

Triage and diagnosis of COVID-19 from medical social media

Abul Hasan, Mark Levene, David Weston, Renate Fromson, Nicolas Koslover,
Tamara Levene

Submitted to: Journal of Medical Internet Research
on: May 13, 2021

Disclaimer: © The authors. All rights reserved. This is a privileged document currently under peer-review/community review. Authors have provided JMIR Publications with an exclusive license to publish this preprint on its website for review purposes only. While the final peer-reviewed paper may be licensed under a CC BY license on publication, at this stage authors and publisher expressly prohibit redistribution of this draft paper other than for review purposes.

Table of Contents

Original Manuscript..... 5

Supplementary Files..... 25

 Figures 26

 Figure 1..... 27

 Figure 2..... 28

 Figure 3..... 29

 Figure 4..... 30

 Figure 5..... 31

 Figure 6..... 32

Triage and diagnosis of COVID-19 from medical social media

Abul Hasan¹ MS; Mark Levene¹ PhD; David Weston¹ PhD; Renate Fromson² MBBS; Nicolas Koslover² MBBS; Tamara Levene² MBBS

¹Birkbeck, University of London London GB

²Barnet General Hospital Wellhouse Lane, London EN5 3DJ, United Kingdom London GB

Corresponding Author:

Abul Hasan MS

Birkbeck, University of London

Malet street, bloomsbury

London

GB

Abstract

Background: The COVID-19 pandemic has created a pressing need for integrating information from disparate sources, in order to assist decision makers. Social media is important in this respect, however, to make sense of the textual information it provides and be able to automate the processing of large amounts of data, natural language processing methods are needed. Social media posts are often noisy, yet they may provide valuable insights regarding the severity and prevalence of the disease in the population. In particular, machine learning techniques for triage and diagnosis could allow for a better understanding of what social media may offer in this respect.

Objective: This study aims to develop an end-to-end natural language processing pipeline for triage and diagnosis of COVID-19 from patient-authored social media posts, in order to provide researchers and other interested parties with additional information on the symptoms, severity and prevalence of the disease.

Methods: The text processing pipeline first extracts COVID-19 symptoms and related concepts such as severity, duration, negations, and body parts from patients' posts using conditional random fields. An unsupervised rule-based algorithm is then applied to establish relations between concepts in the next step of the pipeline. The extracted concepts and relations are subsequently used to construct two different vector representations of each post. These vectors are applied separately to build support vector machine learning models to triage patients into three categories and diagnose them for COVID-19.

Results: We report that Macro- and Micro-averaged $F_{1\setminus}$ scores in the range of 71-96% and 61-87%, respectively, for the triage and diagnosis of COVID-19, when the models are trained on human labelled data. Our experimental results indicate that similar performance can be achieved when the models are trained using predicted labels from concept extraction and rule-based classifiers, thus yielding end-to-end machine learning. Also, we highlight important features uncovered by our diagnostic machine learning models and compare them with the most frequent symptoms revealed in another COVID-19 dataset. In particular, we found that the most important features are not always the most frequent ones.

Conclusions: Our preliminary results show that it is possible to automatically triage and diagnose patients for COVID-19 from natural language narratives using a machine learning pipeline, in order to provide additional information on the severity and prevalence of the disease through the eyes of social media.

(JMIR Preprints 13/05/2021:30397)

DOI: <https://doi.org/10.2196/preprints.30397>

Preprint Settings

1) Would you like to publish your submitted manuscript as preprint?

✓ **Please make my preprint PDF available to anyone at any time (recommended).**

Please make my preprint PDF available only to logged-in users; I understand that my title and abstract will remain visible to all users.

Only make the preprint title and abstract visible.

No, I do not wish to publish my submitted manuscript as a preprint.

2) If accepted for publication in a JMIR journal, would you like the PDF to be visible to the public?

✓ **Yes, please make my accepted manuscript PDF available to anyone at any time (Recommended).**

Yes, but please make my accepted manuscript PDF available only to logged-in users; I understand that the title and abstract will remain visible to all.

Yes, but only make the title and abstract visible (see Important note, above). I understand that if I later pay to participate in [JMIR Publications](#), my work will be made available to all.



Original Manuscript

Original Paper

Triage and diagnosis of COVID-19 from medical social media

Abul Hasan¹, MSc; Mark Levene¹, PhD; David Weston¹, PhD; Renate Fromson², MBBS; Nicolas Koslover², MBBS; Tamara Levene², MBBS

¹Department of Computer Science, Birkbeck, University of London, London WC1E 7HX, United Kingdom.

²Barnet General Hospital, Wellhouse Lane, London EN5 3DJ, United Kingdom.

Corresponding Author:

Abul Hasan, MSc

Department of Computer Science,

Birkbeck, University of London,

London WC1E 7HX,

Email: abulhasan@dcs.bbk.ac.uk

Phone: +44 (0)20 7631 8147

Abstract

Background: The COVID-19 pandemic has created a pressing need for integrating information from disparate sources, in order to assist decision makers. Social media is important in this respect, however, to make sense of the textual information it provides and be able to automate the processing of large amounts of data, natural language processing methods are needed. Social media posts are often noisy, yet they may provide valuable insights regarding the severity and prevalence of the disease in the population. In particular, machine learning techniques for triage and diagnosis could allow for a better understanding of what social media may offer in this respect.

Objective: This study aims to develop an end-to-end natural language processing pipeline for triage and diagnosis of COVID-19 from patient-authored social media posts, in order to provide researchers and other interested parties with additional information on the symptoms, severity and prevalence of the disease.

Materials and Methods: The text processing pipeline first extracts COVID-19 symptoms and related concepts such as severity, duration, negations, and body parts from patients' posts using conditional random fields. An unsupervised rule-based algorithm is then applied to establish relations between concepts in the next step of the pipeline. The extracted concepts and relations are subsequently used to construct two different vector representations of each post. These vectors are applied separately to build support vector machine learning models to triage patients into three categories and diagnose them for COVID-19.

Results: We report that Macro- and Micro-averaged F_1 scores in the range of 71-96% and 61-87%, respectively, for the triage and diagnosis of COVID-19, when the models are trained on human labelled data. Our experimental results indicate that similar performance can be achieved when the models are trained using predicted labels from concept extraction and rule-based classifiers, thus yielding end-to-end machine learning. Also, we highlight important features uncovered by our diagnostic machine learning models and compare them with the most frequent symptoms revealed in another COVID-19 dataset. In particular, we found that the most important features are not always the most frequent ones.

Conclusions: Our preliminary results show that it is possible to automatically triage and diagnose patients for COVID-19 from natural language narratives using a machine learning pipeline, in order to provide additional information on the severity and prevalence of the disease through the eyes of social media.

Key words: COVID-19, Triage and diagnosis, Medical social media, Natural Language processing, Conditional random fields, Support vector machines.

