



Correlation-Based Discovery of Disease Patterns for Syndromic Surveillance

Michael Rapp^{1*}, Moritz Kulesa¹, Eneldo Loza Mencía¹ and Johannes Fürnkranz²

¹ Knowledge Engineering Group, Technical University of Darmstadt, Darmstadt, Germany, ² Computational Data Analysis Group, Johannes Kepler University Linz, Linz, Austria

OPEN ACCESS

Edited by:

Panagiotis Papapetrou,
Stockholm University, Sweden

Reviewed by:

Ioanna Millou,
Stockholm University, Sweden

Denise Voci,
University of Klagenfurt, Austria

Myra Spiliopoulou,
Otto von Guericke University
Magdeburg, Germany

*Correspondence:

Michael Rapp
michael.rapp.ml@gmail.com
orcid.org/0000-0001-8570-8240

Specialty section:

This article was submitted to
Data Mining and Management,
a section of the journal
Frontiers in Big Data

Received: 27 September 2021

Accepted: 21 December 2021

Published: 13 January 2022

Citation:

Rapp M, Kulesa M, Loza Mencía E
and Fürnkranz J (2022)
Correlation-Based Discovery of
Disease Patterns for Syndromic
Surveillance.
Front. Big Data 4:784159.
doi: 10.3389/fdata.2021.784159

Early outbreak detection is a key aspect in the containment of infectious diseases, as it enables the identification and isolation of infected individuals before the disease can spread to a larger population. Instead of detecting unexpected increases of infections by monitoring confirmed cases, syndromic surveillance aims at the detection of cases with early symptoms, which allows a more timely disclosure of outbreaks. However, the definition of these disease patterns is often challenging, as early symptoms are usually shared among many diseases and a particular disease can have several clinical pictures in the early phase of an infection. As a first step toward the goal to support epidemiologists in the process of defining reliable disease patterns, we present a novel, data-driven approach to discover such patterns in historic data. The key idea is to take into account the correlation between indicators in a health-related data source and the reported number of infections in the respective geographic region. In an preliminary experimental study, we use data from several emergency departments to discover disease patterns for three infectious diseases. Our results show the potential of the proposed approach to find patterns that correlate with the reported infections and to identify indicators that are related to the respective diseases. It also motivates the need for additional measures to overcome practical limitations, such as the requirement to deal with noisy and unbalanced data, and demonstrates the importance of incorporating feedback of domain experts into the learning procedure.

Keywords: outbreak detection, syndromic surveillance, rule learning, knowledge discovery, time series analysis

1. INTRODUCTION

Throughout history, major outbreaks of infectious diseases have caused millions of deaths and, therefore, pose a serious threat to public health. Among the most well-known outbreaks is the *Great Influenza Pandemic* between the years 1918 and 1920, which has killed approximately 40 million people worldwide, as well as the recent, still ongoing, pandemic of SARS-CoV-2 (Barro et al., 2020). A fundamental strategy to diminish or even prevent the spreading of infectious diseases is to detect local outbreaks as early as possible in order to identify and isolate infected individuals. For the early detection of unexpected increases in the number of infections, which may be an indicator for an outbreak, infectious diseases are under constant surveillance by epidemiologists.

Besides tracking the number of confirmed infections based on laboratory testing, a promising approach to outbreak detection is *syndromic surveillance* (Henning, 2004), which focuses on monitoring the number of cases with early symptoms. Compared to laboratory testing, which can

take several days until results are available, it allows for a more timely detection of outbreaks. Moreover, a much larger population can be put under surveillance by using health-related data sources that do not depend on confirmed results. For example, the number of antipyretic drug sales in pharmacies could be considered as an indicator for an outbreak of influenza. Or, based on data that is gathered in emergency departments, the number of patients with a fever or other related symptoms could serve as another indicator for this particular disease.

One of the major challenges in syndromic surveillance is the definition of such indicators, also referred to as *syndromes* or *disease patterns*. They highly depend on the infectious disease and the data source under surveillance. Since early symptoms are usually shared among many diseases and because a particular disease can have several clinical pictures at early stages of an infection, it is difficult to obtain reliable syndromes. For this reason, the definition of disease patterns is usually based solely on expert knowledge of epidemiologists, a time-consuming and laborious process (Mandl et al., 2004). This motivates the demand for tools that allow for a user-guided generation and comparison of syndrome definitions. To be useful in practice, such tools should be flexible enough to be applied to different types of data (Hopkins et al., 2017).

In this work, we investigate a data-driven approach that aims at supporting epidemiologists in the process of identifying disease patterns for infectious diseases. It discovers syndrome definitions from health-related data sources, based on their correlation to the reported number of infections in the respective geographical area. As the first contribution of this work, we introduce a formal definition of this correlation-based discovery task. Our second contribution is an algorithm for the automatic extraction of disease patterns that utilizes techniques from the field of inductive rule learning. To provide insight into the data, the syndromes it discovers may be suggested to epidemiologists, who can adjust the input or the parameters of the algorithm to interactively refine the syndromes according to their domain knowledge. To better understand the capabilities and shortcomings of the proposed method, we evaluate its ability to reconstruct randomly generated disease patterns with varying characteristics. Furthermore, we apply our approach to emergency department data to learn disease patterns for Influenza, Norovirus and SARS-CoV-2. To assess the quality of the obtained patterns, we discuss the indicators they are based on and relate them to the number of infections according to publicly available reports, as well as handcrafted syndrome definitions.

2. PRELIMINARIES

In the following, we formalize the problem that we address in the present work, including a definition of relevant notation and an overview of related work.

2.1. Problem Definition

We are concerned with the deduction of patterns from a health-related data source $X = \{\mathbf{x}_1, \dots, \mathbf{x}_N\} \subset \mathcal{X}$. It incorporates information about individual *instances* $\mathbf{x}_n \in \mathcal{X}$ from a population \mathcal{X} , which are represented in terms of a finite set of predefined

