

Original Paper

Extraction of Explicit and Implicit Cause-Effect Relationships in Patient-Reported Diabetes-Related Tweets From 2017 to 2021: Deep Learning Approach

Adrian Ahne^{1,2}, PhD; Vivek Khetan³, MSc; Xavier Tannier⁴, Prof Dr; Md Imbesat Hassan Rizvi⁵, MSc; Thomas Czernichow², PhD; Francisco Orchard², MSc; Charline Bour⁶, MSc; Andrew Fano³, PhD; Guy Fagherazzi⁶, PhD

¹Center of Epidemiology and Population Health, Inserm, Hospital Gustave Roussy, Paris-Saclay University, Villejuif, France

²Epiconcept Company, Paris, France

³Accenture Labs, San Francisco, CA, United States

⁴Laboratoire d'Informatique Médicale et d'Ingénierie des Connaissances pour la e-Santé, Inserm, University Sorbonne Paris Nord, Sorbonne University, Paris, France

⁵Indian Institute of Science, Bengaluru, India

⁶Deep Digital Phenotyping Research Unit, Department of Precision Health, Luxembourg Institute of Health, Strassen, Luxembourg

Corresponding Author:

Adrian Ahne, PhD
Center of Epidemiology and Population Health
Inserm, Hospital Gustave Roussy
Paris-Saclay University
20 Rue du Dr Pinel
Villejuif, 94800
France
Phone: 33 142115386
Email: adrian.ahne@protonmail.com

Abstract

Background: Intervening in and preventing diabetes distress requires an understanding of its causes and, in particular, from a patient's perspective. Social media data provide direct access to how patients see and understand their disease and consequently show the causes of diabetes distress.

Objective: Leveraging machine learning methods, we aim to extract both explicit and implicit cause-effect relationships in patient-reported diabetes-related tweets and provide a methodology to better understand the opinions, feelings, and observations shared within the diabetes online community from a causality perspective.

Methods: More than 30 million diabetes-related tweets in English were collected between April 2017 and January 2021. Deep learning and natural language processing methods were applied to focus on tweets with personal and emotional content. A cause-effect tweet data set was manually labeled and used to train (1) a fine-tuned BERTweet model to detect causal sentences containing a causal relation and (2) a conditional random field model with Bidirectional Encoder Representations from Transformers (BERT)-based features to extract possible cause-effect associations. Causes and effects were clustered in a semisupervised approach and visualized in an interactive cause-effect network.

Results: Causal sentences were detected with a recall of 68% in an imbalanced data set. A conditional random field model with BERT-based features outperformed a fine-tuned BERT model for cause-effect detection with a macro recall of 68%. This led to 96,676 sentences with cause-effect relationships. "Diabetes" was identified as the central cluster followed by "death" and "insulin." Insulin pricing-related causes were frequently associated with death.

Conclusions: A novel methodology was developed to detect causal sentences and identify both explicit and implicit, single and multiword cause, and the corresponding effect, as expressed in diabetes-related tweets leveraging BERT-based architectures and visualized as cause-effect network. Extracting causal associations in real life, patient-reported outcomes in social media data provide a useful complementary source of information in diabetes research.

(*JMIR Med Inform* 2022;10(7):e37201) doi: [10.2196/37201](https://doi.org/10.2196/37201)

KEYWORDS

causality; deep learning; natural language processing; diabetes; social media; causal relation extraction; social media data; machine learning

Introduction

Diabetes distress refers to psychological factors such as emotional burden, worries, frustration, or stress in the day-to-day management of all types of diabetes [1-3]. Diabetes distress is associated with poor quality of life [4], high hemoglobin A_{1c} levels [5,6], and low medication adherence [7]. Reducing diabetes distress may improve hemoglobin A_{1c} levels and reduce the burden of disease among people with diabetes [8]. Social media is a useful observatory resource for patient-reported diabetes issues and could help to contribute directly to public and clinical decision-making from a patient's perspective, given the active online diabetes community [9,10]. Identifying causal relations in expressed text data in social media platforms might help to discover unknown etiological results, specifically, causes of health problems, concerns, and symptoms.

To intervene and potentially prevent diabetes distress, it is necessary to understand the causes of diabetes distress from a patient's perspective to understand how patients see their disease. Causal relation extraction in natural language text has gained popularity in clinical decision-making, biomedical knowledge discovery, or emergency management [11]. In particular, causal relations on Twitter have been examined for diverse factors causing stress and relaxation [12], adverse drug reactions [13], or causal associations related to insomnia or headache [14]. Most approaches examine *explicit* causality in text [14-16], when cause and effect are explicitly stated, for instance, by connective words (eg, so, hence, because, lead to, since, if-then) [11,17]. An example for an *explicit* cause-effect pair is "diabetes causes hypoglycemia." However, *implicit* causality is more complicated to detect such as in "I reversed diabetes with lifestyle changes" with cause "lifestyle changes" and effect "reversed diabetes."

Natural language processing methods explore among other things how computers can be used to extract useful information from natural language documents. In combination with machine learning and deep learning models, which are artificial intelligence algorithms designed to learn from experience, they have also been applied to extract causal relations [18,19]. Machine learning methods are able to explore implicit relations and provide better generalization contrary to rule-based approaches [11,20-22]. An interesting approach leveraging the transfer learning paradigm and addressing both explicit and

implicit cause-effect extraction is provided by Khetan et al [23]. They fine-tuned pretrained transformer-based Bidirectional Encoder Representations from Transformers (BERT) language models [24,25] to detect "cause-effect" relationships by using publicly available data sets such as the adverse drug effect data set [26]. More generally, the idea of transfer learning is to leverage the knowledge of a model that has been trained on an auxiliary domain [27].

In this study, we aimed to extract spans of text as 2 distinct events from diabetes and diabetes-related tweets such that one event directly (explicit) or indirectly (implicit) impacts another event. We categorized these events as cause-event and effect-event depending upon the expressed context of each tweet. The identified cause and effect will then be aggregated into clusters and ultimately visualized in an interactive cause-effect network.