Introduction

Overview

During the ongoing Coronavirus pandemic, hospitals have been continuously at risk of being overwhelmed by the number of people developing serious illness. People have therefore been advised, as far as possible, to self-care at home for COVID-19 symptoms. Consequently, there is an urgent need to develop novel practical approaches to assist medical staff in the diagnosis and triage of patients. A variety of methods have been recently developed that involve natural language processing techniques; see for example [1-3]. In addition, social media search queries related to COVID-19 symptoms in China have been shown to be an effective predictor for the number of infections [4].

Herein, we propose an end-to-end *Natural Language Processing (NLP)* pipeline to automatically triage and diagnose COVID-19 cases from patient-authored medical social media posts. The triage may inform decision-makers about the severity of COVID-19, and the diagnosis could help in gauging the prevalence of infections in the population. A key concern is the production of a high-quality ground truth dataset on which to build our pipeline. In the following we give a brief overview of our pipeline and how we developed our dataset.

The first step in the pipeline is attained by developing an annotation application that detects and highlights COVID-19 related symptoms with their severity and duration in a social media post, henceforth collectively termed as *concepts*. During the second step relations between symptoms and other relevant concepts are also automatically identified and annotated. For example, *breathing hurts* is a symptom which is related to a body part *upper chest area*.

One author manually annotated our data with concepts and relations, allowing us to present posts highlighted with identified concepts and relations to three experts along with several questions, as shown in Figure 1. The first question asked the experts to triage a patient into one of the following three categories: *Stay at home*, *Send to a GP*, and/or *Send to hospital*. The second question asked to diagnose the likelihood of COVID-19 in a *Likert Scale* of 1 to 5 [5].

The three experts are foundation doctors working in the UK who were redeployed to work on COVID-19 wards

during the first wave of the pandemic, between March and July 2020. Their roles involved the diagnosis and management of patients with COVID-19, including patients who were particularly unwell and required either non-invasive or invasive ventilation. There were some training sessions organised for doctors working on COVID-19 wards. However, these were only provided towards the end of the first wave, as there was initially little knowledge of the virus and how to treat it. A lot of learning happened “on the go”, in which doctors would learn from previous experience and pattern recognition of a set of symptoms, that is somewhat comparable to the process of machine learning. The doctors applied a similar process to the analysis of the medical forum posts presented to them.

We also asked the doctors to indicate whether the highlighted text presented is helpful, in order to understand its usefulness when we incorporate them in the annotation interface. The annotations were found to be very helpful in as many as 85% of the posts, on average, as indicated by the doctors' answers to Question 3 in Figure 1.

The posts labelled by the doctors were then used to construct two types of predictive machine learning model using *Support Vector Machines* (SVM) [6] [7]; see Step 4 from section Methods. The *triage models* employ hierarchical binary classifiers, which consider the risk averseness or tolerance of the doctors when making the diagnosis [8]. The *diagnostic models* first calculate the probability of a patient having COVID-19 from doctors' ratings. The probabilities are then used to construct three different decision functions for classifying *COVID* and *NO_COVID* classes; these are detailed in the Problem setting section in Methods.

We trained the SVM models in two different ways, first with ground truth annotations, and second using predictions from the concept and relation extraction step described above. Predictions obtained from the concept extraction step make use of *Conditional Random Fields* (CRF) [9]; see Step 1 of the Methodology Sub Section in Methods for implementation details. Relations are obtained from these predicted concepts using an unsupervised *Rule-Based* (RB) classifier [10]; see Step 2 from section Methods.

Hi im currently the same day 27 since my syptoms started , deep breathing hurts [upper chest area][throat] which is upper chest area into throat , breathing [slightly][laboured] is slightly laboured time to time , dry cough on and off , also have major fatigue weakness took a course of Amoxcillian given by GP which made no change to me , have asthma so take my inhalers which aint making no change , never been so unwell in my life ! ! !

Question 1: Please specify recommendation from one of the options below:

- ☐ Stay at home
- ☐ Send to a GP
- ☐ Send to hospital

Question 2: How would you rate the chance of this person having COVID-19 on a range of 1 to 5?

- ☐ 1 (Very unlikely)
- ☐ 2 (Unlikely)
- ☐ 3 (Uncertain)
- ☐ 4 (Likely)
- ☐ 5 (Very likely)

Question 3: Was the highlighted text sufficient in reaching your decision?

- ☐ Yes
- ☐ No

Figure 1.

A patient-authored social media post is annotated with symptoms (light green), affected body parts (pale blue), duration (light yellow) and severities (pink). The phrases in the square brackets show relations between a symptom and a body part/duration/severity, when the distance was greater than 1. This annotated post was presented to three doctors to triage and diagnose the author of the post by answering Questions 1 and 2, respectively.

We also discuss the feature importance obtained from the constructed COVID-19 diagnostic models, and compare them with the most frequent symptoms from [3] and our dataset. We found that symptoms such as anosmia/ageusia (loss of taste and smell) rank in the top 5 most important features, whereas they do not rank in the top 5 most frequent symptoms; see Discussion. Overall, we make several contributions as follows:

1. We show that it is possible to construct machine learning models to triage and diagnose COVID-19 from patients' natural language narratives. To the best of our knowledge, no other previous work has attempted to triage or diagnose COVID-19 from social media posts.

2. We also build an end-to-end NLP pipeline by making use of automated concept and relation extraction. Our experiments show that the models built using predictions from concept and relation extraction produce similar results to those built using ground truth human concept annotation.

Related work

Respiratory diseases such as influenza were studied previously to determine the likelihood of symptom severity of the disease in clinical settings using Classification and Regression Trees [11]. Machine learning algorithms, such as decision trees, have shown promising results in detecting COVID-19 from blood test analyses [12].

The purpose of this study is to extract as much useful information as possible from social media, which is known to be noisy. In particular, posts from medical social media may provide an additional source of information regarding the symptom severity and prevalence of COVID-19 in the population. Here, we focus on features extracted from a textual source to triage and diagnose COVID-19. The studies related to our work deploy features obtained from online symptom tracker applications, telehealth visits, and structured and unstructured patient/doctors notes from Electronic Health Records (EHR). In general, the COVID-19 clinical prediction models can broadly be categorised into risk, diagnosis and prognosis models [13].

Since the start of the COVID-19 pandemic, a number of mobile app based self-reported symptom tools have emerged, to track novel symptoms [14]. The mobile application in [15] applied Logistic Regression (LR) to predict percentage of probable infected cases among the total app users in the US and UK combined.

In [16], employed a portal-based COVID-19 self-triage and self scheduling tool to segment patients into four risk categories: emergent, urgent, no-urgent and self-care. Whereas, the online telemedicine system in [17] used LR to predict low, moderate and high risk patients, by utilising demographic information, clinical symptoms, blood tests and computed tomography (CT) scan results.