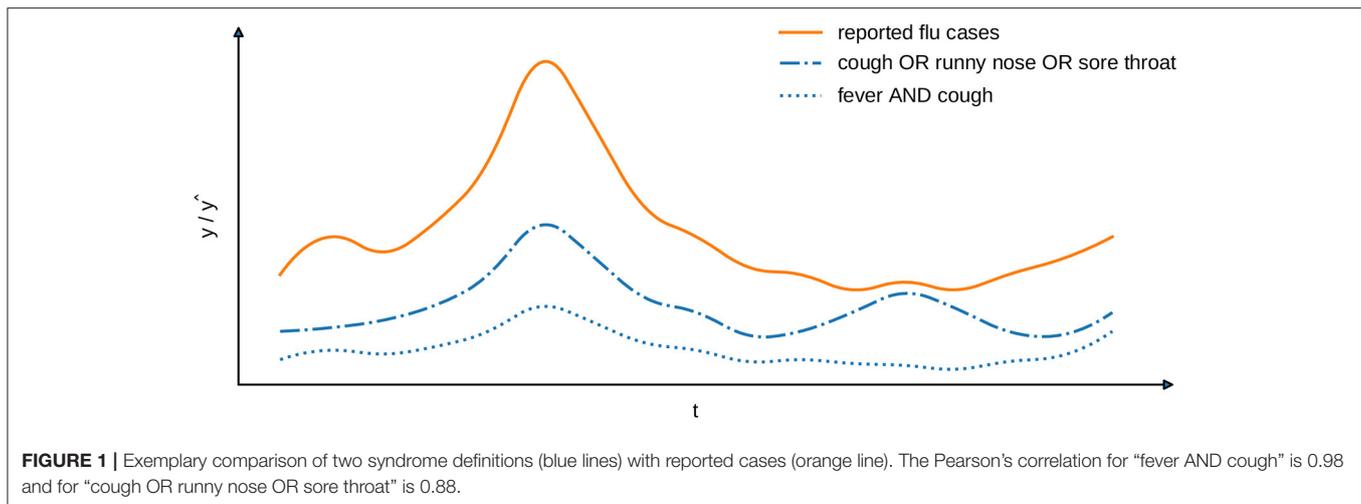
attributes $A = \{a_1, \dots, a_K\}$. An instance $\mathbf{x} = (x_1, \dots, x_K)$, e.g., representing a patient that has received treatment in an emergency department, assigns discrete or numerical values x_k to the k -th attribute a_k . For example, discrete attributes can be used to specify a patient's gender, whereas numerical attributes are suitable to encode continuous values, such as body temperature, blood pressure, or the like. The values for individual attributes may also be missing, e.g., because some medical tests have not been carried out as part of an emergency treatment. In addition, each instance in a data source is subject to a mapping $h: \mathbb{N}_+ \rightarrow \mathbb{N}_+$. It associates the n -th instance with a corresponding period in time, identified by a timestamp $t = h(n)$. Instances that correspond to the same interval, e.g., to the same week, are assigned the same timestamp $t: 1 \leq t \leq T$.

For each timestamp t , the instances in a data source may be associated with, a corresponding *target variable* $y_t \in \mathcal{Y}$ must be provided as part of a secondary data source $\mathbf{y} = (y_1, \dots, y_T) \in \mathcal{Y}$. The target space \mathcal{Y} corresponds to the number of infections that may occur within consecutive periods of time. Consequently, a particular target variable $y_t \in \mathbb{N}_+$ specifies how many cases related to a particular infectious disease have been reported for the t -th time interval.

The learning task, which we address in this work, requires to find an interpretable model $f: \mathcal{X} \rightarrow \mathcal{Y}$. Given a set of instances $X \subset \mathcal{X}$ that are mapped to corresponding time intervals via a function h , it provides an estimate $\hat{\mathbf{y}} = f(X, h) = (\hat{y}_1, \dots, \hat{y}_T) \in \mathcal{Y}$ of the number of infections per time interval. The selection of instances and the number of reported cases, which are provided for the training of such model, must neither originate from the same source, nor comprise information about identical subgroups of the population. As a consequence, the estimates of a model are not obliged to reflect the provided target variables in terms of their absolute values. Instead, we are interested in capturing the correlation between indicators that may be derived from the training instances and the number of infections that have arisen during the considered timespan. To assess the quality of a model, we compare the estimates it provides to the target variables with respect to a suitable correlation coefficient, such as *Spearman's* ρ , *Kendall's* τ , or *Pearson's* correlation. For example, one could align patient data from a medical office with locally reported flu cases. In **Figure 1**, we show the number of patients per timestamp that fulfill two exemplary syndrome definitions of this particular disease. One of the syndromes ("fever AND cough") covers less cases than the other, but has a higher Pearson correlation coefficient (0.98 compared to 0.88) and, therefore, matches the locally reported cases more closely.

2.2. Related Work

Disease patterns for syndromic surveillance are usually defined according to the knowledge of domain experts. This requires a manual examination of the available health-related data to identify indicators that may be related to a particular disease at hand. For example, Edge et al., 2006 and Muchaal et al. (2015) analyze information about the sales of pharmaceuticals to reason about the spread of Norovirus infections, based on their effectiveness against gastrointestinal symptoms. Similarly, the data that is gathered in emergency departments may also



serve as a basis for the definition of disease patterns. In this case, definitions are usually based on the symptoms of individual patients and the diagnoses made by the medical staff. For example, Ivanov et al. (2002) and Suyama et al. (2003) rely on standardized codes for the *International Classification of Diseases* (ICD) (Trott, 1977). Boender et al. (2021) additionally use chief complaints of the patients at the emergency departments. The majority of syndrome definitions are targeted at common infectious diseases, such as gastrointestinal infections, influenza-like illnesses, or respiratory diseases (e.g., Suyama et al., 2003; Heffernan et al., 2004; Boender et al., 2021; Bouchouar et al., 2021). However, they are also used to detect other health-related epidemics, e.g., increased usage of psychoactive substances (Nolan et al., 2017).

The deduction of indicators from unstructured data, such as textual reports of complaints or diagnoses, is particularly challenging. To be able to deal with such data, text documents are often represented in terms of keywords they consist of. For example, (Lall et al., 2017) use syndromes that apply to the keywords contained in medical reports. Similarly, Heffernan et al. (2004) use a list of exclusive keywords to reduce the chance of misclassifications, Bouchouar et al. (2021) utilizes regular expressions to extract symptoms from texts and Ivanov et al. (2002) use a classifier system that takes textual data as an input to assign syndromes to individual patients. In order to train a classifier, the latter approach requires labeled training data that must manually be created by experts. The analysis of textual data is even more profound in approaches to syndromic surveillance that are based on web data. For example, Velardi et al. (2014) analyze Twitter messages to capture indicators for the spread of influenza-like illnesses. Starting with a handcrafted set of medical conditions that are related to the respective disease, they learn a language model that aims to identify closely related terms based on clustering.

The problem of learning syndrome definitions in a data-driven way, without relying on expert knowledge, has for example been addressed by Kalimeri et al. (2019). The authors of this work propose an unsupervised, probabilistic

framework based on matrix factorization. Their goal is to identify patterns of symptoms in structured data that has been obtained from participatory systems. Given a set of 19 symptoms, e.g., fever or vomiting, they construct a matrix that incorporates information about the occurrences of individual symptoms over time. Ultimately, syndromes can be generated from this matrix by extracting latent features that correspond to linear combinations of groups of symptoms.

Another method that relies on structured data is proposed by Goldstein et al. (2011). It is aimed at capturing the likelihood of syndromes for a particular infectious disease. The authors propose to use expectation maximization and deconvolution to identify syndromes, which are highly correlated with the occurrences of symptoms that have been reported in regular time intervals. However, their approach does only allow to evaluate and compare disease patterns that have been specified in advance. Even though the aforementioned algorithms deal with structured data that is less cumbersome to handle than unstructured inputs, they have only been applied to small and pre-selected sets of features.

The problem of learning from assignments of target variables to sets of instances, rather than individual instances, is known as *multiple instance learning* (Carbonneau et al., 2018). Chevalyre and Zucker (2001) tackle such task by adapting the quality criterion used by the well-known rule learning method RIPPER. To be able to deduce classification rules from sets of instances, Bjerring and Frank (2011) incorporate the separate-and-conquer rule induction technique into a tree learner. Both approaches are limited to the assignment of a binary signal to a bag of instances and are not intended to cope with multiple instance regression tasks (Ray and Page, 2001). The mapping of numeric values to bags of instances, as in the syndrome definition learning task at hand, is a much less explored problem in the literature. We are not aware of any existing work that approaches this kind of problem with the goal to obtain rule-based models.

