k e_j} \text{ for } k \in \{f, r, c\}$$

We then use KL-Divergence between the ground truth attention probabilities and the ones generated by the model ($\hat{\mathbf{a}}$) to compute identification loss $\mathcal{L}_i = \sum_{k \in \{f, r, c\}} \text{KL} \left({}^k \mathbf{a} \parallel {}^k \hat{\mathbf{a}} \right)$. Note that \mathcal{L}_i is zero for data-points that do not have span labels. Using these two loss functions, the overall loss $\mathcal{L} = \mathcal{L}_c + \lambda \mathcal{L}_i$.

Model			Span labels	Token-wise extraction F1				LCSF1				Classification F1			
Encoder	Scorer	Projector		freq.	route	change	Avg.	freq.	route	change	Avg.	freq.	route	change	Avg.
<i>Phrase-based baseline</i>			-	41.03	48.57	10.75	33.45	36.26	50.41	11.54	32.73	-	-	-	-
BERT	Additive	Softmax	0	51.22	46.27	22.81	40.10	39.87	46.40	18.92	35.06	51.51	54.06	51.65	52.40
BERT	Additive	Fusedmax	0	47.55	51.31	5.10	34.65	46.39	59.10	4.82	36.77	43.54	42.91	9.19	31.88
BERT	TAScore	Softmax	0	66.53	48.96	27.61	47.70	61.49	47.34	22.49	43.77	44.93	51.34	46.49	47.58
BERT	TAScore	Fusedmax	0	56.35	44.04	22.07	40.82	61.96	50.27	25.25	45.82	51.95	48.37	43.00	47.77
BERT	Additive	Softmax	150	61.56	45.08	33.54	46.73	57.90	48.14	28.28	44.77	55.62	52.42	50.40	52.81
BERT	Additive	Fusedmax	150	47.05	52.49	27.69	42.41	42.37	57.50	30.63	43.50	54.04	48.40	52.28	51.57
BERT	Additive	Fusedmax*	150	65.90	47.30	34.77	49.32	67.15	51.12	36.04	51.30	56.46	42.63	50.68	49.93
BERT	TAScore	Softmax	150	66.53	54.35	34.27	51.72	62.90	53.05	28.33	48.09	50.13	45.86	47.16	47.72
BERT	TAScore	Fusedmax	150	58.24	58.09	25.09	47.32	57.93	64.05	26.70	49.56	51.61	53.95	43.51	49.69
BERT	TAScore	Fusedmax*	150	66.90	54.85	33.28	51.67	70.10	60.05	35.92	55.36	64.26	44.50	51.21	53.32

Table 3: Attribute extraction performance for various combinations of scoring and projection functions. The avg. columns represent the macro average of the corresponding metric across the attributes.

Training Type	Model	Tokenwise Extraction F1				Classification F1			
		freq.	route	change	avg.	freq.	route	change	avg.
Classification only	BERT Classifiers	-	-	-	-	74.72	40.82	55.76	58.48
Classification only	BERT+TAScore+Fusedmax*	58.55	45.00	24.43	42.66	52.45	46.37	43.00	47.27
Extraction only	BERT+TAScore+Fusedmax*	53.79	44.44	14.32	37.18	-	-	-	-
Classification +Extraction	BERT+TAScore+Fusedmax*	66.90	54.85	33.28	51.67	64.26	44.50	51.21	53.32

Table 4: Effect of performing extraction+classification jointly in our proposed model. While the *Extraction Only* training only uses the 150 examples which are explicitly annotated with span labels, the *Classification only* training uses the complete training dataset with classification labels.

5 Metrics

Token-wise F1 (TF1): Each token in text is either part of the extracted span (positive class) for an attribute or not (negative class). Token-wise F1 score is the F1 score of the positive class obtained by considering all the tokens in the dataset as separate binary classification data points. TF1 is calculated separately for each attribute.