In [2], developed various machine learning models to predict patients outcome from clinical, laboratory and demographic features found in EHR [18]. They reported that Gradient Boosting (XGB), Random Forest and SVM are the best performing models for predicting COVID-19 test results, and, hospital and ICU admissions for positive patients, respectively. A detailed list of clinical and laboratory features can be found in [19], where they developed predictive models for the inpatient mortality in Wuhan, using an ensemble of XGB models. Similarly, In [20], predicted mortality and critical events for patients using XGB classifiers. Finally, a critical review on various diagnostic and prognostic models of COVID-19 used in clinical settings, can be found in [13].

In [21], extracted COVID-19 symptoms from unstructured clinical notes in the EHR of patients subjected to COVID-19 PCR testing. The authors in [22] performed a statistical analysis on primary care EHR data to find longitudinal dynamics of symptoms prior to and throughout the infection. In addition, COVID-19 SignSym [23] was designed to automatically extract symptoms and related attributes from free text. Furthermore, the study in [24] utilises radiological text reports from lung CT scans to diagnose COVID-19. Similar to our approach, Lopez et al. [24] first extracted concepts using a popular medical ontology [25] and then constructed a document representation using word embeddings [26] and concept vectors [24]. However, our methodology differs from theirs with respect to the extraction of relations between concepts, and moreover, our dataset, comprising posts obtained from medical social media, is more challenging to work with, since social media posts exhibit greater heterogeneity in language than radiological text reports.

Finally, [3] published a COVID-19 symptom lexicon extracted from Twitter, which we compare our work to in the Discussion section.

Methods

Data

We collected social media posts discussing COVID-19 medical conditions from a forum called *Patient* [27].

This a public forum that was created at the onset of the Coronavirus outbreak in the UK. We obtained permission from the site administrator and collected the posts dated between April and June, 2020. After the posts were anonymised, and the duplicates were removed, we randomly selected 500 distinct posts. The first

author annotated these posts with the classes shown in Figure 2. The class labels represent symptoms and the related concepts: (i) duration, (ii) intensifier, which increases the level of symptom severity, (iii) severity, (iv) negation, which denotes presence or absence of the symptom or severity, and (v) affected body parts. We also annotated relations between a symptom and other concepts that exist at sentence level. For example, the relation between a symptom and a severity concept is denoted as (*SYM*, *SEVERITY*). The posts were then marked with concepts in different colours and the relations were placed right after the symptom in square brackets, as shown in Figure 1. Each marked post was presented to the doctors using a web application, and they were asked three questions independently; see Figure 1. We call the doctors answers to Q1 and Q2 as the COVID-19 symptom triage and diagnosis, respectively. Thus for each post we have three independent answers from three doctors, which we denote as A, B, and C, respectively; these correspond to the last three authors of the paper and have been assigned randomly.




Figure 2. Frequency distribution of annotated classes/concepts from the text are shown. We also show the percentage of each class after discounting the OTHER labels. The average number of tokens per post is 130.17(SD=97.83). Here, SYM, DURATION,

INTENSIFIER, SEVERITY, BPOC and NEGATION denote symptoms, duration, intensifiers, severity, body parts and negations, respectively.

Measurement of agreement

We measured the agreement and association between the answers (recommendations and ratings) of the three doctors to Q1 and Q2 of Figure 1, using Cohen's κ [28] and Pearson's correlation coefficient [29]; therein the strength of agreement for different ranges of κ and the interpretation of the correlation coefficient, are respectively given. The correspondence of answers is measured for the three pairs: AB, AC, and BC. We observe a *poor* agreement between BC for Q1 and the same for AC for Q2 in the case of linear weighted κ , see Table 1. We further calculated the quadratic weighted κ for both questions, which resulted in *fair* agreements for all pairs as shown in Table 1. For Q2, we adjusted the linear and quadratic weights so that we do not distinguish between neighbouring diagnoses, producing a *fair* agreement in linear adjusted weighted κ for AC of Q2. For the correlations between the pairs we observe associations from *weak* to *moderate* [29].

We did not observe a high agreement between doctors' triage and diagnosis using the standard measures of Cohen's κ and Pearson's correlation. This is likely due to the fact that COVID-19 is a novel virus, for which the doctors did not have prior experience or training before the first wave of the pandemic. Moreover, there are probable differences in risk tolerances between the doctors, which would lead to potentially different decisions and diagnoses. We note that, when we calculate the proportion of observed agreement, ρ_o , as suggested by [30] who stipulate that κ is a measure of reliability rather than agreement, we do in fact observe a high value.

Table 1. Pair-wise agreement between pairs of doctors' answers for Question 1 and 2; see Figure 1 for an example. Linear Adj., Quad., Quad. Adj., and Corr. are acronyms for Linear adjusted, Quadratic, Quadratic adjusted weights and Correlation, respectively.

Question1					Question 2					
P air	ρ_o	Lin ear	Quadr atic	Correl ation	ρ_o	Lin ear	Linear Adj.	Qu ad.	Quad. Adj.	Corr.
		κ	κ	Pearso n		κ	κ	κ	κ	Pear son
A B	0. 77	0.3 5	0.46	0.48	0. 72	0.2 1	0.38	0.3 3	0.39	0.42
A C	0. 74	0.3 0	0.34	0.39	0. 71	0.1 2	0.22	0.2 2	0.24	0.41
B C	0. 80	0.1 7	0.21	0.29	0. 74	0.2 7	0.45	0.4	0.37	0.45

Problem setting

Triage classification for Question 1

We map the doctors' recommendation from Q1 to ordinal values; the options *Stay at home*, *Send to a GP*, or *Send to hospital* are transformed to the values 1, 2, and 3, respectively. In order to combine recommendations from two or more doctors, we first take their average. This result is rounded to an integer in one of two ways, either by taking the floor or the ceiling. Considering the risk attitude prevalent among medical practitioners [8], we categorise the ceiling of the average to be *risk averse*, denoted by e.g. AB(R-a), and the floor to be *risk tolerant*, denoted by e.g. AB(R-t). Thus for each patient's post, we have in total eleven recommendations from three doctors for Q1. We construct a hierarchical classification model for each of these recommendations, where the goal is to classify a post into one of the three options.

Diagnosis classification for Question 2

To diagnose whether a patient has COVID-19 from their post, we first estimate the probability of having the disease by normalising the rating, i.e. given a rating, r , the probability of COVID-19, $Pr(COVID \vee r)$, which we term as the *ground truth probability (GTP)*, is simply:

$$Pr(COVID \vee r) = \frac{r-1}{4}$$

Given our ground truth probability estimates are discrete we investigate three decision boundaries based on a threshold value of 0.5 to classify a post as follows:

\hat{r} : if $Pr(COVID|r) \leq 0.5$, then NO_COVID else COVID.

\hat{r} : if $Pr(COVID|r) < 0.5$, then NO_COVID else COVID.

\hat{r} : if $Pr(COVID|r) < 0.5$, then NO_COVID, elseif $Pr(COVID|r) > 0.5$ then COVID.