3. LEARNING OF SYNDROME DEFINITIONS

In the following, we propose an algorithm for the automatic induction of syndrome definitions, based on the indicators that can be constructed from a health-related data source. Each indicator c_m , which is included in such a model, refers to a certain attribute that is present in the data. It compares the values, which individual instances assign to this particular attribute, to a constant using relational operators, such as = if the attribute is discrete, or \leq and $>$ if it is numerical. By definition, if an indicator is concerned with an attribute for which an instance's value is missing, the indicator is not satisfied. We strive for a combination of different indicators via logical AND (\wedge) and OR (\vee) operators. The model that is eventually produced is given in *disjunctive normal form*, i.e., as a disjunction of conjunctions. Such a logical expression $r = r_1 \vee \dots \vee r_L$ with $r_l = c_{l,1} \wedge \dots \wedge c_{l,M}$ evaluates to $r(\mathbf{x}_n) = 1$ (*true*) or $r(\mathbf{x}_n) = 0$ (*false*), depending on whether it is satisfied by a given instance \mathbf{x}_n or not. If the context is clear, we abbreviate $c_{l,i}$ with c_i . For each time interval $t \in [1, \dots, T]$, the number of infected cases, as estimated by a logical expression r , calculate as

$$\hat{\mathbf{y}} = r(X) = \left(\sum_{\mathbf{x}_n \in X} \llbracket h(n) = t \rrbracket r(\mathbf{x}_n) \right)_{1 \leq t \leq T}, \quad (1)$$

where $\llbracket p \rrbracket = 1$ if the predicate p is true, and 0 otherwise. That is, for a particular time interval, the number of infections is given as the total number of instances that belong to the time interval according to the mapping h and match the indicators that are included in the logical expression. We refer to a logical expression r that does not consist of any indicators as the *empty hypothesis*. In such case, where $r(\mathbf{x}_n) = 0, \forall \mathbf{x}_n$, the estimates for all time intervals are zero.

The representation of syndromes introduced above is closely related to sets of conjunctive rules r_l as commonly used in *inductive rule learning*—an established and well-researched area of machine learning (e.g., Fürnkranz et al., 2012 provide an extensive overview on the topic). Consequently, we rely on commonly used techniques from this particular field of research to learn the definitions of syndromes. We use a sequential algorithm that starts with an empty hypothesis to which new conjunctions of indicators r_1, \dots, r_L are added step by step. Given a data source that incorporates many features, the number of possible combinations of indicators can be very large. For this reason, we rely on *top-down hill climbing* to search for suitable combinations. With such an approach, conjunctions of indicators that can potentially be added to a model are constructed greedily. At first, single indicators are taken into account individually. They are evaluated by computing the overall quality of a model that results from the addition of an indicator. The one that results in the greatest quality is ultimately selected. Afterwards, combinations that possibly result from a conjunction of already chosen indicators with an additional one are evaluated in the same way. The search continues to add more indicators, resulting in more restricted patterns that apply to fewer instances, as long as an improvement of the model's quality can be achieved.

Optionally, the maximum number of indicators per conjunction M can be limited via a parameter. If $M = 1$, the algorithm is restricted to learn disjunctions of indicators. Furthermore, we enforce a *minimum support* $s \in \mathbb{R}$ with $0 < s < 1$, which specifies the number of instances $N \cdot s$ a conjunction of indicators must apply to. Once it has decided for a conjunction of indicators to be included in the model, the algorithm attempts to learn another conjunction to deal with instances that have not yet been adequately addressed by the model. The training procedure terminates as soon as it is unable to find a new pattern that improves upon the quality of the model. In addition, an upper bound can be imposed on the number of disjunctions L by the user.

The search for suitable indicators and combinations thereof is guided by a target function to be optimized at each training iteration. It assesses the quality that results from adding an additional conjunction of indicators to an existing model in terms of a numeric score. We denote the estimates that are provided by a model after the l -th iteration as $\hat{\mathbf{y}}^{(l)}$. When adding a conjunction of indicators r_l to an existing model, the estimates of the modified model for each available time interval can be computed incrementally as

$$\hat{\mathbf{y}}^{(l)} = r^{(l)}(X) = r^{(l-1)}(X) + r_l(X). \quad (2)$$

Typically, the addition of indicators to a model does only affect the estimates for certain time intervals. The estimates for time intervals that are unaffected by the indicators remain the same as in the previous training iteration. We compare a model's estimates for consecutive time intervals to the numbers infections that are reported for each time interval, referred to as the *ground truth*, and assess their quality in terms of the *Pearson correlation coefficient*. In case of a positive correlation coefficient, the numbers of infections increase over time, whereas they decrease in case of a negative coefficient. Regardless of the overall trend in the number of infections, the provided estimates are strongly correlated with the ground truth, if the measure indicates a strong positive or negative correlation. We, therefore, assess the quality of a model in terms of the absolute Pearson correlation coefficient. At a particular training iteration, it can be computed in a single pass over the target time series \mathbf{y} and the current estimates $\hat{\mathbf{y}}^{(l)}$ according to the formula

$$m_p(\mathbf{y}, \hat{\mathbf{y}}^{(l)}) := \left| \frac{T \sum_t y_t \hat{y}_t \sum_t y_t \sum_t \hat{y}_t}{\sqrt{T \sum_t y_t^2 - \left(\sum_t y_t\right)^2} \sqrt{T \sum_t \hat{y}_t^2 - \left(\sum_t \hat{y}_t\right)^2}} \right|. \quad (3)$$

If the score that is computed for a potential modification according to the target function m_p is greater than the quality of the current model, it is considered an improvement. Among all possible modifications that are considered during a particular training iteration, the one with the greatest score is preferred. By using a measure of correlation, such as the Pearson correlation coefficient, we ensure that a model's estimate for individual time intervals must not necessarily be close to the corresponding number of infections that are reported for the same time interval

in terms of their absolute value. Instead, the estimated time series should correlate to the reported number of infections over time, i.e., it should replicate temporal patterns in the data, such as seasonal peaks with high numbers of infections. Relying on a correlation measure, rather than comparing a model's estimates to the ground truth in terms of absolute values, enables to learn from training instances that do not necessarily describe all cases that are included in the ground truth.

4. EVALUATION

To evaluate the previously proposed learning approach, we have implemented the methodology introduced above by making use of the publicly available source code of the BOOMER rule learning algorithm (Rapp et al., 2020). In adherence to the principles of reproducible research, our implementation can be accessed online¹. A major goal of the empirical study, which is outlined in the following, is to investigate whether the proposed methodology is able to deduce patterns from health-related data that correlate with the number of infections supplied via a secondary data source.