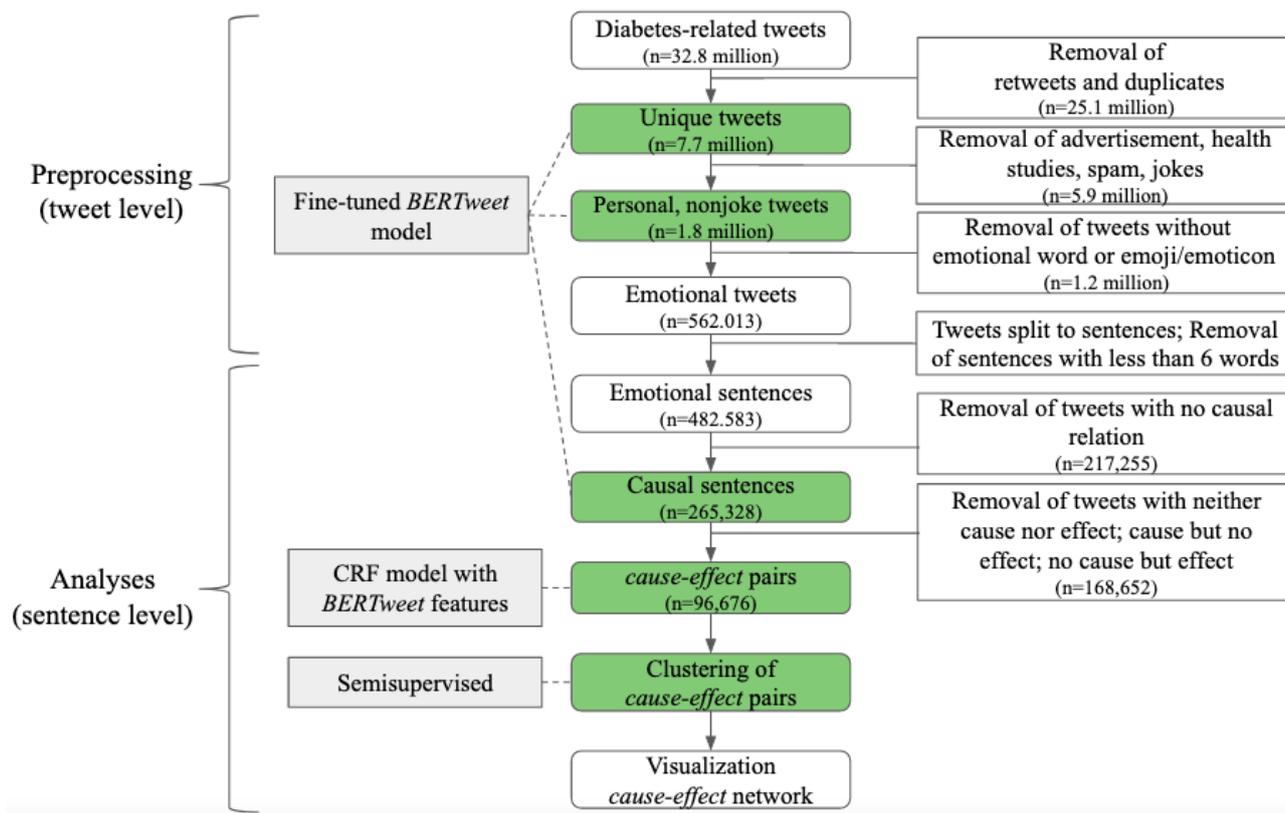
This work is realized in the frame of the World Diabetes Distress Study, which aims to analyze what is shared on social media worldwide to better understand what people with diabetes and diabetes distress are experiencing [28,29]. The social network "Twitter" is a popular data resource among diabetes researchers owing to its public character and its active online diabetes community compared to other social media [30,31]. Recent studies suggest an overrepresentation of people with type 1 diabetes compared to those with type 2 diabetes who are active on Twitter [9,31].

Methods

Overview

On the basis of diabetes-related tweets, we first preprocessed tweets to only focus on personal, nonjoke, and emotional content. Second, after this preprocessing step, we split tweets into sentences for our analyses, as we aimed to identify the cause-effect relationships between events within a sentence (sentence level) and not across multiple sentences (tweet level). This also simplifies model training and helps with easier learning. Third, we identified sentences in which causal information (opinion, observation, etc) is communicated. In the fourth step, causes and their corresponding effects were extracted. Lastly, those cause-effect pairs were aggregated, described, and visualized. The entire workflow is illustrated in Figure 1.

Figure 1. Workflow. The steps shown in green include machine learning methods. CRF: conditional random field.



Data Collection and Ethical Considerations

Via Twitter's streaming application programming interface, 32 million diabetes-related tweets in English were collected between April 2017 and January 2021 based on a list of diabetes-related keywords such as *diabetes*, *hypoglycemia*, *hyperglycemia*, and *insulin* from all over the world (see [Multimedia Appendix 1](#) for the full list of keywords used). This is an extended data set of the one used in earlier works [9]. All data collected in this study were publicly posted on Twitter. Therefore, according to the privacy policy of Twitter, users agree to have this information available to the general public [30].

Data Preprocessing

Tweets are noisy and unstructured. They contain many misspelled or nonstandard English words. To reduce noise in the data set, we applied a preprocessing pipeline similar to that in earlier works, the details of which are summarized in [Figure 1](#) [9]. First, retweets and duplicates were removed to obtain a database with 7.7 million unique tweets. Second, we determined only tweets with *personal* content where feelings, emotions, and opinions could be shared by people with or talking about diabetes and excluded *institutional* tweets referring to commercial, news, or health information. To identify *personal* content in tweets, we leveraged the transfer learning paradigm and fine-tuned the already pretrained transformer-based language model *BERTweet*, which was pretrained on 850 million English tweets (16 billion word tokens ~ 80 GB) [25,32]. To use the model and fine-tune it for a binary sentence classification, a linear layer was added on top of the last transformer layer of the *BERTweet* model by using the

transformers package of HuggingFace [33]. The model was then fine-tuned with an extended data set of one used in earlier works, leading to a total of 4303 tweets (1539 *personal* and 2764 *institutional*) to account for a possible temporal divergence of the way people tweet [9]. The model performance to identify tweets with personal content had accuracy of 91.2%, precision of 86.2%, recall of 90.9%, and F1 score of 88.5%. The trained model was then applied to all unique tweets, resulting in a total of 2.5 million tweets with personal content. Moreover, jokes around diabetes are common on Twitter and were considered out of scope for this study as well. Similar to the *personal* content classifier, *BERTweet* was fine-tuned to detect if a tweet is a joke. For this purpose, a joke tweet data set from earlier works was extended to 1648 tweets (486 jokes, 1162 nonjokes) [9]. The performance to identify if a tweet is a joke had accuracy of 90.4%, precision of 78.5%, recall of 90.8%, and F1 score of 84.2%. Applying the joke classifier on all tweets with personal content led to a data set of 1.8 million personal nonjoke tweets.

A particular focus of this study was on studying diabetes distress and thus, the psychological factors and emotions. To capture these factors in tweets, only tweets containing an emotional element such as emojis/emoticons or emotional words were kept. Emotional words were identified based on a combination of the psychologist Parrot's hierarchical classification of emotions with the 6 primary emotions (*joy*, *love*, *surprise*, *sadness*, *anger*, *fear*) and emotional words present in common questionnaires to study diabetes distress such as the Problem Areas in Diabetes scale and Diabetes Distress Scale [34-36]. This led to 562,013 tweets containing personal, nonjoke, and emotional content. More details on the preprocessing pipeline are summarized in [Multimedia Appendix 2](#) [9,25,32-40].

Data Annotation

In order to identify causal sentences and *cause-effect* association, 5000 randomly chosen diabetes-related tweets were selected, preprocessed, split into sentences, and then manually labeled. We did not restrict ourselves to a specific area of diabetes-related causal relationships, and we included potentially

all types. Table 1 illustrates some example sentences. Only causal relationships related to diabetes were labeled as positive samples, whereas non-diabetes-related or unclear cause-effect relationships were labeled as negative samples. For a more detailed explanation on the annotation, please refer to our annotation guidelines in Multimedia Appendix 3.

Table 1. Sample sentences in different label scenarios. The examples are fictive to ensure privacy.

Sentences	Cause	Effect	Causal association	Explanation
Diabetes causes me to have mood swings	Diabetes	mood swings	1	Possible causal association
I just want to eat, I hate #diabetes	#diabetes	hate	1	Possible causal association related to diabetes distress
Scary, have a diabetic daughter but I read thousands of people a year die in the United Kingdom just from flu so why panic over corona.	— ^a	—	0	Nondiabetes or diabetes distress-related relationship. “Flu” is not diabetes-related
Had two strokes and recover now and also have high blood pressure and diabetes. 😊	—	—	0	Unclear cause-effect relationship. Not clear if “high blood pressure” or “diabetes” caused the stroke
Not sure if I've been up since 3:30 to watch Titanic or because of my anxiety over my glucose test is what keeps me up 😊	glucose test	anxiety	1	Chaining cause-effect relationship (A->B->C) Event A: glucose test Event B: anxiety Event C: been up since 3:30 => label the relationship which is closest to our study objective: diabetes and diabetes distress
My 14-year-old daughter is type 1 = malfunctioning pancreas, meaning not enough insulin being made to regulate 😊	type 1	malfunctioning pancreas; not enough insulin	1	Negation in a cause/effect is considered being part of the cause/effect as it does not alter the meaning
It is not true to think that insulin makes you feel so bad 😊	insulin	feel so bad	0	Negation is not part of cause/effect and alters the meaning

^aNot available.