Note \hat{r} ignores cases on the 0.5 boundary.

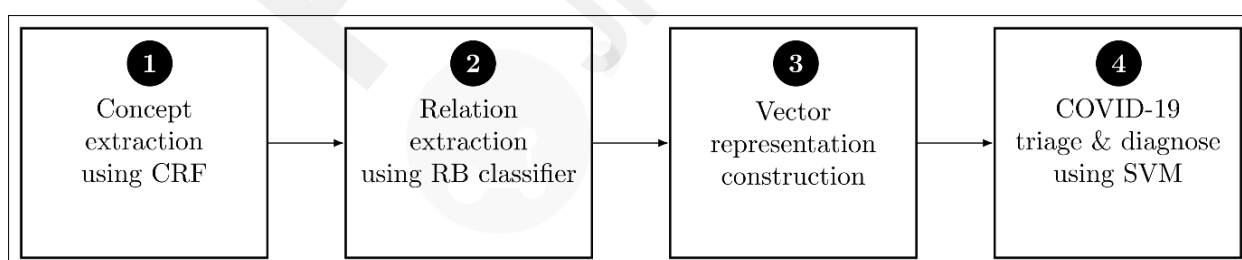


Figure 3. A block diagram of COVID-19 triage and diagnosis text processing pipeline. Here, CRF, RB classifier and SVM are acronyms for Conditional Random Fields, Rule-Based classifier and Support Vector Machine, respectively.

Methodology

A schematic of our methodology to triage and diagnose patients from their social posts is shown in Figure 3.

Here, the circles denote the steps followed in the pipeline. We now detail each of these steps.

Step 1: Concept extraction

In the first step, we pre-process each patient's post by splitting it into sentences and tokens using the GATE software [31] built-in *Natural Language Processing* (NLP) pipeline. For each token in a sentence we build discrete features that signal whether the token is a member of one of the following dictionaries: (i) Symptom, (ii) Severity, (iii) Duration, (iv) Intensifier, and (v) Negation. The dictionaries were built by analysing the posts while annotating them. We also utilise the MetaMap system [25], assuming that it contains all the necessary technical terms, to map tokens to three useful semantic categories: *Sign or Symptom*; *Disease or Syndrome*; *Body Part, Organ, or Organ Component*. Due to the assumption regarding medical terms, the system does not expect any new additional terms, and thus we are justified in extracting concepts and relations in pre-processing steps. The pre-processed text is then used to build a concept extraction module to recognise the classes, shown in Figure 2, by applying a CRF [9]. A detailed description of our CRF training methodology can be found in [32]. The extracted concepts are then used for our next step to recognise the relations between concepts.

Step 2: Relation extraction

The semantic relation between a symptom and other concepts, which we formally termed as *modifiers*, is resolved using an unsupervised RB classifier algorithm. We first filter all symptom and modifier pairs from a sentence within a predefined distance and then select the closest modifier to a symptom to construct a relation. In total, we extracted five kinds of relations as follows: (SYM, SEVERITY), (SYM, DURATION), (SYM, BPOC), (SYM, NEGATION) and (SYM, ?).

The severity modifiers are mapped to a scale of 1-5. The semantic meaning of the scale is: *very mild*, *mild*, *moderate*, *severe* and *very severe*, respectively. The duration modifiers are also mapped to real values in chunks of weeks. So, for example, *10 days* is mapped to the value 1.42.

Step 3: Vector representation

Fixed length vector representations suitable as input for SVM classifiers are built as follows.

Symptom-only vector representation

Let $\langle S_0, S_1, \dots, S_n \rangle$ be a vector of symptoms constructed from the symptom vocabulary, for our dataset the

number of unique symptom words/phrases $n=871$. To construct the vector representation for a post, we

extract the concept, SYM, and the relation (SYM, NEGATION), and set S_i to 1, 0, or -1, according to whether the symptom is present, not present, or negated, respectively.

Symptom-modifier relation vector representation

The symptom-modifier relation vector is a much larger vector than the symptom-only and comprises three appended vectors containing: (i) the absence or presence of 110 unique body parts, (ii) the absence or value of a symptom duration, and (iii) the absence, negation or value or a symptom severity.

Step 4: Triage and diagnosis

We utilise SVM classification and regression models to triage and diagnose patients' posts, respectively, from the vector representations described above. For Q1, the recommendation from a doctor or combination of doctors is the class label of the post; see section Problem setting in Methods for a description. To build a binary classifier, we first combine the *Send to a GP* and *Send to hospital* recommendations to represent a

single class, *Send*. The SVM is trained to distinguish between the *Stay at home* and the *Send* options; we call this *SVM Classifier 1*. Next, the posts labelled as *Stay at home* are discarded and *SVM Classifier 2* is built utilising the remaining posts to classify the *Send to GP* and *Send to hospital* recommendations. This results in a hierarchical classifier for COVID-19 triage.

For diagnosing COVID-19 cases, we deploy a variant of SVM, called *Support Vector Regression* (SVR) [6], to estimate the probability of COVID-19. We use the GTP that is derived from answers to Q2, as the dependent variable. SVR takes as input a high dimensional feature vector such as a *symptom-only* or *symptom-modifier relation* vector representation, as described above. Classification is performed using the three decision functions, f_1 , f_2 , and f_3 , described previously.

Results

Evaluation

We evaluate the performance of the CRF and SVM classification algorithms using the standard measures of precision (P), recall (R) and Macro and Micro-averaged F_1 scores [33]. Macro-averaged scores are computed by considering the score independently for each class and then taking the average, while Micro-averaged scores are computed by considering all the classes together. As our dataset is relatively balanced with *COVID* and *NO_COVID* classes, we report Micro-averaged scores for SVR classification. On the other hand, in case of concept extraction, the *Other* class dominates. So, we report the Macro-averaged scores for the CRF classification results.

Experimental setup

For the CRF we report 3-fold cross validated Macro-averaged results. Specifically, we trained each fold by a Python wrapper [34] for CRFsuite, see [35]. For relation extraction, we ran our unsupervised rule-based algorithm on the 500 posts and calculated the F_1 scores by varying distances considering the two cases with and without stop words.

We constructed SVM binary classifiers, *SVM Classifier 1* and *SVM Classifier 2*, using the Python wrapper for LIBSVM [36] implemented in Sklearn [37] with both Linear and Gaussian *Radial Basis Function* (RBF) kernels [7]. Similarly, the SVR [38], implemented using LIBSVM, is built with both Linear and RBF kernels. The hyperparameters ($C=10$ for the penalty, $\gamma=0.01$ for the RBF kernel, and $\epsilon=0.5$ for the threshold) were discovered using grid search [37].