In a first step, we conducted a series of experiments using synthetic syndrome definitions. The objective was to validate the algorithm and to better understand its capabilities and limitations when it comes to the reconstruction of known disease patterns in a controlled environment. On the one hand, we considered synthetic syndromes with varying characteristics and complexity. On the other hand, we investigated the impact that the temporal granularity of the available data has on the learning approach. As elaborated below, the health-related data used in this work is available on a daily basis. By using synthetic syndromes, we were able to validate the algorithm's behavior when dealing with a broader or more fine-grained granularity as well. The use of synthetic syndromes also allows to investigate the ability of the proposed approach independently of the negative effects of artifacts that may be present in real data. This includes delays of reports, inaccuracies in the reported dates or instances that are present in one data source, but not in the other. For example, cases may have been reported in one of the considered districts, but have not been treated in one of the emergency departments included in our dataset. Vice versa, it is also possible that cases have been treated at one of the considered departments but have not been reported to the public agencies.

Such artifacts almost certainly play a role in our second experiment, where we tried to discover patterns that correlate with the publicly reported cases. We selected cases from the notifiable diseases of *Influenza* and *Norovirus*, which have extensively been studied in existing work (e.g., Heffernan et al., 2004; Muchaal et al., 2015; Kalimeri et al., 2019), as well as of the recently emerged SARS-CoV-2, which has for example been analyzed by Bouchouar et al. (2021). To evaluate whether the algorithm is able to identify meaningful indicators that are related to these particular diseases, we provide a detailed discussion of the discovered syndromes and compare them to manually defined disease patterns.

¹<https://github.com/mrapp-ke/SyndromeLearner>

TABLE 1 | Attributes included in the emergency department data.

Name	Type	#Values	Missing values in %
① Diagnosis			
MTS presentation	Discrete	57	0.01
MTS indicator	Discrete	179	5.10
ICD code	Discrete	5901	65.45
ICD code (short)	Discrete	1509	65.45
② Demographic Information			
Gender	Discrete	3	0.00
Age	Discrete	21	0.00
③ Vital Parameters			
Blood pressure systolic	Numeric	–	57.19
Blood pressure diastolic	Numeric	–	57.22
Temperature	Numeric	–	59.31
Respiration rate	Numeric	–	59.55
Pulse frequency	Numeric	–	91.91
Oxygen saturation	Numeric	–	57.18
④ Contextual Information			
No isolation	Discrete	11	1.81
Transport	Discrete	6	59.74
Disposition	Discrete	13	90.56

4.1. Experimental Setup

4.1.1. Health-Related Data

In line with related work on syndromic surveillance (e.g., Ivanov et al., 2002; Suyama et al., 2003; Boender et al., 2021), we relied on routinely collected and fully anonymized data from 12 German emergency departments, which capture information about patients that have consulted these institutions between January 2017 and April 2021. Although the emergency department data that we have used for our experiments provides valuable insights into the clinical symptoms of thousands of patients, it is restricted to cases that required medical treatment. As it does not include patients with early or mild symptoms that did not demand for medical attention, the data does only entail information about a small subpopulation. We consider the limited availability of training data, which is inherent to many machine learning problems and is not restricted to the type of data used in this work, as one of the main challenges of the learning task at hand. We hope to address this practical limitation by focusing on the correlation with the reported number of infections according to a secondary data source, which provides a more exact estimate of the overall population's size, rather than modeling the exact number of infectious cases that are present in the health-related data.

As shown in **Table 1**, we have extracted 15 attributes from the emergency department data. Each of the available attributes corresponds to one out of four categories. The first category, *diagnosis*, includes an initial assessment in terms of the *Manchester Triage System* (MTS) (Gräff et al., 2014). It is obtained for each patient upon arrival at an emergency department. Besides, this first category also comprises an ICD code (Trott,

TABLE 2 | Pearson correlation between cases identified by automatically learned syndromes on different feature categories and actually reported cases, as well as cases that match the handcrafted syndrome definitions.

	Feature categories				Reported cases	Handcrafted syndromes
	①	②	③	④		
Influenza	✓				0.9354	0.9917
	✓	✓			0.9357	0.9796
	✓		✓		0.9480	0.9768
	✓			✓	0.9366	0.9948
	✓	✓	✓	✓	0.9493	0.9800
SARS-CoV-2	✓				0.9399	0.9473
	✓	✓			0.9454	0.9219
	✓		✓		0.9528	0.8689
	✓			✓	0.9464	0.9506
	✓	✓	✓	✓	0.9528	0.8689
Norovirus	✓				0.7669	0.2761
	✓	✓			0.7669	0.2761
	✓		✓		0.7303	0.1470
	✓			✓	0.7167	0.1608
	✓	✓	✓	✓	0.7242	0.1672

1977) that represents a physician's assessment. In addition to the full ICD code, we also consider a more general variant that consists of the leading character and the first two digits (e.g., *U07* instead of *U07.1*). Features that belong to second category, *demographic information*, indicate the gender and age of patients, whereas *vital parameters* correspond to measurement data, such as blood pressure or pulse frequency, that may have been registered by medical staff. Features of the last category, *contextual information*, may provide information about why a patient was possibly quarantined (*isolation*), the means of transport used to get to the emergency department (*transport*), and the status when exiting the department (*disposition*).

In accordance with the findings of Hartnett et al. (2020), we observed a reduced number of emergency department visits during the first weeks of the SARS-CoV-2 pandemic. However, preliminary experiments suggested that this anomaly has no effect on the operation of our algorithm. To obtain a single dataset, we have merged the data from the considered emergency departments. It consists of approximately 1,900,000 instances. Each of the instances corresponds to a particular week (i.e., around 8,500 instances per week). Additional information about the emergency data used in this work is provided by Boender et al. (2021), who used a slightly different subset of the data set to evaluate their handcrafted syndrome definitions.

In contrast to existing work on the detection of disease patterns (e.g., Goldstein et al., 2011; Kalimeri et al., 2019), we have not applied any pre-processing techniques to the health-related data. As a consequence, the data contains a lot of noise, e.g., diagnoses related to injuries, and many missing values (cf. **Table 2**). Manual preparation of the data, such as

the selection of symptoms that are known to be related to an infectious disease, can be expected to reduce the level of noise and therefore reduce the risk of finding patterns that are irrelevant to a particular disease. Nevertheless, we have decided against such a pre-processing step, because it demands for a manual analysis of the data by domain experts. When dealing with large amounts of data this process can become very time consuming and must be repeated for different data sources and diseases. Instead, we aim to develop a tool that helps experts in the process of finding syndrome definitions and keeps the need for manual inspection of the data at a low level. We therefore strive for a machine learning method that is able to deal with different types of data sources and works independently of any particular disease without the need for costly pre-processing techniques. Ideally, it should be able to identify patterns that are most relevant to a particular disease on its own. Compared to a manual analysis of the underlying data source, an inspection of the resulting syndrome definitions, which we consider indispensable to testify the correctness of syndromes and to identify issues that may result from noise or other anomalies in the data, is less complex and time consuming. However, we focus on the conceptual and algorithmic fundamentals of a data-driven approach to syndromic surveillance in our experiments and leave the discussion of how to incorporate feedback that may be provided by domain experts into the learning procedure to the analysis of opportunities and limitations in section 5.