Labeling cause-effect pairs is a complex task. To verify the reliability of the labeling, 2 authors labeled 500 sentences independently and we calculated Cohen κ score, a statistical measure expressing the level of agreement between 2 annotators [41]. We obtained a score of 0.83, which is interpreted as an *almost perfect* agreement according to Altman [42] and Landis and Koch [43]. Disagreements were discussed between 2 authors, and 1 author labelled the other samples, resulting in 8235 labelled sentences (7218 noncausal sentences and 1017 causal sentences) from 5000 tweets.

Models

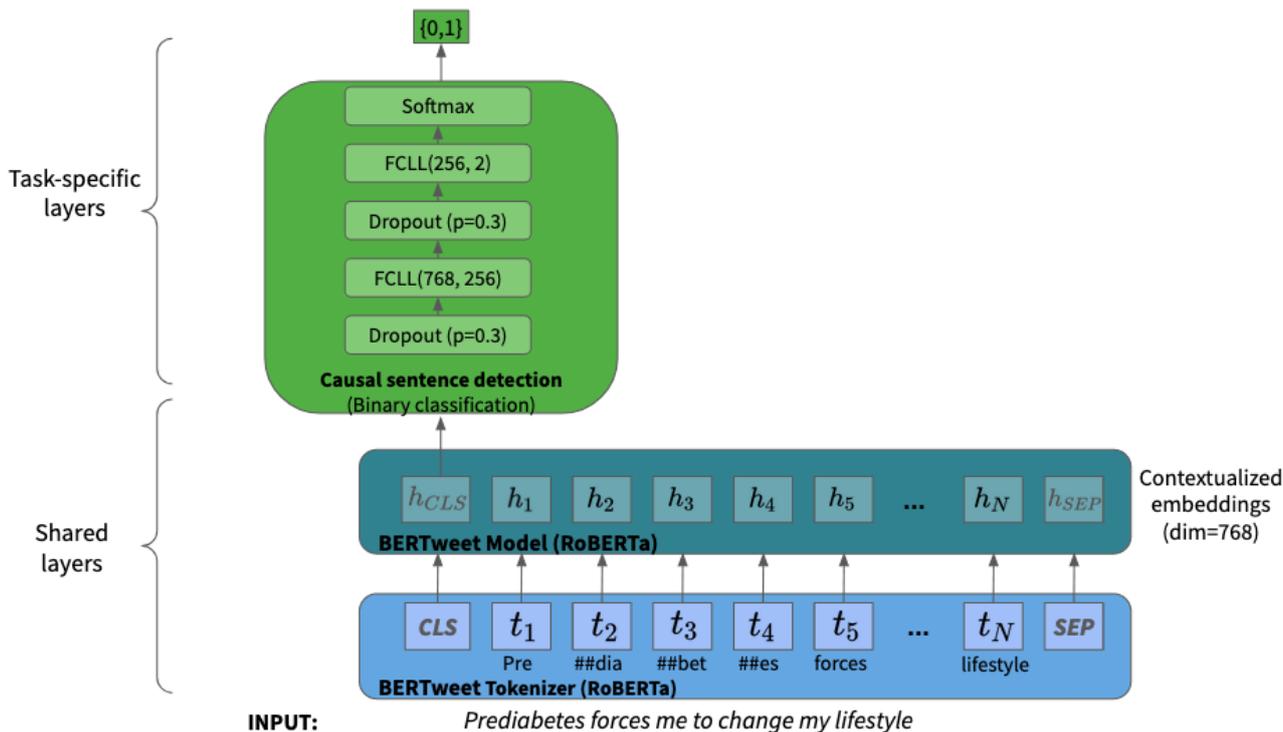
The first model was trained to predict if a sentence contains a potential cause-effect association (causal sentence), and the second model extracted the specific cause and the associated effect from the causal sentence. Thus, the first model acts like a barrier and filters noncausal sentences out. These sentences may have either a cause, an effect, none of them, but not both. To simplify the model training, we hypothesized that cause-effect pairs only occur in the same sentence and we removed all sentences with less than 6 words owing to a lack

of context. For this reason, we operated on a sentence level and not at the tweet level. Additional challenges in our setting were that *causes* and *effects* could be multiword entities and the language used on Twitter is nonstandard with frequent slang and misspelled words.

Causal Sentence Detection

The identification of causal sentences is a binary classification task. The pretrained language model *BERTweet* served as a foundation for the model architecture capable of handling the nonstandard nature of Twitter data [32]. A feed-forward network is built on top of the *BERTweet* [32] architecture consisting of 2 fully connected layers with dropout layers with a probability of 0.3, finalized by a softmax layer, which translates the model predictions into probabilities (Figure 2). To adjust for the class imbalance in the labeled data, class weights were included as parameters in the categorical cross-entropy loss function to penalize mispredictions for causal sentences strongly. Initially, labelled data were stratified, and 10% of it was kept as test set. The remaining 90% of the samples were further separated into training and validation sets with 80:20 split.

Figure 2. Model architecture for causal sentence detection. FCLL: fully connected linear layer; p: probability of an element to be zeroed.



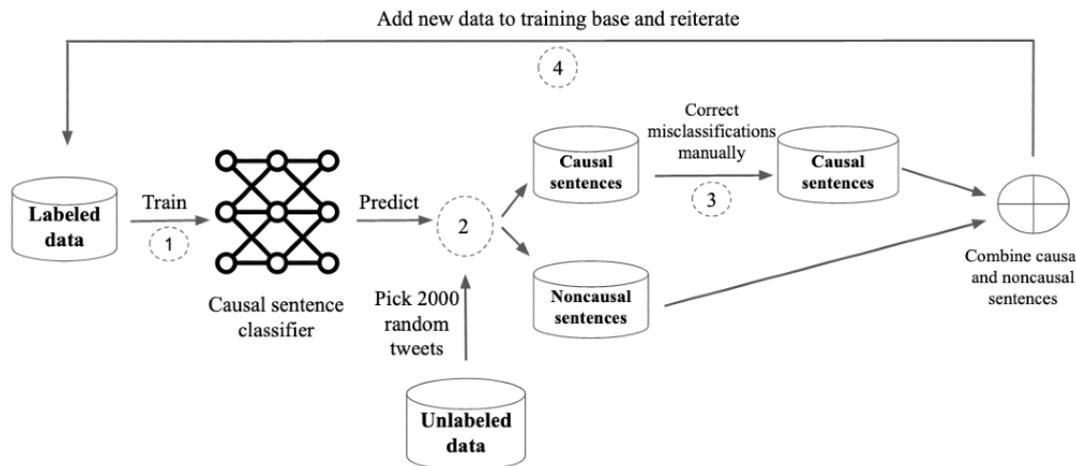
Data Augmentation Through Active Learning

Data imbalance on the one hand and the limited number of positive training examples for each cause-effect pair on the other hand (as causes and effects could potentially be related to any concept in the diabetes domain) drove us to adopt an active learning approach to increase the training data. Active learning is a sample selection approach aiming to minimize the annotation cost while maximizing the performance of machine learning-based models [44]. It has been widely applied on textual data [45,46]. The training data were increased in several iterations, as illustrated in Figure 3.

The first iteration started by training the causal sentence classifier on sentences from the 5000 tweets. The trained classifier was then applied on 2000 randomly selected unlabeled

tweets, which were preprocessed and split into sentences, resulting in a set of causal sentences and a set of noncausal sentences. The sentences predicted as causal sentences were examined manually, and possible misclassifications were corrected to ensure clean positive training samples. The noncausal sentence set remained untouched. As a consequence, potential misclassifications remained in the noncausal sentence set, which should then be considered noisy. Both the causal and noncausal sentence set were then combined and added as new training data to the already labeled data, leading to an updated training set of 7000 tweets. This process was iterated 4 times and allowed us to augment the labelled data much faster and more efficiently than that without active learning, as it enables us to focus on the few positive samples. The final training set was used to train the classification model and the cause-effect extraction model.