We simulated two cases for COVID-19 triage and diagnosis. First SVM and SVR models trained with the ground truth examine the predictive performance when they are deployed as stand-alone applications. Second, when trained with the predictions from CRF and RB classifier, they resemble an end-to-end NLP application. To get a comparable result, the models were always tested with the ground truth. As a measure of performance, we report Macro and Micro-averaged F_1 scores for SVM classifiers and SVR, respectively.

Table 2. The concept extraction using CRF and relation extraction using RB classifier

results on 3-fold cross validation.

Concept extraction using CRF					Relation extraction using RB classifier						
Label	P	R	F_1	Support	Distance	With stop words			Without stop words		
						P	R	F_1	P	R	F_1
SYM	0.94	0.97	0.95	1300	2	0.74	0.63	0.68	0.74	0.64	0.69
SEVERITY	0.80	0.79	0.79	437	3	0.75	0.67	0.71	0.75	0.67	0.71
BPOC	0.92	0.83	0.87	356	4	0.75	0.69	0.72	0.75	0.69	0.72
DURATION	0.87	0.91	0.89	667	5	0.75	0.71	0.73	0.74	0.71	0.73
INTENSIFIER	0.88	0.97	0.92	494	6	0.74	0.72	0.73	0.74	0.72	0.73
NEGATION	0.83	0.89	0.86	338	7	0.73	0.73	0.73	0.73	0.73	0.73
OTHER	0.99	0.98	0.98	16892							
Macro-average	0.89	0.89	0.89								




Figure 4. Support ratio of triage classes across models for Question 1 classification tasks. Absolute numbers for the Send to hospital class in test sets are as follows: A=10, B=12, AB(R-a)=14, AB(R-t)=5, BC(R-a)=6, AC(R-a)=5, ABC(R-a)=9; the value for the remaining models is zero.




Figure 5. Support ratio of diagnosis classes across models and three decision functions for Question 2 classification tasks.

Table 3. Question 1: Hierarchical classification results for RBF kernel using the symptom-modifier relation vector.

Symptom-modifier relation vector												
Trained on the ground truth							Trained on the CRF predictions					
Model	SVM Classifier 1			SVM Classifier 2			SVM Classifier 1			SVM Classifier 2		
	P	R	F_1	P	R	F_1	P	R	F_1	P	R	F_1
A	0.82	0.91	0.86	0.73	0.95	0.83	0.81	0.89	0.85	0.72	0.91	0.80
B	0.73	0.77	0.75	0.81	0.99	0.89	0.74	0.74	0.74	0.81	0.99	0.89
C	0.85	0.98	0.91	—	—	—	0.85	0.96	0.90	—	—	—
AB(R-a)	0.70	0.75	0.72	0.80	0.96	0.88	0.73	0.71	0.71	0.81	0.96	0.88
AB(R-t)	0.84	0.96	0.89	0.85	1.00	0.92	0.84	0.94	0.88	0.84	1.00	0.92
BC(R-a)	0.72	0.75	0.73	0.92	1.00	0.96	0.74	0.71	0.72	0.92	1.00	0.96
BC(R-t)	0.86	0.99	0.92	—	—	—	0.88	0.98	0.93	—	—	—
AC(R-a)	0.79	0.87	0.83	0.89	1.00	0.94	0.81	0.85	0.83	0.89	1.00	0.94
AC(R-t)	0.88	0.98	0.93	—	—	—	0.88	0.98	0.93	—	—	—
ABC(R-a)	0.70	0.76	0.73	0.89	0.99	0.93	0.72	0.72	0.72	0.89	1.00	0.94
ABC(R-t)	0.88	0.99	0.93	—	—	—	0.89	0.98	0.93	—	—	—
Symptom-only vector												
Trained on the ground truth							Trained on the CRF predictions					
Model	SVM Classifier 1			SVM Classifier 2			SVM Classifier 1			SVM Classifier 2		
	P	R	F_1	P	R	F_1	P	R	F_1	P	R	F_1
A	0.83	0.91	0.87	0.74	0.85	0.79	0.84	0.89	0.87	0.74	0.82	0.78

B	0.71	0.81	0.76	0.81	0.98	0.89	0.74	0.79	0.77	0.82	0.98	0.89
C	0.87	0.97	0.92	—	—	—	0.86	0.95	0.90	—	—	—
AB(R-a)	0.69	0.75	0.72	0.83	0.96	0.89	0.72	0.76	0.73	0.83	0.92	0.87
AB(R-t)	0.85	0.94	0.89	0.85	1.00	0.92	0.87	0.93	0.90	0.84	0.98	0.90
BC(R-a)	0.71	0.79	0.75	0.92	0.99	0.95	0.72	0.78	0.75	0.92	0.99	0.95
BC(R-t)	0.88	0.98	0.93	—	—	—	0.87	0.97	0.92	—	—	—
AC(R-a)	0.80	0.86	0.83	0.89	1.00	0.94	0.80	0.86	0.83	0.89	1.00	0.94
AC(R-t)	0.90	0.98	0.94	—	—	—	0.89	0.95	0.92	—	—	—
ABC(R-a)	0.68	0.74	0.71	0.90	1.00	0.95	0.71	0.76	0.73	0.89	0.99	0.93
ABC(R-t)	0.90	0.98	0.94	—	—	—	0.90	0.95	0.92	—	—	—

Table 4. Question 2: Micro-averaged F_1 results for different models and decision functions.

Here A, B, C are three medical doctors (abbreviated as Dr) who took part in the experiment.

Model	Trained on the ground truth						Trained on the CRF predictions					
	Symptom-modifier			Symptom-only			Symptom-modifier			Symptom-only		
	1	2	3	1	2	3	1	2	3	1	2	3
A	0.72	0.61	0.78	0.70	0.59	0.74	0.68	0.64	0.76	0.50	0.79	0.74
B	0.78	0.61	0.76	0.78	0.62	0.77	0.76	0.64	0.77	0.78	0.57	0.74
C	0.87	0.75	0.87	0.88	0.75	0.87	0.86	0.75	0.87	0.87	0.74	0.86
AB	0.72	0.66	0.74	0.74	0.65	0.75	0.70	0.65	0.73	0.71	0.66	0.74
BC	0.84	0.76	0.84	0.85	0.79	0.86	0.83	0.76	0.83	0.85	0.78	0.86
AC	0.81	0.73	0.81	0.83	0.74	0.83	0.80	0.74	0.82	0.80	0.73	0.81
ABC	0.74	0.67	0.76	0.75	0.67	0.77	0.72	0.69	0.76	0.74	0.69	0.77

Evaluation outcomes

The concept and relation extraction phases produce excellent and very good predictive performances, respectively; see Table 2. The triage classification results from Q1 are shown in Table 3. When we trained the models with the *Symptom-modifier* vector representations from the ground truth, the results of SVM Classifier 1 and 2 are in the range of 72-93% and 83-96%, respectively. The Symptom-only vector representations produces results in the range of 71-94% and 79-95%. These results suggest that we can achieve very good predictive performance for classifying *Stay at home* and *Send*, and for

Send to a GP and *Send to hospital*. In general, risk-tolerant models achieve better performance than the risk-averse models. However, since, in the test set, posts with the label *Send to hospital* are missing for some models (as can be seen from Figure 4) we cannot report them. We report Macro-averaged F_1 results since Q1 is framed as a decision problem, where weights for the classes are a priori equal. The results obtained after training with CRF predictions are in similar ranges for both representations and classifiers. This is important, because it indicates that an end-to-end NLP application is likely to produce similar predictive performance.