Longest Common Substring F1 (LCSF1): LCSF1 measures if the extracted spans, along with being part of the gold spans, are contiguous or not. Longest Common Substring (LCS) is the longest overlapping contiguous span of tokens between the predicted and gold spans. LCSF1 is defined as the harmonic mean of LCS-Recall and LCS-Precision which are defined per extraction as:

$$\text{LCS-Recall} = \frac{\text{\#tokens in LCS}}{\text{\#tokens in gold span}}$$

$$\text{LCS-Precision} = \frac{\text{\#tokens in LCS}}{\text{\#tokens in predicted span}}$$

6 Results and Analysis

Table 3 shows the results obtained by various combinations of attention scoring and projection functions on the task of MR attribute extraction in terms of the metrics defined in Section 5. It also shows

the classification F1 score to emphasize how the attention bottleneck affects classification performance. The first row shows how a simple phrase based extraction system would perform on the task.¹⁰

6.1 Effect of Span labels

In order to see if having a small number of extraction training data-points (containing explicit span labels) helps the extraction performance, we annotate 150 (see Section 2 for how we sampled the datapoints) of the training data-points with span labels. As seen from Table 3, even a small number of examples with span labels ($\approx 0.3\%$) help a lot with the extraction performance for all models. We think this trend might continue if we add more training span labels. We leave the finding of the right balance between annotation effort and extraction performance as a future direction to explore.

6.2 Effect of classification labels

In order to quantify the effect of performing the auxiliary task of classification along with the main task of extraction, we train the proposed model in three different settings. (1) The *Classification*

¹⁰The details about the phrase based baseline are presented in Appendix A.4

Only uses the complete dataset (~45k) but only with the classification labels. (2) The *Extraction Only* setting only uses the 150 training examples that have span labels. (3) Finally, the *Classification+Extraction* setting uses the 45k examples with classification labels along with the 150 examples with the span labels to train the model. Table 4 (rows 2, 3 and 4) shows the effect of having classification labels and performing extraction and classification jointly using the proposed model. The model structure and the volume of the classification data (~45k examples) makes the auxiliary task of classification extremely helpful for the main task of extraction, even with the presence of label noise.

It is worth noting that the classification performance of the proposed method is also improved by explicit supervision to the extraction portion of the model (row 2 vs 4, Table 4). In order to set a reference for classification performance, we train strong classification only models, one for each attribute, using pretrained BERT. These *BERT Classifiers*, are implemented as described in Devlin et al. (2019) with input consisting of the text and medication name separated by a [SEP] token (row 1). Based on the improvements achieved in the classification performance using span annotations, we believe that having more span labels can further close the gap between the classification performance of the proposed model and the BERT Classifiers. However, this work focuses on extraction performance, hence improving the classification performance is left to future work.

6.3 Effect of projection function

While softmax with post-hoc threshold tuning achieves consistently higher TF1 compared to fusedmax (which does not require threshold tuning), the later achieves better LCSF1. We observe that while the attention function using softmax projection focuses on the correct portion of the text, it drops intermediate words, resulting in multiple discontinuous spans. Fusedmax on the other hand almost always produces contiguous spans. Figure 4 further illustrates this point using a test example. The training trick which we call fusedmax* swaps the softmax projection function with fusedmax during the final few epochs to combine the strengths of both softmax and fusedmax. This achieves high LCSF1 as well as TF1.

```

... ..
DR: Okay, Fosamax, you take one pill on Monday and one on Thursday.
DR: Do you use much caffeine?
PT: No, none.
DR: Okay, this is Actonel and it's one tablet once a month.
DR : Do you get a one month or a three months supply in your prescriptions?

```

(a) BERT+TAScore+Fusedmax*

```

... ..
DR: Okay, Fosamax, you take one pill on Monday and one on Thursday.
DR: Do you use much caffeine?
PT: No, none.
DR: Okay, this is Actonel and it's one tablet once a month.
DR : Do you get a one month or a three months supply in your prescriptions?

```

(b) BERT+TAScore+Softmax

Figure 4: Difference in extracted spans for MR attributes with models that uses Fusedmax* and Softmax, for the medication Actonel. Blue: *change*, green: *route* and yellow: *frequency*. Refer Figure 1 for ground-truth annotations.