Figure 3. Active learning loop to augment the training set in a time-efficient fashion.



Cause-Effect Pairs

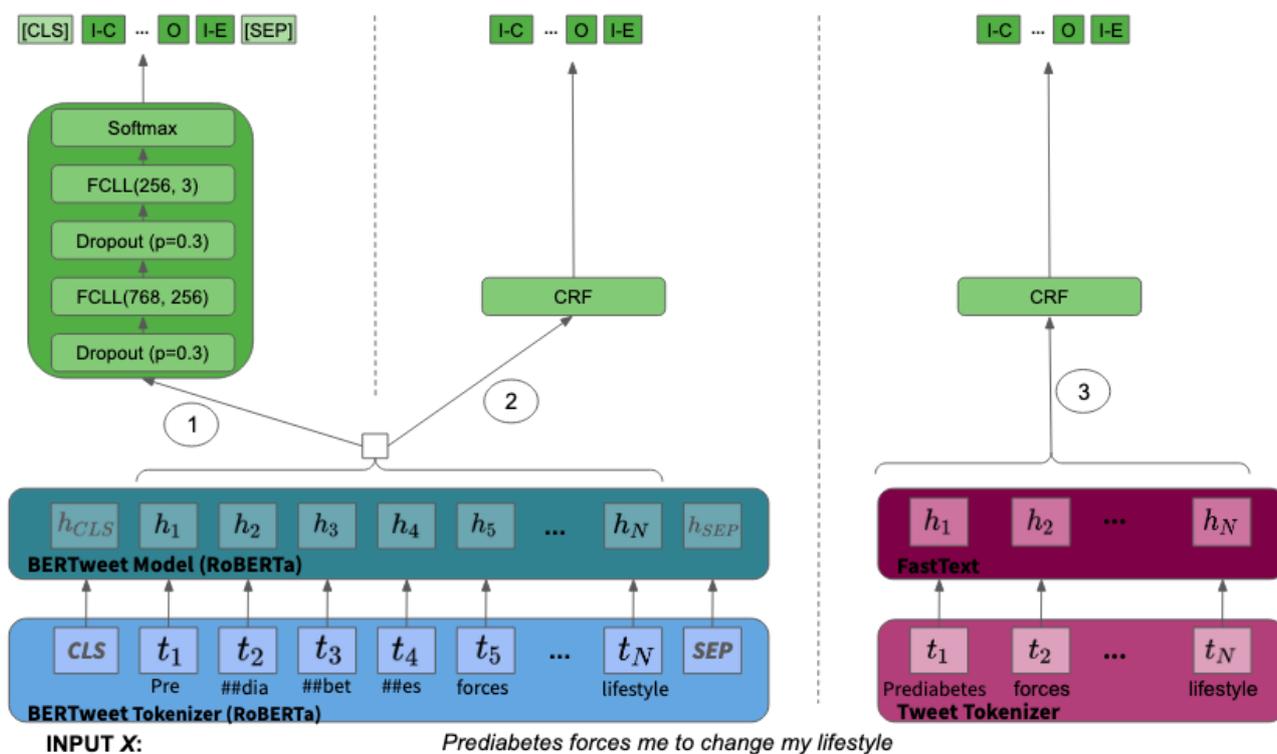
After having trained the causal sentence classifier to detect sentences with causal information, we identified the specific cause-effect pairs in the causal sentences. The identification of cause-effect pairs was casted as an event extraction or named-entity recognition task, that is, assigning a label cause or effect to a sequence of words. The manually labeled causes and effects were encoded in an IO tagging format based on the common tagging format BIO (Beginning, Inside, Outside), introduced by Ramshaw and Marcus [47]. Here, “I-C” denotes inside the cause and “I-E” inside the effect. Those 2 tags were completed by the outside tag “O,” symbolizing that the word is neither cause nor effect. The IO tagging scheme for the example sentence with cause “prediabetes” and effect “change my lifestyle” is summarized:

Sentence: Prediabetes, forces, me, to, change, my, lifestyle

IO tags: I-C, O, O, O, I-E, I-E, I-E

Note that a word can be both cause or effect depending on the context. For instance “prediabetes” in “Prediabetes forces me to change my lifestyle” takes the role of a cause, whereas in “Limited exercising may lead to prediabetes,” it is a possible effect. IO tagging was preferred over BIO tagging to simplify the model learning by reducing the number of class from 5 to 3. Moreover, the task is complex and considered open domain, as causes and effects are not restricted to 1 specific topic but can be related to any concept in our target domain (diabetes). As a consequence, the creation of a representative training set is challenging, as most cause-effect pairs occur rarely. This complexity drove us to test several model architectures; refer to Figure 4 for an overview.

Figure 4. Model architectures of cause-effect identification. CRF: conditional random field; FCLL: fully connected linear layer; p: probability of an element to be zeroed.



1. BERT_FFL: Pretrained BERTweet language model and on top, 2 feed forward layers with a dropout of 0.3, followed by a softmax layer. For the model training, the cross-entropy loss function is selected and weighted by the class weights to penalize mispredictions for causes and effects stronger.
2. WE_BERT_CRF: Single conditional random field (CRF) layer with BERTweet embeddings as features augmented by discrete features such as if the word is lowercase, digit, or the word length. CRFs are a standard statistical sequential classification method to identify entities in a text [48]. The CRF function is implemented with the python package sklearn-crfsuite [49] based on CRFsuite [50]. As parameters for the CRF function, the default algorithm “Gradient descent using the Limited Memory Broyden-Fletcher-GoldfarbShanno method” was chosen, and the coefficient for L1 and L2 regularization was 0.1.
3. FastText_CRF: Similar to WE_BERT_CRF, with the difference that BERTweet embeddings were replaced by FastText embeddings in the feature vector for each word. FastText vectors trained on similar diabetes-related tweets, which were well adapted to our use case [9].

Clustering of Causes and Effects

A large part of *causes* and *effects* can be regrouped into similar concepts (clusters) to facilitate analyses and allow effective network analyses. We chose a semisupervised, time-efficient approach in which 1000 *causes* and 1000 *effects* were randomly chosen and 2 researchers manually grouped these into clusters such as “diabetes,” “death,” “family,” and “fear,” hereinafter referred to as “parent clusters” to simplify understanding. The remaining *causes* and *effects* were then automatically compared to each element of all the clusters based on *FastText* vectors

and cosine similarity and associated with the cluster containing the most similar element. Experimentally, a similarity threshold of 0.55 was determined; if a cause/effect had a similarity smaller than this threshold for all elements, a new cluster was created for this cause/effect. These clusters were also visualized in an interactive cause-effect network, developed in D3, to enable further exploration of the cause-effect association about diabetes distress communication in social media. Python (version 3.8.8) and the deep learning framework PyTorch (version 1.8.1) were used to implement the abovementioned methods. The algorithms are open sourced under [51].

Results

The following results were obtained from 482,583 sentences, which were obtained from splitting the 562,013 personal, emotional, and nonjoke tweets into sentences, excluding questions and including only sentences with more than 5 words.

Model Training and Performance

Causal Sentences

Hyperparameters for the model training were optimized, and the best model was trained with an Adam optimizer with a learning rate of 1e-3 among [1e-2, 1e-3, 1e-4] and a scheduler with linearly decreasing learning rate with 0 warmup steps. The optimal batch size was obtained for 16 among [8,16,32], and we trained for 35 epochs with early stopping. The performances to detect causal sentences for the imbalanced data set are illustrated in Table 2 for each round of the active learning loop, with each round having been trained on more data. The highest accuracy was reached in round 4 with 71%. We applied the model of round 4 on all the remaining tweets, as it was trained on the largest training data set, including difficult causal examples missed by earlier models and is thus better at identifying complex causal sentences. The active learning strategy led us to increase the training data much quicker than that without active learning and without loss in performance. This led to a clean database of 265,328 causal sentences with the most noisy sentences removed.