Regarding Q2, when we trained the models with the *Symptom-modifier* vector representation from ground truth, the results of COVID-19 diagnosis are in the range of 72-87%, 61-76%, and 74-87% for the 1, 2, and 3

decision functions, respectively; see Table 4. The Symptom-only vector representation produce results in the range of 70-88%, 59-79%, and 74-87%.

In general, \hat{f} models perform better due to the omission of borderline cases where the GTPs are exactly 0.5. The support ratios for each model for different decision functions, is shown in Figure 5. When we trained the models with the *Symptom-modifier* vector representation from the CRF predictions, the results are in the range of 68-86%, 64-76%, and 73-87% for the \hat{f} , \hat{f} , and \hat{f} decision functions, respectively. This indicates that, for diagnosis as well as triage, an end-to-end NLP application is likely to perform similarly to standalone applications. Here, we report Micro-averaged F_1 scores since, in our dataset, *NO_COVID* cases dominate; this largely resembles the natural distribution in the population, where people tested positive for Coronavirus are relatively a low percentage in the whole population even when the prevalence of the virus is high.

Finally, we trained our models using a Linear kernel, but found that RBF dominates in most of the cases; however, Linear kernels are useful in finding feature importance [39].

Discussion

Comparison with prior work

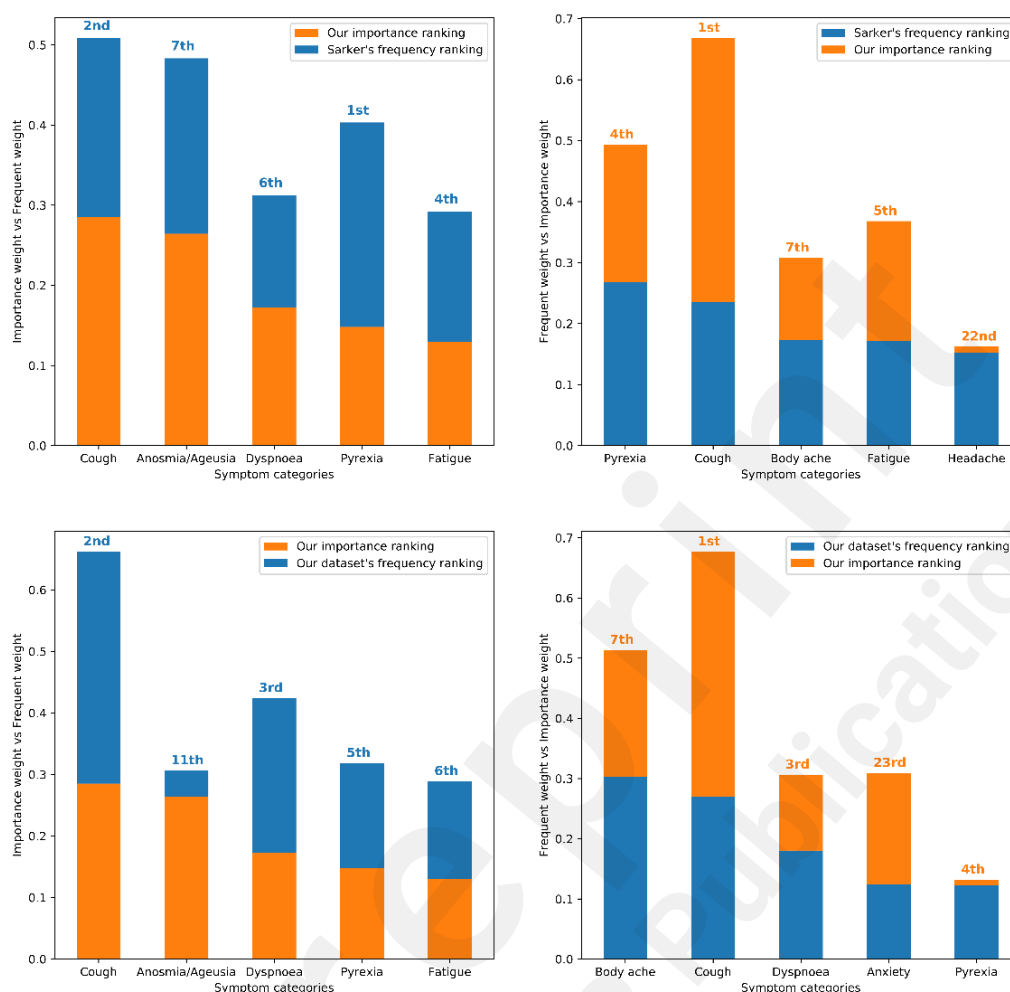
To quantify the important predictive features in the training set, we experimented with COVID-19 diagnosis using Linear Kernel SVR regression. More specifically, we used the Symptom-only vector representation constructed from the ground truth. We summed feature weights for each S_i in $\langle S_0, S_1, \dots, S_n \rangle$ from the seven models, and the three decision functions; see Methods. The features are then mapped to the categories found in the Twitter COVID-19 lexicon compiled by [3]. The top 5 important features in our dataset are: *Cough*, *Anosmia/Agusia*, *Dyspnoea*, *Pyrexia* and *Fatigue*. In [22], quoted 4 of these symptoms as the most prevalent Coronavirus symptoms, strongly correlating with our findings.

To compare our importance ranking with that of Sarker et al.'s [3] frequent categories, we compiled the corresponding frequencies of our 5 most important symptoms. Normalised weights and frequencies are then plotted in Figure 6. The top-left stacked bar chart compares our 5 most important features with Sarker's frequencies. *Cough* is the most important symptom from our dataset, where it is the 2nd most frequent. *Anosmia/Agusia* ranks 2nd in our importance list, while it is 7th in the most frequent list. *Pyrexia* comes 1st and 4th in both the frequent and importance lists, respectively.

The top-right chart in Figure 6 shows a comparison between Sarker's frequent ranking and our importance ranking. Here, we select top 5 most frequent symptoms from Sarker's frequency list and normalise them. They are: *Pyrexia*, *Cough*, *Body ache*, *Fatigue*, and *Headache*. We took the corresponding importance weights of those symptoms and plotted the stacked bar chart. Here, *Headache* ranks 22nd in our importance ranking, while it is 5th in the frequent ranking. We find a large difference between the two rankings, implying that the top most frequent symptoms are not necessarily the most important ones.