6.4 Effect of scoring function

Table 5 shows the percent change in the extraction F1 if we use TAScore instead of additive scoring (everything else being the same). As seen, there is a significant improvement irrespective of the projection function being used.

Scorer	TF1 ($\Delta\%$)			LCSF1 ($\Delta\%$)		
	MM (77.3)	SM (22.7)	All (100)	MM (77.3)	SM (22.7)	All (100)
softmax	+11.1	+10.6	+10.6	+6.5	+6.6	+6.3
fusedmax	+12.1	+8.3	+11.5	+16.4	+15.5	+13.9
fusedmax*	+5.4	+1.9	+4.7	+9.25	+1.1	+7.9

Table 5: MR extraction improvement (%) brought by TAScore over additive scorer in the full test set (All=100%), and test subset with single medication (SM=22.7%) and multiple medications (MM=77.3%) in the *text*.

The need for TAScore stems from the difficulty of the additive scoring function to resolve references between spans when there are multiple medications present. In order to measure the efficacy of TAScore for this problem, we divide the test set into two subsets: data-points which have multiple distinct medications in their text (MM) and data-points that have single medication only. As seen from the first two columns for both the metrics in Table 5, using TAScore instead of additive results in more improvement in the MM-subset compared to the SM-subset, showing that using transformer scorer does help with resolving references when multiple medications are present in the text.

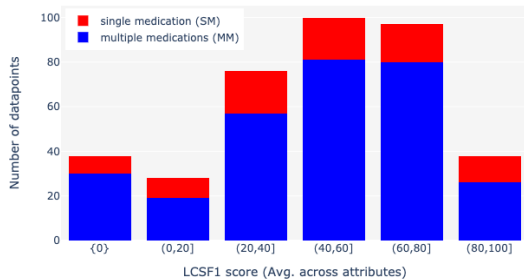


Figure 5: Distribution of the Avg. LCSF1 for the best performing model (BERT+TAScore+Fusedmax*). A significant number ($\approx 10\%$) of datapoints with multiple medication in their *text* get LCSF1 of zero (1st bar).

Figure 5 shows the distribution of Avg. LCSF1 (average across all three attributes). It can be seen that there are a significant number of datapoints in the MM subset which get LCSF1 of zero, showing that even when the transformer scorer achieves improvement on MM subset, it gets quite a lot of these data-points completely wrong. This shows that there is still room for improvement.

6.5 Discussion

In summary, our analysis reveals that Fusedmax/Fusedmax* favors contiguous extraction spans which is a necessity for our task. Irrespective of the projection function used, the proposed scoring function TAScore improves the extraction performance when compared to the popular additive scoring function. The proposed model architecture is able to establish a synergy between the classification and span extraction tasks where one improves the performance of the other. Overall, the proposed combination of TAScore and Fusedmax* achieves a 22 LCSF1 points improvement over the phrase-based baseline and 10 LCSF1 points improvement over the naive additive and softmax combination.

7 Related Work

Existing literature directly related to our work can be bucketed into two categories – related methods and related tasks.

Methods: The recent work on generating rationales/explanations for deep neural network based classification models (Lei et al., 2016; Bastings et al., 2020; Paranjape et al., 2020) is closely related to ours in terms of the methods used. Most of these works use binary latent variables to perform extraction as an intermediate step before classification. Our work is closely related to (Jain et al., 2020; Zhong et al., 2019), who use attention scores to

generate rationales for classification models. These works, however, focus on generating *faithful* and *plausible* explanation for classification as opposed to extracting the spans for attributes of an entity, which is the focus of our work. Moreover, our method can be generalized to any number of attributes while all these methods would require a separate model for each attribute.

Tasks: Understanding doctor-patient conversations is starting to receive attention recently (Rajkumar et al., 2019; Schloss and Konam, 2020). Selvaraj and Konam (2019) performs MR extraction by framing the problem as a generative question answering task. This approach is not efficient at inference time – it requires one forward pass for each attribute. Moreover, unlike a span extraction model, the generative model might produce hallucinated facts. Du et al. (2019) obtain MR attributes as spans in text; however, they use a fully supervised approach which requires a large dataset with span-level labels.