Table 2. Performance measures (macro) for each round of more training data.

Round	Sentences in training set (n)	Sentences in test set (n)	Accuracy (%)	Precision (%)	Recall (%)
0	6024	837	64.5	58.0	67.4
1	7536	1047	67.7	61.2	71.6
2	8804	1223	67.7	60.3	66.3
3	10,284	1429	65.4	60.0	68.8
4	11,861	1648	71.0	61.0	67.8

Cause and Effect Detection

After having identified the causal sentences, the cause-effect models were trained to extract the specific cause-effect pairs. The active learning strategy led to an extended data set of 2118 causal sentences, that is, containing both cause and effect, of which 10% were used as a test set while the remaining 90% were further used to create a training and validation set with an 80:20 split. The performances of the different cause-effect models are listed in Table 3. The best performing model was the CRF model with BERT-embedding features (WE_BERT_CRF) with a precision, recall, and F1 score of 0.68. Surprisingly, it outperforms fine-tuning a BERT model, which is considered the gold standard of current named-entity

recognition tasks. A potential explanation for this is that BERT-based models make local decisions at every point of the sequence taking the neighboring words into account before its decision. In a situation like ours, with strong uncertainty on all elements, owing to the complexity of the task, a single CRF layer model leveraging BERT features, making global decisions using the local context of each word, maximizes the probability of the whole sequence of the decision better. Moreover, the CRF model with simpler FastText models achieved strong results as well with one reason being probably that the word embeddings were specifically trained on this diabetes corpus.

Consequently, the WE_BERT_CRF model was applied on all causal sentences leading to a data set of 96,676 sentences with the *cause* and associated *effect* predicted.

Table 3. Performance measures for each of the 4 architectures.

Models	Precision	Recall	F1 score
BERT_FFL			
I-C	0.48	0.46	0.47
I-E	0.20	0.48	0.29
O	0.91	0.77	0.83
macro	0.53	0.57	0.53
WE_BERT_CRF			
I-C	0.63	0.61	0.62
I-E	0.49	0.49	0.49
O	0.93	0.93	0.93
macro	0.68	0.68	0.68
FastText_CRF			
I-C	0.59	0.57	0.58
I-E	0.45	0.38	0.41
O	0.92	0.94	0.93
macro	0.65	0.63	0.64

Cause-Effect Description

The semisupervised clustering led to 1751 clusters. To remove noisy clusters through potential misclassifications, only clusters with a minimal number of 10 cause/effect occurrences were considered for the following analyses, resulting in 763 clusters. Note that the order of documents might affect the results, as different clusters might have been created. Please refer to [Multimedia Appendix 4](#) for an overview over the 100 largest clusters (automatically added clusters have “other” as “parent cluster”).

[Table 4](#) provides an overview over the largest clusters, containing either cause or effect. [Table 5](#) provides the most frequent cause-effect associations, excluding the largest cluster “diabetes,” as it will be studied separately. The cluster “diabetes” is the largest one with 66,775 occurrences of “diabetes” as either cause or effect (eg, diabetes, #diabetes, diabetes mellitus) followed by “death” with 16,989 (eg, passed away, killed, died, suicide) and “insulin” (eg, insulin, insulin hormone) with 14,148 occurrences. From the 30 largest clusters, 6 refer to nutrition, 4 to diabetes, and 3 to each of insulin, emotions, and the health care system. The most frequent cause-effect is “unable to afford insulin,” which causes “death” expressed in 1246 cases,

followed by “insulin” causing “death” with 1156 cases and “type 1 diabetes” causing “fear” with 1054 cases.

The largest cluster “diabetes” mainly occurs as a cause and its 10 most frequent effects are death (n=7446), fear (n=4836), sick (n=2799), neuropathy (n=2477), hypoglycemia (n=2062), anger (n=1908), suffer (n=1808), insulin (n=1605), overweight (n=1506), and reduce weight (n=1487). From the 30 most numerous effects for “diabetes,” 6 were related to “nutrition” and 5 to “complications and comorbidities” and 3 to each of “diabetes distress,” “emotions,” and “health care system.”

The interactive visualization in D3 with filter options is published in [52]. [Figure 5](#) provides an example graph of this visualization showing only cause-effect relationships with at least 250 occurrences to ensure readability. It is striking that “death” seems to play such a central role as *effect* with various causes (unable to afford insulin, rationing insulin, finance, insulin, type 1 diabetes, overweight) pointing at it. Other central nodes are type 1 diabetes acting as cause for insulin pump, insulin, hypoglycemia (hypo), sickness, finance, and anger, and fear emotions, where the latter has the strongest association, or the node “insulin” mostly relating as cause for sickness, medication, finance, death, or hypoglycemia and fear and anger.

Table 4. The most frequent clusters (causes and effects) with the number of occurrences.

Parent cluster	Cluster	Value (n)
Diabetes	diabetes	66,775
Death	death	16,989
Insulin	insulin	14,148
Diabetes	type 1 diabetes	11,693
Emotions	fear	10,160
Glycemic variability	hypoglycemia	9547
Symptoms	sick	6549
Nutrition	overweight	5186
Diabetes	type 2 diabetes	4909
Complications and comorbidities	neuropathy	4481
Health care system	medication	4389
Diabetes Technology	insulin pump	4307
Nutrition	nutrition	4230
Emotions	anger	4149
Health	oral glucose tolerance test	4053
Blood pressure	hypertension	3782
Health care system	finance	3767
Nutrition	reduce weight	3589
Insulin	unable to afford insulin	3381
Nutrition	diet	3325
Emotions	sadness	3153
Glycemic variability	hyperglycemia	3144
Diabetes	suffer	3132
Diabetes Distress	depression	2810
Health care system	hospital	2721
Diabetes Distress	stress	2681
Nutrition	sugar	2369
Nutrition	fasting	2363
Insulin	rationing insulin	2244
Health	gestational diabetes	2076

Table 5. The most frequent cause-effect relationships excluding the cluster “diabetes” with the number of occurrences.

Cause	Effect	Value (n)
unable to afford insulin	death	1246
insulin	death	1156
type 1 diabetes	fear	1054
type 1 diabetes	death	999
rationing insulin	death	805
type 1 diabetes	insulin	751
oral glucose tolerance test	sick	584
type 1 diabetes	hypoglycemia	578
insulin	hypo	545
insulin	fear	534
type 1 diabetes	insulin pump	436
finance	death	423
type 1 diabetes	sick	400
insulin	sick	385
insulin	finance	367
type 1 diabetes	anger	356
insulin	medication	305
insulin	anger	296
oral glucose tolerance test	fear	293
type 2 diabetes	death	293
type 2 diabetes	fear	290
hypertension	death	286
overweight	death	280
type 1 diabetes	finance	277
hypoglycemia	insulin	272
hypoglycemia	sick	263
affordable insulin	death	262
insulin	insulin pump	255
complications	death	248
insulin	sadness	240

the first paper investigating both explicit and implicit cause-effect relationships on diabetes-related Twitter data.

Strengths and Limitations

This study demonstrates various strengths. First, by leveraging powerful language models, we were able to identify a large number of tweets containing *cause-effect* relationships, which enabled us to detect cause-effect associations in 20% (96,676/482,583) of the sentences, contrary to other approaches that were able to identify causality in less than 2% of tweets [14]. Second, contrary to most previous work, we tackled both explicit and implicit *causal relationships*, an additional explanation for the higher number of *cause-effect* associations we obtained, compared to other studies focusing only on explicit associations [14]. Third, relying fully on automatic machine learning algorithms avoided us from defining manually crafted patterns to detect causal associations. Fourth, operating on social media data that are expressed spontaneously and in real time offers the opportunity to gain knowledge from an alternative data source and, in particular, from a patient's perspective, which might complement traditional epidemiological data sources. Lastly, the data-driven approach to identify cause-effect relationships, as reported from Twitter users, can be used in the next step to generate new hypotheses that can be tested in a more clinical setting, for example, in a clinical trial.