Next, we compare our most important feature weights with our dataset's frequency ranking using the methods described above. From the bottom left stacked bar chart of Figure 6, we observe that *Anosmia/Agusia* are a relatively low in order in the frequency ranking, i.e. 11th. Like Sarker's, *Cough* comes 2nd in our dataset's frequency ranking.

Finally, the bottom-right chart in Figure 6 refers to the comparison between our dataset's frequency and importance rankings of the corresponding symptoms. We observe that *Anxiety* ranks 4th in the most frequent list, where it is very low, i.e. 23rd, in the most importance ranking.



Figure

6. Feature comparison between our most important features and Sarker's most frequent symptoms (top row), and between our most important features and our most frequent symptoms (bottom row). The feature importance rankings are obtained from an SVM linear kernel using the symptom-only vector representation.

Principal findings

This study demonstrates the potential to triage and diagnose COVID-19 patients from their social media posts. We have presented a proof of concept system to predict a patient's health state by building machine learning models from their narrative. The models are trained in two ways; using (i) ground truth labels, and (ii) predictions obtained from the NLP pipeline. Trained models are always tested on the ground truth labels. We obtained good performances in both cases which indicates that an automated NLP pipeline could be used to triage and diagnose patients from their narrative; see Evaluation outcomes in the Results section. In general, health professionals and researchers could deploy, triage models to determine the severity of COVID-19 cases in the population, and diagnostic models to gauge the prevalence of the pandemic.

Limitations

It is worth reiterating that social media posts, which are known to be noisy, are not on a par with the consultation that a patient would have with a doctor. Our manually annotated dataset has two main limitations. First having only three experts limits the quality of our labelling, although we deem this study to be a proof of concept. A larger number of experts, including more senior doctors would be beneficial in a follow-up study. The robustness of our results could be further improved by both increasing the size of our dataset and

introducing posts from several alternate sources. Given that the posts come from social media, it is not clear whether the results could be used as such in a diagnostic system, without combining them with actual consultations. However, it is worth noting that medical social media such as the posts we used herein, may uncover novel information regarding COVID-19.

Conclusion

The Coronavirus pandemic has drawn a spotlight on the need to develop automated processes to provide additional information to researchers, health professionals and decision-makers. Medical social media comprises a rich resource of timely information that could fit this purpose, for example, using outputs from our triage and diagnosis system to indicate the severity and estimate the prevalence of the disease in the population. We have demonstrated that it is possible to create an automated triage and diagnosis system, despite the heterogeneous nature of typical social media posts.

Author contributions

All authors were involved in the design of the work. The first author wrote the code. The first three authors drafted the article and all authors critically revised the article.

Conflict of interest statement

None declared.

Abbreviations

CRF: conditional random fields
CT: computed tomography
EHR: electronic health records
LR: logistic regression
NLP: natural language pipeline
RBF: radial basis function
SVM: support vector machine
SVR: support vector regressions
XGB: gradient boosting

References

1. Obeid J. S, Davis M, Turner M et al. An artificial intelligence approach to COVID-19 infection risk assessment in virtual visits: A case report. J Am Med Inform Assoc 2020; 27 (8):1321–1325.
2. Schwab P, DuMont Schütte A, Dietz B et al. Clinical Predictive Models for COVID-19: Systematic Study. J Med Internet Res 2020; 22 (10):e21439.
3. Sarker A, Lakamana S, Hogg-Bremer W et al. Self-reported COVID-19 symptoms on Twitter: An analysis and a research resource. J Am Med Inform Assoc 2020; 27 (8):1310–1315.

4. Qin L, Sun Q, Wang Y et al. Prediction of number of cases of 2019 novel coronavirus (COVID-19) using social media search index. *Int J Environ Res Public Health* 2020; 17 (7):2365.
5. Norman G. Likert scales, levels of measurement and the “laws” of statistics. *Adv Health Sci Educ* 2010; 15 (5):625–632.
6. Drucker H, Burges C. J, Kaufman L et al. Support Vector Regression Machines. *Adv Neural Inf Process Syst* 1996; 9:155–161.
7. Marsland S. *Machine Learning: An Algorithmic Perspective* (second ed.), Chapman & Hall/CRC (2014):169-188.
8. Arrieta A, García-Prado A, Gonzalez P and Pinto-Prades J. L. Risk attitudes in medical decisions for others: An experimental approach. *Health Economics* 2017; 26:97–113.
9. Sutton C and McCallum A. An Introduction to Conditional Random Fields. *Found Trends Mach Learn* 2012; 4 (4):267–373.
10. Bach N and Badaskar S. A Review of Relation Extraction. *Lit. Rev. Lang. Stat. II* 2007; 2:1–15.
11. Zimmerman R. K, Balasubramani G, Nowalk M. P et al. Classification and Regression Tree (CART) analysis to predict influenza in primary care patients. *BMC Infect Dis* 2016; 16 (1):1–11.
12. Brinati D, Campagner A, Ferrari D et al. Detection of COVID-19 Infection from Routine Blood Exams with Machine Learning: A Feasibility Study. *J Med Syst* 2020; 44 (8):1–12.
13. Wynants L, Van Calster B, Collins G. S et al. Prediction models for diagnosis and prognosis of COVID-19: systematic review and critical appraisal. *British Medical Journal* 2020;369.
14. Zens M, Brammertz A, Herpich J et al. App-Based Tracking of Self Reported COVID-19 Symptoms: Analysis of Questionnaire Data. *J Med Internet Res* 2020; 22 (9):e21956.
15. Menni C, Valdes A. M, Freidin M. B et al. Real-time tracking of self-reported symptoms to predict potential COVID-19. *Nat Med* 2020; 26 (7):1037–1040.
16. Judson T. J, Odisho A. Y, Neinstein A. B et al. Rapid design and implementation of an integrated patient self-triage and self-scheduling tool for COVID-19. *J Am Med Inform Assoc.* 2020; 27 (6):860–866.
17. Liu Y, Wang Z, Tian Y et al. A COVID-19 Risk Assessment Decision Support System for General Practitioners: Design and Development Study. *J Med Internet Res* 2020; 22 (6):e19786.
18. Einstein Data4u. Diagnosis of COVID-19 and its clinical spectrum AI and Data Science supporting