4.1.2. Number of Infections

The number of cases corresponding to the infectious diseases Influenza, Norovirus and SARS-CoV-2 have been retrieved from the *SurvStat*² platform. It is provided by the *Robert Koch-Institut*, which is a German federal government agency and research institute that is responsible for disease control and prevention. It maintains a database of cases notifiable diseases, reported under the German "Act on the Prevention and Control of Infectious Diseases in Man" ("Infektionsschutzgesetz.") The *SurvStat* platform allows to retrieve aggregated data from this database. Its use comes with limited control over the temporal and spatial aggregation of the data, which is only available at a weekly basis and is aggregated across German districts ("Landkreise" and "Stadtkreise.") To match the information in the health-related dataset, we use the weekly reported numbers for the districts where the considered emergency departments are located.

4.1.3. Parameter Setting

For all experiments that are discussed in the following, we have set the minimum support to $s = 0.0001$. With respect to the approximately 1,900,000 instances contained in the training dataset, this means that each conjunction of indicators considered by the algorithm must apply to at least 190 patients. The parameter s is necessary, because we assess the quality of syndromes in terms of their correlation with the ground truth, rather than taking the absolute number of estimated infections into account. By enforcing a minimum support, conjunctions of

²<https://survstat.rki.de>

indicators that apply to very few instances are discarded. On the one hand, this reduces the training time, as infrequent indicators can be ignored. On the other hand, this ensures that the syndromes apply to minimum fraction of the training instances and therefore are more general. This is necessary, because we assess the quality of syndromes in terms of their correlation to the ground truth, rather than taking the number of estimated infections into account. Larger values for s restrict the freedoms of the algorithm and may prevent it from learning syndromes that correlate with the ground truth if chosen too restrictively. In preliminary experiments, we have found a minimum support of $s = 0.0001$ to produce reasonable results, while keeping the training time at an acceptable level (typically under 1 min). In addition, we have limited the maximum number of disjunctions in a model to $L = 50$, which ensures that complex patterns can potentially be learned. However, the algorithm usually terminates before this threshold is reached.

4.2. Reconstruction of Synthetic Syndromes

In our first experiment, we validated the ability of our algorithm to discover disease patterns under the assumption that the reported cases are actually present in the data. For this purpose, we defined synthetic syndromes with varying characteristics from the emergency department data. For each syndrome, we determined the number of instances they apply to over time. The goal of the algorithm was to reconstruct the original syndrome definitions, exclusively based on the correlation with the corresponding number of cases. For this experiment, we focused on syndromes that use ICD codes and MTS representations, since these indicators are most commonly used in related work (e.g., Ivanov et al., 2002; Suyama et al., 2003; Boender et al., 2021). We have not used short versions of the ICD codes due to their overlap with the full codes. The following three different types of synthetic syndromes were considered:

1. Conjunctions of indicators (AND):

$$r = c_1 \wedge \dots \wedge c_M, \text{ where } M \in \{2, 3\}$$

2. Disjunctions of indicators (OR):

$$r_1 \vee \dots \vee r_L, \text{ where } r_l = c \text{ and } L \in [2, 9]$$

3. Disjunctions of conjunctions (AND-OR):

$$r_1 \vee \dots \vee r_L, \text{ where } r_l = c_1 \wedge c_2 \text{ and } L \in [2, 5]$$

For each syndrome type, we generated 100 artificial definitions by randomly selecting indicators that are present in the data, such that each indicator and each conjunction of indicators applies to at least 200 patients. This ensures that the syndromes that are ultimately generated apply to this particular number of patients at minimum. In addition, we have considered three temporal granularities to determine the number of cases different syndromes apply to. Experiments have been conducted with counts that are available on a daily, weekly, or monthly basis. To quantify to which extent our approach is able to reconstruct

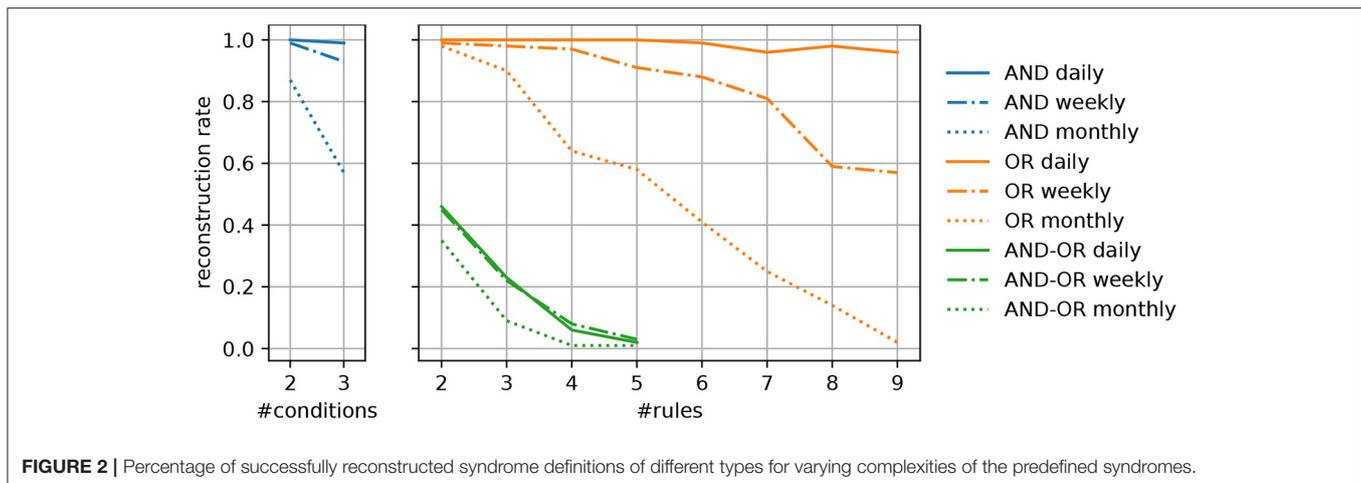
the original syndrome definitions, we compute the percentage of correctly identified patterns, i.e., syndromes that use the exact same indicators, referred to as the *reconstruction rate*. A visualization of the experimental results is given in **Figure 2**.

Generally, we can observe that the algorithm's ability to capture the predefined disease patterns benefits from a more fine-grained granularity of the available data (e.g., daily instead of weekly reported numbers). This meets our expectations, as a greater temporal resolution results in more specific patterns of covered cases, given a particular syndrome. As a result, it is easier to identify the indicators that allow to replicate a certain disease pattern and separate them from unrelated ones. In particular, syndromes that are exclusively based on disjunctions (OR) or conjunctions (AND), regardless of their complexity, can reliably be captured when supplied with daily numbers. When dealing with a broader temporal granularity, the uniqueness of disease patterns vanishes and they become more likely to interfere with the numbers resulting from similar syndromes.