8 Conclusion and Future work

We provide a framework to perform MR attribute extraction from medical conversations with weak supervision using noisy classification labels. This is done by creating an attention bottleneck in the classification model and performing extraction using the attention weights. After experimenting with several variants of attention scoring and projection functions, we show that the combination of our transformer-based attention scoring function (TAScore) combined with Fusedmax* achieves significantly higher extraction performance compared to the other attention variants and a phrase-based baseline.

While our proposed method achieves good performance, there is still room for improvement, especially for text with multiple medications. Data augmentation by swapping or masking medication names is worth exploring. An alternate direction of future work involves improving the naturalness of extracted spans. Auxiliary supervision using a language modeling objective would be a promising approach for this.

Acknowledgments

We thank University of Pittsburgh Medical Center (UPMC) and Abridge AI Inc. for providing access to the de-identified data corpus.

References

- J. L. Anderson, Sally Dodman, M. Kopelman, and A. Fleming. 1979. Patient Information Recall in a Rheumatology Clinic. *Rheumatology*, 18:18–22.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. *CoRR*, abs/1409.0473.
- Joost Bastings, Wilker Aziz, and Ivan Titov. 2020. Interpretable neural predictions with differentiable binary variables. In *ACL 2019 - 57th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference*, pages 2963–2977. Association for Computational Linguistics.
- Lukas Biewald. 2020. Experiment tracking with weights and biases. Software available from wandb.com.
- Mathieu Blondel, Andre F.T. Martins, and Vlad Niculae. 2020. Learning with fenchel-young losses. *Journal of Machine Learning Research*, 21(35):1–69.
- Laurent Condat. 2013. A direct algorithm for 1-D total variation denoising. *IEEE Signal Processing Letters*, 20(11):1054–1057.
- 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.
- Nan Du, Mingqiu Wang, Linh Tran, Gang Li, and Izhak Shafran. 2019. Learning to infer entities, properties and their relations from clinical conversations. *arXiv preprint arXiv:1908.11536*.
- John Duchi, Shai Shalev-Shwartz, Yoram Singer, and Tushar Chandra. 2008. Efficient projections onto the l_1 -ball for learning in high dimensions. In *Proceedings of the 25th International Conference on Machine Learning, ICML '08*, page 272–279, New York, NY, USA. Association for Computing Machinery.
- Matt Gardner, Joel Grus, Mark Neumann, Oyvind Taffjord, Pradeep Dasigi, Nelson F. Liu, Matthew Peters, Michael Schmitz, and Luke S. Zettlemoyer. 2017. Allennlp: A deep semantic natural language processing platform.
- Stuart W Grande, Mary Ganger Castaldo, Elizabeth Carpenter-Song, Ida Griesemer, and Glyn Elwyn. 2017. A digital advocate? reactions of rural people who experience homelessness to the idea of recording clinical encounters. *Health Expectations*, 20(4):618–625.
- Ruidan He, Wee Sun Lee, Hwee Tou Ng, and Daniel Dahlmeier. 2017. An unsupervised neural attention model for aspect extraction. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 388–397, Vancouver, Canada. Association for Computational Linguistics.
- Sarthak Jain and Byron C. Wallace. 2019. Attention is not explanation. In *NAACL-HLT*.
- Sarthak Jain, Sarah Wiegrefe, Yuval Pinter, and Byron C Wallace. 2020. Learning to Faithfully Rationalize by Construction.
- Shailesh Kumar. 2016. Burnout and doctors: prevalence, prevention and intervention. In *Healthcare*, volume 4, page 37. Multidisciplinary Digital Publishing Institute.
- Tao Lei, Regina Barzilay, and Tommi Jaakkola. 2016. Rationalizing neural predictions. In *EMNLP 2016 - Conference on Empirical Methods in Natural Language Processing, Proceedings*, pages 107–117.
- Yankai Lin, Shiqi Shen, Zhiyuan Liu, Huanbo Luan, and Maosong Sun. 2016. Neural relation extraction with selective attention over instances. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2124–2133, Berlin, Germany. Association for Computational Linguistics.
- Lisa C. McGuire. 1996. Remembering what the doctor said: Organization and adults' memory for medical information. *Experimental Aging Research*, 22:403–428.
- Vlad Niculae and Mathieu Blondel. 2017. A regularized framework for sparse and structured neural attention. In *Advances in neural information processing systems*, pages 3338–3348.
- Bhargavi Paranjape, Mandar Joshi, John Thickstun, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2020. An Information Bottleneck Approach for Controlling Conciseness in Rationale Extraction. Technical report.
- Alvin Rajkomar, Anjuli Kannan, Kai Chen, Laura Vardoulakis, Katherine Chou, Claire Cui, and Jeffrey Dean. 2019. Automatically charting symptoms from patient-physician conversations using machine learning. *JAMA internal medicine*, 179(6):836–838.
- Benjamin Schloss and Sandeep Konam. 2020. Towards an automated soap note: Classifying utterances from medical conversations. *Machine Learning for Health Care, 2020*, arXiv:2007.08749. Version 3.
- Sai P Selvaraj and Sandeep Konam. 2019. Medication regimen extraction from clinical conversations. *arXiv preprint arXiv:1912.04961*.