A strong limitation is that *cause-effect* relations are expressed in tweets and this cannot be used for causal inference as the Twitter data source is uncertain and the information shared can be an opinion or an observation. Another shortcoming is that the performance of our algorithms to detect *cause-effect* pairs is not perfect. However, the overall process and the vast amount of data minimize this issue. The lack of recall is counterbalanced by the sheer amount of data, and the lack of precision is counterbalanced by the clustering approach in which nonfrequent causes or effects are discarded [56]. Labeling causes and effects in a data set is a highly complicated task, and we would like to emphasize that mislabeling in the data set may occur. Here, the actual prevalence of causal sentences is lower,

as we wanted to catch as many causal sentences as possible, which led to also having captured some noncausal sentences. Enhancing data quality certainly is a strong point to address to further improve performance. The causal association structures learnt by the model from the training set might not generalize completely when applied on the large amount of Twitter data. Besides, the active learning strategy certainly added noise to the model, as only positive samples were corrected, which could be improved in future investigations. Moreover, we would like to highlight that the diabetes-related information shared on Twitter may not be representative for all people with diabetes. For instance, we observed a bigger cluster of causes/effects related to type 1 diabetes compared to that related to type 2 diabetes, which is contrary to that in the real world [57]. A potential explanation for that is the age distribution of Twitter users [58]. However, owing to the large number of tweets analyzed, a significant variability in the tweets could be observed.

Conclusion

In this work, we developed an innovative methodology to identify possible cause-effect relationships among diabetes-related tweets. This task was challenging owing to addressing both explicit and implicit causality, multiword entities, the fact that a word could be both cause or effect, the open domain of causes and effects, the biases occurring during labeling of causality, and the relatively small data set for this complex task. We overcame these challenges by augmenting the small data set via an active learning loop. The feasibility of our approach was demonstrated using modern BERT-based architectures in the preprocessing and causal sentence detection. A combination of BERT features and CRF layer were leveraged to extract causes and effects in diabetes-related tweets, which were then aggregated to clusters in a semisupervised approach. The visualization of the cause-effect network based on Twitter data can deepen our understanding of diabetes, in a way of directly capturing patient-reported outcomes from a causal perspective. The fear of death owing to the inability to afford insulin was the main concern expressed.

Acknowledgments

This work was supported by the MSDAVENIR Foundation, the French Speaking Diabetes Society, and the Luxembourg Institute of Health. These study sponsors had no role in the design or the interpretation of the results of this study. AA, FO, and TC are supported by Epiconcept Company. Epiconcept was involved in the data collection and writing of the report. No study sponsor influenced the decision to submit the paper for publication.

Conflicts of Interest

None declared.

Multimedia Appendix 1

List of diabetes-related keywords for the Twitter application programming interface tweet extraction.
[\[PDF File \(Adobe PDF File\), 47 KB-Multimedia Appendix 1\]](#)

Multimedia Appendix 2

Preprocessing pipeline.
[\[PDF File \(Adobe PDF File\), 64 KB-Multimedia Appendix 2\]](#)

Multimedia Appendix 3

Annotation guidelines.

[\[PDF File \(Adobe PDF File\), 120 KB-Multimedia Appendix 3\]](#)

Multimedia Appendix 4

Most frequent clusters.

[\[PDF File \(Adobe PDF File\), 75 KB-Multimedia Appendix 4\]](#)

References

1. Fisher L, Hessler DM, Polonsky WH, Mullan J. When is diabetes distress clinically meaningful?: establishing cut points for the Diabetes Distress Scale. *Diabetes Care* 2012 Feb;35(2):259-264. [doi: [10.2337/dc11-1572](https://doi.org/10.2337/dc11-1572)] [Medline: [22228744](https://pubmed.ncbi.nlm.nih.gov/22228744/)]
2. Fisher L, Polonsky WH, Hessler DM, Masharani U, Blumer I, Peters AL, et al. Understanding the sources of diabetes distress in adults with type 1 diabetes. *J Diabetes Complications* 2015;29(4):572-577. [doi: [10.1016/j.jdiacomp.2015.01.012](https://doi.org/10.1016/j.jdiacomp.2015.01.012)] [Medline: [25765489](https://pubmed.ncbi.nlm.nih.gov/25765489/)]
3. Cocco E, Lazarus S, Joseph J, Wyne K, Drossos T, Phillipson L, et al. Emotional Regulation and Diabetes Distress in Adults With Type 1 and Type 2 Diabetes. *Diabetes Care* 2021 Jan;44(1):20-25. [doi: [10.2337/dc20-1059](https://doi.org/10.2337/dc20-1059)] [Medline: [33444157](https://pubmed.ncbi.nlm.nih.gov/33444157/)]
4. Carper MM, Traeger L, Gonzalez JS, Wexler DJ, Psaros C, Safren SA. The differential associations of depression and diabetes distress with quality of life domains in type 2 diabetes. *J Behav Med* 2014 Jun;37(3):501-510. [doi: [10.1007/s10865-013-9505-x](https://doi.org/10.1007/s10865-013-9505-x)] [Medline: [23515932](https://pubmed.ncbi.nlm.nih.gov/23515932/)]
5. Cummings DM, Lutes L, Littlewood K, DiNatale E, Hambidge B, Schulman K, et al. Regimen-Related Distress, Medication Adherence, and Glycemic Control in Rural African American Women With Type 2 Diabetes Mellitus. *Ann Pharmacother* 2014 Aug;48(8):970-977. [doi: [10.1177/1060028014536532](https://doi.org/10.1177/1060028014536532)] [Medline: [24904183](https://pubmed.ncbi.nlm.nih.gov/24904183/)]
6. Fisher L, Mullan J, Skaff M, Glasgow R, Areal P, Hessler D. Predicting diabetes distress in patients with Type 2 diabetes: a longitudinal study. *Diabet Med* 2009 Jun;26(6):622-627. [doi: [10.1111/j.1464-5491.2009.02730.x](https://doi.org/10.1111/j.1464-5491.2009.02730.x)] [Medline: [19538238](https://pubmed.ncbi.nlm.nih.gov/19538238/)]
7. Pandit AU, Bailey SC, Curtis LM, Seligman HK, Davis TC, Parker RM, et al. Disease-related distress, self-care and clinical outcomes among low-income patients with diabetes. *J Epidemiol Community Health* 2014 Jun;68(6):557-564. [doi: [10.1136/jech-2013-203063](https://doi.org/10.1136/jech-2013-203063)] [Medline: [24489044](https://pubmed.ncbi.nlm.nih.gov/24489044/)]
8. Schmidt CB, van Loon BJP, Vergouwen ACM, Snoek FJ, Honig A. Systematic review and meta-analysis of psychological interventions in people with diabetes and elevated diabetes-distress. *Diabet. Med* 2018 Jun 30;35(9):1157-1172. [doi: [10.1111/dme.13709](https://doi.org/10.1111/dme.13709)] [Medline: [29896760](https://pubmed.ncbi.nlm.nih.gov/29896760/)]
9. Ahne A, Orchard F, Tannier X, Perchoux C, Balkau B, Pagoto S, et al. Insulin pricing and other major diabetes-related concerns in the USA: a study of 46 407 tweets between 2017 and 2019. *BMJ Open Diabetes Res Care* 2020 Jun;8(1):e001190 [FREE Full text] [doi: [10.1136/bmjdr-2020-001190](https://doi.org/10.1136/bmjdr-2020-001190)] [Medline: [32503810](https://pubmed.ncbi.nlm.nih.gov/32503810/)]
10. Balkhi AM, Reid AM, McNamara JP, Geffken GR. The diabetes online community: the importance of forum use in parents of children with type 1 diabetes. *Pediatr Diabetes* 2014 Sep 25;15(6):408-415. [doi: [10.1111/vedi.12110](https://doi.org/10.1111/vedi.12110)] [Medline: [24372986](https://pubmed.ncbi.nlm.nih.gov/24372986/)]
11. Yang J, Han S, Poon J. A survey on extraction of causal relations from natural language text. *ArXiv*. 2021 Nov 01. URL: <http://arxiv.org/abs/2101.06426> [accessed 2022-05-29]
12. Doan S, Ritchart A, Perry N, Chaparro JD, Conway M. How Do You #relax When You're #stressed? A Content Analysis and Infodemiology Study of Stress-Related Tweets. *JMIR Public Health Surveill* 2017 Jun 13;3(2):e35 [FREE Full text] [doi: [10.2196/publichealth.5939](https://doi.org/10.2196/publichealth.5939)] [Medline: [28611016](https://pubmed.ncbi.nlm.nih.gov/28611016/)]
13. Cocos A, Fiks AG, Masino AJ. Deep learning for pharmacovigilance: recurrent neural network architectures for labeling adverse drug reactions in Twitter posts. *J Am Med Inform Assoc* 2017 Jul 01;24(4):813-821. [doi: [10.1093/jamia/ocw180](https://doi.org/10.1093/jamia/ocw180)] [Medline: [28339747](https://pubmed.ncbi.nlm.nih.gov/28339747/)]
14. Doan S, Yang EW, Tilak SS, Li PW, Zisook DS, Torii M. Extracting health-related causality from twitter messages using natural language processing. *BMC Med Inform Decis Mak* 2019 Apr 04;19(Suppl 3):79 [FREE Full text] [doi: [10.1186/s12911-019-0785-0](https://doi.org/10.1186/s12911-019-0785-0)] [Medline: [30943954](https://pubmed.ncbi.nlm.nih.gov/30943954/)]
15. Khoo C, Chan S, Niu Y. Extracting causal knowledge from a medical database using graphical patterns. 2000 Presented at: Proceedings of the 38th Annual Meeting on Association for Computational Linguistics; October; Hong Kong p. 336-343. [doi: [10.3115/1075218.1075261](https://doi.org/10.3115/1075218.1075261)]
16. Kayesh H, Islam M, Wang J. On event causality detection in tweets. *ArXiv*. 2019. URL: <http://arxiv.org/abs/1901.03526> [accessed 2022-01-31]
17. Khoo C, Chan S, Niu Y. The many facets of the cause-effect relation. In: *The Semantics of Relationships*. Dordrecht: Springer; 2002.
18. Chowdhury GG. Natural language processing. *Ann Rev Info Sci Tech* 2005 Jan 31;37(1):51-89. [doi: [10.1002/aris.1440370103](https://doi.org/10.1002/aris.1440370103)]
19. El Naqa I, Murphy M. What is machine learning? In: *Machine Learning in Radiation Oncology*. Cham: Springer International Publishing; 2015.