Regarding the different types of predefined syndromes, it can be seen that their reconstruction becomes more difficult as their complexity increases. Especially when dealing with syndromes that include both, disjunctions and conjunctions (AND-OR), the reconstruction rate mostly depends on the number of indicators, whereas the temporal resolution plays a less important role. On the one hand, complex syndrome definitions apply to fewer training instances, which increases the imbalance of the learning task at hand. To overcome the practical limitations that result from imbalanced data, techniques for over- and undersampling are commonly used in machine learning. However, such techniques demand for the availability of labeled data. On the other hand, the results do also show the limitations of a greedy hill climbing strategy when it comes to the reconstruction of complex patterns. To overcome this shortcoming, approaches for the re-examination of previously induced rules, such as pruning techniques, could be considered. It is also possible to extend the search space that is explored by the training algorithm, e.g., by conducting a beam search, where several promising solutions are explored instead of focusing on a single one at each step. However, if the patterns, which have been found by the algorithm, only slightly differ from the predefined syndromes (e.g., by omitting or including infrequent ICD codes). While we did not evaluate this in depth, we believe they could still comprise useful information, e.g., by providing alternative, but nearly equivalent, descriptions of the syndrome.

4.3. Discovery of Syndrome Definitions From Real-World Data

In our second experiment, we used the proposed algorithm to obtain syndrome definitions for the infectious diseases Influenza, Norovirus, and SARS-CoV-2. In the literature, the quality of syndromes is either evaluated by experts (e.g., Ivanov et al., 2002; Heffernan et al., 2004; Lall et al., 2017; Bouchouar et al., 2021) or by measuring the correlation with reported infections, reported deaths or expert definitions (e.g., Suyama et al., 2003; Edge et al., 2006; Velardi et al., 2014; Muchaal et al., 2015; Nolan et al., 2017; Kalimeri et al., 2019). We follow the latter approach by



reporting the Pearson correlation coefficient of the automatically discovered disease patterns with the publicly reported number of infections supplied for training, as well as syndromes that have been handcrafted by ourselves. In addition, we provide a detailed discussion of the indicators included in our models.

Inspired by the expert syndrome definitions for Influenza and SARS-CoV-2 used by Boender et al. (2021), we created a set of similar, but much simpler, definitions solely based on ICD codes. They incorporate the ICD codes that correspond to suspected or confirmed cases of a particular disease, i.e., *J10* (Influenza due to identified seasonal influenza virus) or *J11* (Influenza, virus not identified) for Influenza, *A08* (viral and other specified intestinal infections) for Norovirus and *U07.1* (COVID-19, virus identified) or *U07.2* (COVID-19, virus not identified) for SARS-CoV-2. We have found the number of cases, these ICD codes apply to, to be very similar to those matched by the aforementioned expert definitions.

For each of the considered diseases, we trained several models using different sets of features. First of all, for a fair comparison with the handcrafted syndromes, we provided our algorithm with the features that belong to the first category in **Table 1**, i.e., ICD codes and MTS representations. A visualization of the number of infections that correspond to the disease patterns that have been discovered with respect to these features is shown in **Figure 3**. Each one of them includes a comparison with the reported number of infections supplied for training and the number of cases our handcrafted syndromes apply to, respectively. Note, that the numbers that correspond to the syndrome definitions are generally much lower than the reported numbers, as only a small fraction of detected cases have actually been treated in emergency departments. A detailed discussion of the discovered disease patterns is provided in the following section.

In addition to ICD codes and MTS representations, we have also conducted experiments, where we provided the algorithm with one additional set of features, as well as with all features available. To validate whether the availability of additional features comes with an advantage for an accurate reproduction of the infected cases, we rely on the Pearson correlation coefficients that result from different feature selections in **Table 2**. For all

experiments, we report the correlation of the autonomously learned syndromes with both, the number of reported cases used for training and the cases captured by the handcrafted syndromes. In the case of Influenza and SARS-CoV-2, the inclusion of vital parameters introduces a minor advantage for matching the reported numbers. Understandably, the use of additional features typically reduces correlation with the handcrafted syndromes, as they do not make use of these features. In the case of Norovirus, the Pearson correlation does not benefit from the availability of vital parameters. Regardless of any specific disease, this does also apply to the contextual and demographic information. We consider the absence of demographic indicators as positive, as none of the diseases appears to be specific to gender or age.

4.4. Discussion of Discovered Syndrome Definitions

As the use of ICD codes and MTS representations is sufficient in most cases to match the reported number of infections, we mostly focus on models that have been trained with respect to these features in the following discussion. A selection of exemplary syndromes that have been learned by our algorithm is shown in **Table 3**.

4.4.1. Influenza

The disease pattern that has been learned by our algorithm for modeling Influenza strongly correlate with both, the reported number of infections and the numbers that result from our handcrafted syndromes. In **Figure 3**, one can clearly observe an increase of infections during the first months of each year. The indicators that have been selected by our algorithm for modeling the number of Influenza cases include the ICD codes *J10* and *J11* that are also included in our handcrafted definition. These indicators have been selected during the first iterations of the algorithm and therefore are considered more important than the subsequent ones. As indicated by using different shades of blue in **Figure 3**, patterns found during early iterations (dark blue) mostly focus on the strongly pronounced seasonal peaks. Indicators that have been selected at later iterations (lighter blue)