- Christine Sinsky, Lacey Colligan, Ling Li, Mirela Prgomet, Sam Reynolds, Lindsey Goeders, Johanna Westbrook, Michael Tutty, and George Blike. 2016. Allocation of physician time in ambulatory practice: a time and motion study in 4 specialties. *Annals of internal medicine*, 165(11):753–760.
- Maka Tsulukidze, Marie-Anne Durand, Paul J Barr, Thomas Mead, and Glyn Elwyn. 2014. Providing recording of clinical consultation to patients—a highly valued but underutilized intervention: a scoping review. *Patient Education and Counseling*, 95(3):297–304.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, R’emi Louf, Morgan Funtowicz, and Jamie Brew. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *ArXiv*, abs/1910.03771.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Łukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2016. [Google’s neural machine translation system: Bridging the gap between human and machine translation](#).
- Bowen Yu, Zhenyu Zhang, Tingwen Liu, Bin Wang, Sujian Li, and Quangang Li. 2019. Beyond word attention: Using segment attention in neural relation extraction. In *IJCAI*, pages 5401–5407.
- Ruiqi Zhong, Steven Shao, and Kathleen McKeown. 2019. [Fine-grained sentiment analysis with faithful attention](#).

A Appendices

A.1 Data

The complete set of *normalized* classification labels for all three medication attributes and their meaning is shown in Table 7.

Average statistics about the dataset are shown in Table 6.

	min	max	mean	σ
#utterances in <i>text</i>	3	20	7.8	2.3
#words in <i>text</i>	12	565	80.8	41.0
#words in <i>freq</i> span	1	21	4.4	2.6
#words in <i>route</i> span	1	9	1.5	1.0
#words in <i>change</i> span	1	34	6.8	4.9

Table 6: Statistics of extraction labels (#words) and the corresponding *text*

A.2 Hyperparameters

We use AllenNLP (Gardner et al., 2017) to implement our models and Weights&Biases (Biewald, 2020) to manage our experiments. Following is the list of hyperparameters used in our experiments:

- Contextualized Token Embedder:** We use 1024-dimensional 24-layer bert-large-cased obtained as a pre-trained model from HuggingFace¹¹. We freeze the weights of the embedder in our training. The max sequence length is set to 256.
- Speaker embedding:** 2-dimensional trainable embedding with vocabulary size of 4 as we only have 4 unique speakers in our dataset: doctor, patient, caregiver and nurse.
- Softmax and Fusedmax:** The temperatures of softmax and fusedmax are set to a default value of 1. The sparsity weight of fusedmax is also set to its default value of 1 for all attributes.
- TAScore:** The transformer used in TAScore is a 2-layer transformer encoder where each layer is implemented as in Vaswani et al. (2017). Both the hidden dimensions inside the transformer (self-attention and feedforward) are set to 32 and all the dropout probabilities

are set to 0.2. The linear layer for the query has input and output dimensions of 1024 and 32, respectively. Due to the concatenation of speaker embedding, the linear layer for keys has input and output dimensions of 1026 and 32, respectively. The feedforward layer (which generates scalar scores for each token) on top of the transformer is 2-layered with relu activations and hidden sizes (16, 1).