- clinical decision ((from 28th Mar to 3rd Apr)). <https://www.kaggle.com/einsteindata4u/covid19>. Accessed: 2021-02-24.
19. Wang K, Zuo P, Liu Y et al. Clinical and Laboratory Predictors of In-hospital Mortality in Patients With Coronavirus Disease-2019: A Cohort Study in Wuhan, China. *Clinical infectious diseases* 2020; 71 (16):2079–2088.
 20. Vaid A, Somani S, Russak A. J et al. Machine Learning to Predict Mortality and Critical Events in a Cohort of Patients With COVID-19 in New York City: Model Development and Validation. *J Med Internet Res* 2020; 22 (11):e24018.
 21. Wagner T, Shweta F, Murugadoss K et al. Augmented curation of clinical notes from a massive EHR system reveals symptoms of impending COVID-19 diagnosis. *Elife* 2020; 9:e58227.
 22. Mizrahi B, Shilo S, Rossman H et al. Longitudinal symptom dynamics of COVID-19 infection. *Nat Commun* 2020; 11 (1):1–10.
 23. Wang J, Abu-el Rub N, Gray J et al. COVID-19 SignSym – A fast adaptation of general clinical NLP tools to identify and normalize COVID-19 signs and symptoms to OMOP common data model. *J Am Med Inform Assoc* 2021.
 24. López-Úbeda P, Díaz-Galiano M. C, Martín-Noguerol T et al. COVID-19 detection in radiological text reports integrating entity recognition. *Comput Biol Med* 2020; 127:104066.
 25. Bodenreider O. The Unified Medical Language System (UMLS): integrating biomedical terminology. *Nucleic Acids Res* 2004; 32(1):267–270.
 26. Mikolov T, Sutskever I, Chen K et al. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems*, pages 3111–3119, May 2013.
 27. Patient. <https://patient.info/forums/discuss/browse/coronavirus-covid-19--4541>. Accessed: 2021-01-18.
 28. Landis J. R and Koch G. G. The Measurement of Observer Agreement for Categorical Data. *Biometrics* 1977: 159–174.
 29. Schober P, Boer C and Schwarte L. A. Correlation Coefficients: Appropriate Use and Interpretation. *Anesthesia & Analgesia* 2018; 126 (5):1763–1768.
 30. de Vet H. C, Mokkink L. B, Terwee C. B et al. Clinicians are right not to like Cohen's κ . *British Medical*

- Journal 2013; 346:f2125.
31. Cunningham H, Maynard D and Bontcheva K. Text Processing with GATE (Version 6). CA: Gateway Press 2011.
 32. Hasan A, Levene M and Weston D. Learning structured medical information from social media. J Biomed Inform, 2020; 110:103568.
 33. Manning C and Schutze H. Foundations of Statistical Natural Language Processing. MIT press, 1999.
 34. Python-crfsuite. <https://python-crfsuite.readthedocs.io/en/latest/>. Accessed: 2018-03-14.
 35. Okazaki N. CRFSuite: A Fast Implementation of Conditional Random Fields (CRFs) 2007. <http://www.chokkan.org/software/crfsuite/>.
 36. Chang C.-C and Lin C.-J. LIBSVM: A Library for Support Vector Machines. ACM Trans Intell Syst Technol 2011; 2(3):1–27.
 37. Pedregosa F, Varoquaux G, Gramfort A et al. Scikit-learn: Machine Learning in Python. J Mach Learn Res 2011; 12:2825–2830.
 38. Support Vector Machines. <https://scikit-learn.org/stable/modules/svm.html>. Accessed: 2021-01-19.
 39. Weston J, Mukherjee S, Chapelle O et al. Feature selection for SVMs. In Advances in Neural Information Processing Systems 2000.

Supplementary Files

Figures

A patient-authored social media post is annotated with symptoms (light green), affected body parts (pale blue), duration (light yellow) and severities (pink). The phrases in the square brackets show relations between a symptom and a body part/duration/severity, when the distance was greater than 1. This annotated post was presented to three doctors to triage and diagnose the author of the post by answering Questions 1 and 2, respectively.

Hi im currently the same **day 27** since my symtoms started , deep **breathing hurts** [upper chest area][throat] which is **upper chest area** into **throat** , **breathing** [slightly][laboured] is **slightly laboured** time to time , **dry cough** **on and off** , also have **major** **fatigue weakness** took a course of Amoxcillian given by GP which made no change to me , have **asthma** so take my inhalers which aint making no change , never been **so unwell** in my life ! ! !

Question 1: Please specify recommendation from one of the options below:

- ☐ Stay at home
- ☐ Send to a GP
- ☐ Send to hospital


Question 2: How would you rate the chance of this person having COVID-19 on a range of 1 to 5?

- ☐ 1 (Very unlikely)
- ☐ 2 (Unlikely)
- ☐ 3 (Uncertain)
- ☐ 4 (Likely)
- ☐ 5 (Very likely)

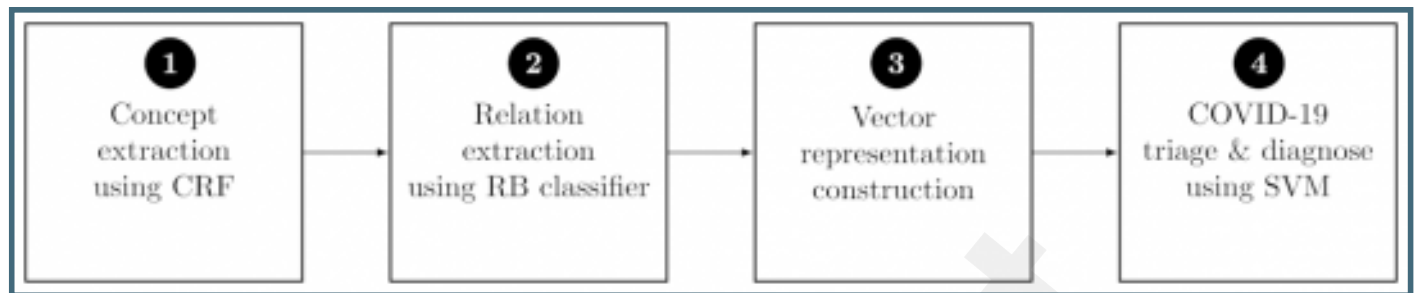
Question 3: Was the highlighted text sufficient in reaching your decision?

- ☐ Yes
- ☐ No


Frequency distribution of annotated classes/concepts from the text are shown. We also show the percentage of each class after discounting the OTHER labels. The average number of tokens per post is 130.17(SD=97.83). Here, SYM, DURATION, INTENSIFIER, SEVERITY, BPOC and NEGATION denote symptoms, duration, intensifiers, severity, body parts and negations, respectively.




A block diagram of COVID-19 triage and diagnosis text processing pipeline. Here, CRF, RB classifier and SVM are acronyms for Conditional Random Fields, Rule-Based classifier and Support Vector Machine, respectively.



Support ratio of triage classes across models for Question 1 classification tasks. Absolute numbers for the Send to hospital class in test sets are as follows: A=10, B=12, AB(R-a)=14, AB(R-t)=5, BC(R-a)=6, AC(R-a)=5, ABC(R-a)=9; the value for the remaining models is zero.



Support ratio of diagnosis classes across models and three decision functions for Question 2 classification tasks.



Feature comparison between our most important features and Sarker's most frequent symptoms (top row), and between our most important features and our most frequent symptoms (bottom row). The feature importance rankings are obtained from an SVM linear kernel using the symptom-only vector representation.