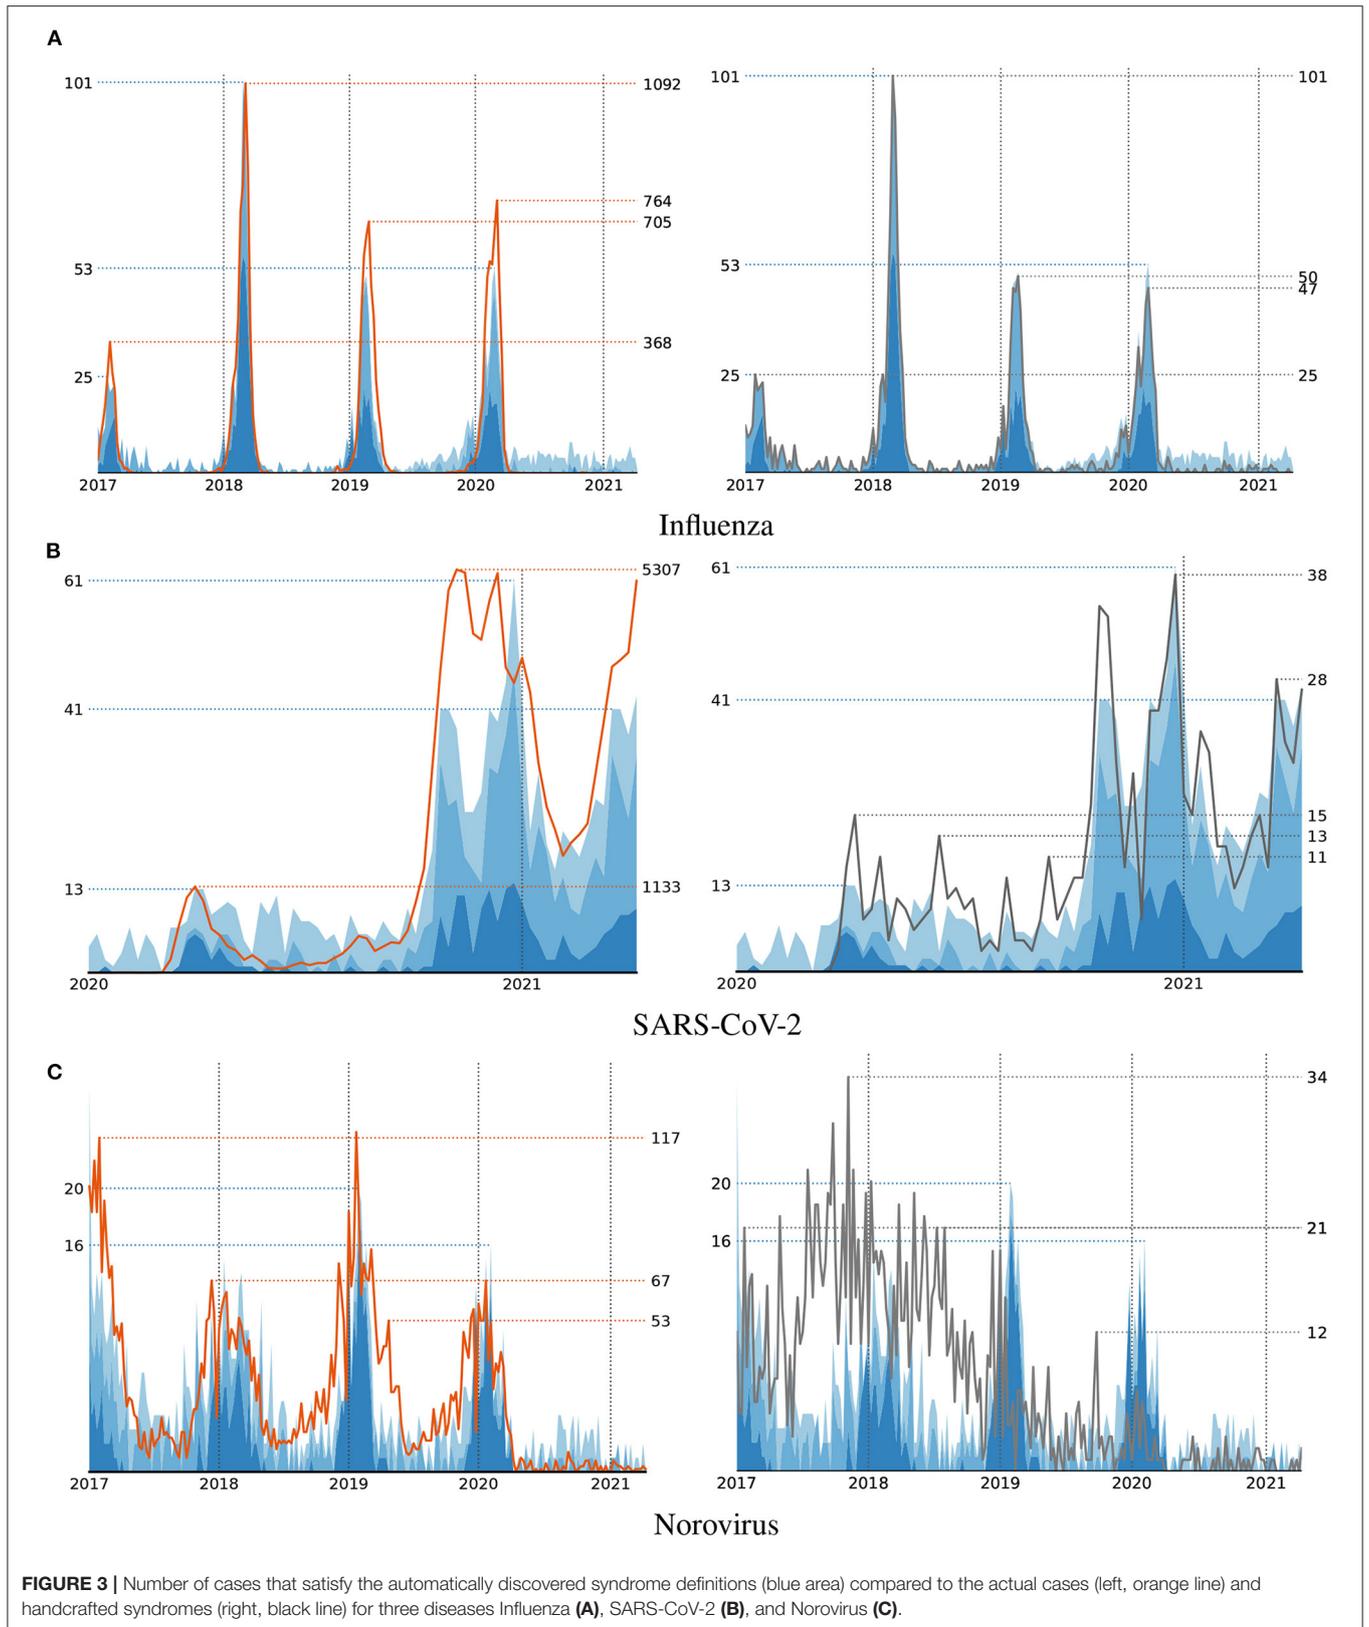


TABLE 3 | Exemplary automatically induced syndrome definitions.

① Influenza
J10 ∨ J11 ∨ "new confusion condition" ∨ Z96.0 ∨ ...
① SARS-CoV-2
(J12 ∧ "breathing problems") ∨ U07.1 ∨ "pain in lower abdomen" ∨ ...
① Norovirus
J21.0 ∨ D40 ∨ (J34 ∧ "recent problem")
① ③ Influenza
J10
∨ (J11 ∧ diastolic ≤ 92.5 ∧ systolic ≤ 156.5 ∧ temperature > 38.5)
∨ (temperature ≤ 40.5 ∧ diastolic ≤ 108.5 ∧ systolic ≤ 162 ∧ 187.5 ≤ heart rate ≤ 207.5)
∨ ...
① ② ③ ④ Influenza
J10
∨ (J11 ∧ diastolic ≤ 92.5 ∧ systolic ≤ 156.5 ∧ temperature > 38.5)
∨ (temperature ≤ 40.5 ∧ diastolic ≤ 110 ∧ systolic ≤ 162 ∧ 187.5 ≤ heart rate ≤ 212.5)
∧ no isolation ∧ patient sent home)
∨ ...

D40, Neoplasm of uncertain/unknown behaviour of male genital organs; J21.0, Acute bronchiolitis due to respiratory syncytial virus; J10, Influenza due to identified seasonal influenza virus; J34, Other disorders of nose and nasal sinuses; J11, Influenza, virus not identified; U07.1, COVID-19, virus identified; J12, Viral pneumonia, not elsewhere classified; Z96.0, Presence of urogenital implants.

are more likely to match irrelevant cases and hence are often unrelated to the respective disease. In the case of Influenza, this includes clearly irrelevant ICD codes, such as *Z96.0* (presence of urogenital implants) or *S53.1* (dislocation of elbow, unspecified) as fourth and fifth indicator, but also codes that may be related to Influenza-like illnesses, such as *J18.8* (other pneumonia) or *J34.2* (deviated nasal septum) at positions 10 and 15. When the algorithm has access to vital parameters, the indicator *J11* is combined with information about blood pressure and body temperature as follows:

$$\begin{aligned} J11 \wedge \text{blood pressure diastolic} \leq 92.5 \\ \wedge \text{blood pressure systolic} \leq 156.5 \\ \wedge \text{temperature} > 38.5 \end{aligned}$$

Due to the lack of domain knowledge, we are not qualified to decide whether such a pattern is in fact characteristic of Influenza. However, it shows the demand for experts, who are indispensable for the evaluation of machine-learned models and may use such a pattern as a starting point for a more detailed analysis of the underlying data.