- Classifiers:** The final classifier for each attribute is a 2-layer feedforward network with hidden sizes (512, “number of classes for the attribute”) and dropout probability of 0.2.


A.3 Examples: Projection Functions

Figures 6 and 7 show examples of outputs of projection functions softmax and fusedmax on random input scores.


A.4 Phrase based extraction baseline

We implement a phrase based extraction system to provide a baseline for the extraction task. A lexicon of relevant phrases is created for each class for each attribute as shown in Table 8. We then look for string matches within these phrases and the text for the data-point. If there are matches then the longest match is considered as an extraction span for that attribute.


¹¹<https://huggingface.co/bert-large-cased>



(a) Positive and negative scores




(b) Positive scores only




(c) More uniformly distributed positive scores


Figure 6: Sample outputs (right column) of softmax function on random input scores (left column).



(a) Positive and negative scores



(b) Positive scores only



(c) More uniformly distributed positive scores

Figure 7: Sample outputs (right column) of fusedmax function on random input scores (left column).

Attribute	class	Meaning	class proportion
<i>frequency</i>	Daily	Take the medication once a day (specific time not mentioned).	8.0
	Every morning	Take the medication once every morning.	0.9
	At Bedtime	At Bedtime	1.7
	Twice a day	Twice a day	6.5
	Three times a day	Three times a day	1.6
	Every six hours	Every six hours	0.2
	Every week	Every week	0.9
	Twice a week	Twice a week	0.2
	Three times a week	Three times a week	0.3
	Every month	Every month	0.3
	Other	Other	1.5
	None	None	77.9
<i>route</i>	Pill	Pill	6.8
	Injection	Injection	3.5
	Topical cream	Topical cream	1.0
	Nasal spray	Nasal spray	0.5
	Medicated patch	Medicated patch	0.2
	Ophthalmic solution	Ophthalmic solution	0.2
	Inhaler	Inhaler	0.2
	Oral solution	Oral solution	0.1
	Other	Other	2.1
	None	None	85.5
<i>change</i>	Take	Take	83.1
	Stop	Stop	6.5
	Increase	Increase	5.2
	Decrease	Decrease	2.0
	None	None	1.6
	Other	Other	1.4

Table 7: Complete set of normalized classification labels for all three medication attributes and their explanation

Attribute	Class	Phrases
freq	Every Morning	everyday in the morning every morning morning
	At Bedtime	everyday before sleeping everyday after dinner every night after dinner at bedtime before sleeping
	Twice a day	twice a day 2 times a day two times a day 2 times per day two times per day
	Three times a day	3 times a day 3 times per day 3 times every day
	Every six hours	every 6 hours every six hours
	Every week	every week weekly once a week
	Twice a week	twice a week two times a week 2 times a week twice per week two times per week 2 times per week
	Three times a week	3 times a week 3 times per week
	Every month	every month monthly once a month
	Other	
None		
route	Pill	tablet pill capsule mg
	Injection	pen shot injector injection inject
	Topical cream	cream gel ointment lotion
	Nasal spray	spray nasal conversation transcript.
	Medicated patch	patch
	Ophthalmic solution	ophthalmic drops drop
	Oral solution	oral solution
	Other	
	None	
change	Take	take start put you on continue
	Stop	stop off
	Increase	increase
	Decrease	reduce decrease
	Other	
	None	

Table 8: Phrases used in the phrase based baseline. These are also the most frequently occurring phrases in the free-form annotations.