20. Xu Y, Mou L, Li G, Chen Y, Peng H, Jin Z. Classifying relations via long short term memory networks along shortest dependency paths. 2015 Presented at: Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing; September; Lisbon p. 1785-1794. [doi: [10.18653/v1/d15-1206](https://doi.org/10.18653/v1/d15-1206)]
21. Wang L, Cao Z, Melo GD, Liu Z. Relation classification via multi-level attention CNNs. 2016 Presented at: Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics; August; Berlin p. 1298-1307. [doi: [10.18653/v1/P16-1123](https://doi.org/10.18653/v1/P16-1123)]
22. Ponti E, Korhonen A. Event-related features in feedforward neural networks contribute to identifying causal relations in discourse. In: 2017 Presented at: LSDSem 2017 - 2nd Workshop on Linking Models of Lexical, Sentential and Discourse-Level Semantics, Proceedings of the Workshop; April; Valencia, Spain p. 25-30. [doi: <https://doi.org/10.17863/CAM.9725>]
23. Khetan V, Ramnani R, Anand M, Sengupta S, Fano A. Causal BERT: language models for causality detection between events expressed in text. In: Arai K, editor. Intelligent Computing. Cham: Springer International Publishing; Jul 2021:965-980.
24. Devlin J, Chang M, Lee K, Toutanova K. BERT: pretraining of deep bidirectional transformers for language understanding. ArXiv. 2018. URL: <http://arxiv.org/abs/1810.04805> [accessed 2021-05-31]
25. Vaswani A, Shazeer N, Parmar N. Attention is all you need. ArXiv. 2017. URL: <http://arxiv.org/abs/1706.03762> [accessed 2021-09-09]
26. Gurulingappa H, Rajput AM, Roberts A, Fluck J, Hofmann-Apitius M, Toldo L. Development of a benchmark corpus to support the automatic extraction of drug-related adverse effects from medical case reports. J Biomed Inform 2012 Oct;45(5):885-892 [FREE Full text] [doi: [10.1016/j.jbi.2012.04.008](https://doi.org/10.1016/j.jbi.2012.04.008)] [Medline: [22554702](https://pubmed.ncbi.nlm.nih.gov/22554702/)]
27. Weiss K, Khoshgoftaar TM, Wang D. A survey of transfer learning. J Big Data 2016 May 28;3(1):1-40. [doi: [10.1186/s40537-016-0043-6](https://doi.org/10.1186/s40537-016-0043-6)]
28. Fagherazzi G. Challenges and perspectives for the future of diabetes epidemiology in the era of digital health and artificial intelligence. Diabetes Epidemiology and Management 2021 Jan;1:100004. [doi: [10.1016/j.deman.2021.100004](https://doi.org/10.1016/j.deman.2021.100004)]
29. Fagherazzi G, Ahne A, Guillot C, Riveline J, Bonnet F, Mebarki A, et al. Étude mondiale de la détresse liée au diabète : le potentiel du réseau social Twitter pour la recherche médicale. Revue d'Épidémiologie et de Santé Publique 2018 Jun;66:S197-S198. [doi: [10.1016/j.respe.2018.04.002](https://doi.org/10.1016/j.respe.2018.04.002)]
30. Twitter privacy policy. Twitter. 2021. URL: <https://twitter.com/en/privacy> [accessed 2021-07-06]
31. Liu Y, Mei Q, Hanauer DA, Zheng K, Lee JM. Use of Social Media in the Diabetes Community: An Exploratory Analysis of Diabetes-Related Tweets. JMIR Diabetes 2016 Nov 07;1(2):e4 [FREE Full text] [doi: [10.2196/diabetes.6256](https://doi.org/10.2196/diabetes.6256)] [Medline: [30291053](https://pubmed.ncbi.nlm.nih.gov/30291053/)]
32. Nguyen D, Vu T, Nguyen A. BERTweet: A pretrained language model for English tweets. ArXiv. 2020. URL: <http://arxiv.org/abs/2005.10200> [accessed 2021-12-12]
33. Wolf T, Debut L, Sanh V. HuggingFace's transformers: state-of-the-art natural language processing. ArXiv. 2020. URL: <http://arxiv.org/abs/1910.03771> [accessed 2021-12-12]
34. Parrott W. Emotions in Social Psychology: Essential Readings. Hove, East Sussex, United Kingdom: Psychology Press; 2001.
35. Polonsky W, Anderson B, Lohrer P, Welch G, Jacobson AM, Aponte JE, et al. Assessment of diabetes-related distress. Diabetes Care 1995 Jun;18(6):754-760. [doi: [10.2337/diacare.18.6.754](https://doi.org/10.2337/diacare.18.6.754)] [Medline: [7555499](https://pubmed.ncbi.nlm.nih.gov/7555499/)]
36. Polonsky WH, Fisher L, Earles J, Dudl RJ, Lees J, Mullan J, et al. Assessing psychosocial distress in diabetes: development of the diabetes distress scale. Diabetes Care 2005 Mar;28(3):626-631. [doi: [10.2337/diacare.28.3.626](https://doi.org/10.2337/diacare.28.3.626)] [Medline: [15735199](https://pubmed.ncbi.nlm.nih.gov/15735199/)]
37. Beguerisse-Díaz M, McLennan AK, Garduño-Hernández G, Barahona M, Ulijaszek SJ. The 'who' and 'what' of #diabetes on Twitter. Digit Health 2017;3:2055207616688841 [FREE Full text] [doi: [10.1177/2055207616688841](https://doi.org/10.1177/2055207616688841)] [Medline: [29942579](https://pubmed.ncbi.nlm.nih.gov/29942579/)]
38. Johnsen JAK, Eggesvik TB, Rørvik TH, Hanssen MW, Wynn R, Kummervold PE. Differences in Emotional and Pain-Related Language in Tweets About Dentists and Medical Doctors: Text Analysis of Twitter Content. JMIR Public Health Surveill 2019 Feb 06;5(1):e10432 [FREE Full text] [doi: [10.2196/10432](https://doi.org/10.2196/10432)] [Medline: [30724738](https://pubmed.ncbi.nlm.nih.gov/30724738/)]
39. Liu Y, Ott M, Goyal N, Du J, Joshi M, Chen D, et al. RoBERTa: A robustly optimized BERT pretraining approach. Arxiv. 2019. URL: <https://arxiv.org/abs/1907.11692> [accessed 2022-06-27] [doi: [10.48550/ARXIV.1907.11692](https://doi.org/10.48550/ARXIV.1907.11692)]
40. Princeton University. About WordNet. 2010. URL: <https://wordnet.princeton.edu/> [accessed 2019-04-10]
41. Cohen J. A Coefficient of Agreement for Nominal Scales. Educational and Psychological Measurement 2016 Jul 02;20(1):37-46. [doi: [10.1177/001316446002000104](https://doi.org/10.1177/001316446002000104)]
42. Altman D. Practical Statistics for Medical Research. United Kingdom: Chapman & Hall; 1990.
43. Landis JR, Koch GG. The Measurement of Observer Agreement for Categorical Data. Biometrics 1977 Mar;33(1):159. [doi: [10.2307/2529310](https://doi.org/10.2307/2529310)]
44. Settles B. Active Learning Literature Survey. 2010. URL: http://burrsettles.com/pub/settles.activelearning.pdf?source=post_page [accessed 2021-08-08]
45. Zhang Y, Lease M, Wallace B. Active discriminative text representation learning. ArXiv. 2016. URL: <http://arxiv.org/abs/1606.04212> [accessed 2021-09-12]

46. Tong S, Koller D. Support vector machine active learning with applications to text classification. *J Mach Learn Res* 2001;2:45-66 [FREE Full text] [doi: [10.1162/153244302760185243](https://doi.org/10.1162/153244302760185243)]
47. Ramshaw L, Marcus M. Text chunking using transformation-based learning. ArXiv. 1999. URL: <http://arxiv.org/abs/cmp-lg/9505040> [accessed 2021-09-07]
48. Lafferty J, McCallum A, Pereira F. Conditional random fields: probabilistic models for segmenting and labeling sequence data. University of Pennsylvania. 2001. URL: https://repository.upenn.edu/cgi/viewcontent.cgi?article=1162&context=cis_papers [accessed 2022-06-27]
49. Korobov M. sklearn-crfsuite. URL: <https://sklearn-crfsuite.readthedocs.io/en/latest/index.html> [accessed 2021-09-14]
50. Okazaki N. A fast implementation of conditional random fields (CRFs). CRFsuite. 2007. URL: <http://www.chokkan.org/software/crfsuite/> [accessed 2021-10-11]
51. WDDS/Causal-associations-diabetes-twitter. GitHub. URL: <https://github.com/WDDS/Causal-associations-diabetes-twitter/> [accessed 2022-06-27]
52. Cause and effect associations in diabetes-related tweets. Adahne. URL: <https://observablehq.com/@adahne/cause-and-effect-associations-in-diabetes-related-tweets> [accessed 2022-06-27]
53. Bollegala D, Maskell S, Sloane R, Hajne J, Pirmohamed M. Causality Patterns for Detecting Adverse Drug Reactions From Social Media: Text Mining Approach. *JMIR Public Health Surveill* 2018 May 09;4(2):e51 [FREE Full text] [doi: [10.2196/publichealth.8214](https://doi.org/10.2196/publichealth.8214)] [Medline: [29743155](https://pubmed.ncbi.nlm.nih.gov/29743155/)]
54. Dasgupta T, Saha R, Dey L, Naskar A. Automatic Extraction of Causal Relations from Text using Linguistically Informed Deep Neural Networks. 2018 Presented at: Proceedings of the 19th Annual SIGdial Meeting on Discourse and Dialogue; April; Melbourne p. 306-316 URL: <https://www.superlectures.com/sigdial2018/automatic-extraction-of-causal-relations-from-text-using-linguistically-informed-deep-neural-networks> [doi: [10.18653/v1/w18-5035](https://doi.org/10.18653/v1/w18-5035)]
55. Khetan V, Rizvi M, Huber J, Bartusiak P, Sacaleanu B, Fano A. MIMICause: Defining, identifying and predicting types of causal relationships between biomedical concepts from clinical notes. ArXiv. URL: <http://arxiv.org/abs/2110.07090> [accessed 2021-12-12]
56. Tannier X. NLP-driven Data Journalism: Time-Aware Mining and Visualization of International Alliances. 2016. URL: <https://hal.archives-ouvertes.fr/hal-02407145/document> [accessed 2022-06-27]
57. International Diabetes Federation Diabetes Atlas, 9th edn. 2019. URL: <https://www.diabetesatlas.org> [accessed 2021-12-09]
58. Percentage of US adults who use Twitter as of February 2021, by age group. Statista. 2021. URL: <https://www.statista.com/statistics/265647/share-of-us-internet-users-who-use-twitter-by-age-group/> [accessed 2021-10-23]

Abbreviations

- BERT:** Bidirectional Encoder Representations from Transformers
BIO: Beginning, Inside, Outside
CRF: conditional random field

Edited by C Lovis; submitted 10.02.22; peer-reviewed by A Hulman, C Pollack; comments to author 21.04.22; revised version received 17.05.22; accepted 04.06.22; published 19.07.22

Please cite as:

Ahne A, Khetan V, Tannier X, Rizvi MIH, Czernichow T, Orchard F, Bour C, Fano A, Fagherazzi G
Extraction of Explicit and Implicit Cause-Effect Relationships in Patient-Reported Diabetes-Related Tweets From 2017 to 2021: Deep Learning Approach
JMIR Med Inform 2022;10(7):e37201
URL: <https://medinform.jmir.org/2022/7/e37201>
doi: [10.2196/37201](https://doi.org/10.2196/37201)
PMID:

©Adrian Ahne, Vivek Khetan, Xavier Tannier, Md Imbesat Hassan Rizvi, Thomas Czernichow, Francisco Orchard, Charline Bour, Andrew Fano, Guy Fagherazzi. Originally published in *JMIR Medical Informatics* (<https://medinform.jmir.org>), 19.07.2022. This is an open-access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work, first published in *JMIR Medical Informatics*, is properly cited. The complete bibliographic information, a link to the original publication on <https://medinform.jmir.org/>, as well as this copyright and license information must be included.