4.4.2. SARS-CoV-2

Similar to Influenza, the number of infections with SARS-CoV-2 according to the ground truth, the handcrafted syndromes and the machine-learned definitions are strongly correlated. As seen in **Figure 3**, the different peaks of SARS-CoV-2 infections according to the publicly reported numbers are replicated by both, the handcrafted syndromes and the automatically learned patterns. When used to learn patterns for SARS-CoV-2, our

algorithm considers the MTS presentation "breathing problem," as well as the ICD codes *J12* (viral pneumonia) and *U07.1* (COVID-19, virus identified), as most relevant. The latter of these ICD codes is also included in the handcrafted syndrome definition. Besides clearly irrelevant indicators, it further selects the ICD code *J34.2* (deviated nasal septum) at a later stage of training that may be related to this particular illness. When provided with vital parameters, the algorithm decides to use the ICD code *J12* in combination with data about a patient's blood pressure and temperature in its most relevant pattern:

$$\begin{aligned} J12 \wedge 81.5 < \text{blood pressure systolic} \leq 149.5 \\ \wedge \text{blood pressure diastolic} \leq 77.5 \\ \wedge \text{temperature} > 36.5 \end{aligned}$$

4.4.3. Norovirus

As depicted in **Figure 3**, the correlation between syndromes and reported numbers is less strong with respect to Norovirus. However, compared to the handcrafted syndromes, the automatically discovered patterns appear to better capture the seasonal outbreaks of this particular disease. Unfortunately, the algorithm fails to identify any ICD codes that are related to this particular illness, such as the ones included in our manual definition or codes related to symptoms like diarrhea. Instead, it uses indicators like *J21.0* (Acute bronchiolitis due to respiratory syncytial virus) or *J34* (Other disorders of nose and nasal sinuses) in combination with other indicators to match the reported numbers. This is most probably due to the similar seasonality of Norovirus and Influenza-like illnesses. This illustrates another difficulty one may encounter when pursuing a data-driven approach to syndromic surveillance. If high numbers of infections with respect to multiple diseases occur during a similar timespan, the algorithm is not able to distinguish between indicators that relate to different types of infections. In such case it is necessary to provide additional knowledge to the learning algorithm, as it is unable to grasp the semantics of individual features on its own. In particular, this motivates the need for an interactive learning approach, where a human expert interacts with the computer in order to guide the construction of models. For example, by prohibiting the use of certain indicators or features that have been identified to be irrelevant to the problem at hand.

5. DISCUSSION AND LIMITATIONS

Our experimental evaluation using both, synthetic and real-world data, provided several insights into the problem domain addressed in this work. First of all, we were able to demonstrate that a correlation-based learning approach for the extraction of disease patterns is indeed capable of identifying meaningful indicators that are closely related to a particular disease under surveillance. In particular, the learned definitions showed a similar fit to the real distributions as handcrafted expert definitions (**Figure 3**). Also, the experiments with synthetic syndrome definitions showed a good reconstruction rate, and the discovered real-world syndrome definitions contained plausible features.

Nevertheless, the frequent inclusion of unrelated indicators revealed some challenges and limitations of such an approach. Most of them relate to the fact that the training procedure has only limited access to the target information associated with each patient. In contrast to fully labeled data, where information about each patient's medical conditions are available, the learning method is restricted to broad information about a large group of individuals. In addition, the use of temporally aggregated data, depending on its granularity, introduces ambiguity into the learning process. As a result of these constraints, several solutions that satisfy the evaluation criterion to be optimized by the learner exist, even though many of them are undesirable from the perspective of domain experts. This is evident from the fact that the tested algorithm, regardless of the disease and the features used for training, was always able to find strongly correlated patterns, despite the use of unrelated indicators. As another source of problems, we identified the noise, potential inconsistencies and missing pieces of information that may be encountered when dealing with unprocessed and unfiltered real-world data. The consequences become most obvious when taking a look at the results with respect to Norovirus, where the algorithm failed to detect meaningful syndrome descriptions due to the overlap to other, more frequent, diseases with a similar seasonality and more pronounced patterns in the reported numbers.

So far, we were only interested in the identification of patterns that match the target variables as accurate as possible. However, the goal of machine learning approaches usually is to obtain predictions for unseen data. To be able to generalize well beyond the provided training data, this requires models to be resistant against noise and demands for techniques that effectively prevent overfitting. The incorporation of such techniques into our learning approach may improve its ability to find useful patterns despite the noise and ambiguities that are present in the data. For example, successful rule learning algorithms often come with pruning techniques that aim at removing problematic clauses from rules after they have been learned. This requires to split up the training data into multiple partitions in order to be able to obtain unbiased estimates of a rule's quality, independent of the data used for its induction. By splitting up the time series data, the quality of indicators that are taken into account for the construction of syndromes could more reliably be assessed in terms of multiple, independent estimates determined on different portions of the data. Despite such technical solutions, we believe that the active participation of domain experts is indispensable for the success of machine-guided syndromic surveillance. An interactive learning approach, where the syndromes that are discovered by an algorithm are suggested to epidemiologists and feedback is fed back into the system, may prevent the inclusion of undesired patterns and would most likely help to increase the acceptance of machine learning methods among healthcare professionals.

Furthermore, we consider the use of the Pearson correlation coefficient as a limitation of our approach. When modeling the outbreak of a disease, it is especially important to properly reflect the points in time that correspond to high numbers of infections. Other correlation measures, like weighted variants

of the Pearson correlation coefficient, may provide advantages in this regard. We expect this aspect to be particularly relevant when modeling rather infrequent diseases with generally low incidences. Another problem are possible discrepancies between the data obtained from the emergency departments and the data that incorporates information about the number of infections, e.g., resulting from reporting delays. To circumvent potential issues that may result from such inconsistencies, approaches that have specifically been designed for measuring the similarity between temporal sequences, like dynamic time warping (Müller, 2007), could be used in the future. They allow for certain static, and even dynamic, displacements of the sequences to compare.

6. CONCLUSION

In this work, we have presented a novel approach for the automatic induction of syndrome definitions from health-related data sources. As it aims at finding patterns that correlate with the reported numbers of infections, as provided by publicly available data sources, there is no need for labeled training data. This reduces the burdens imposed on domain experts, who otherwise must manually create labeled data in a laborious and time consuming process. Although the proposed algorithm is able to identify meaningful indicators, due to artifacts in the data and technical limitations, we have found that autonomously created syndromes are likely to include indicators that are unrelated to the disease under surveillance. As a result, the knowledge of experts is still indispensable for the evaluation and supervision of such a machine learning method. Nevertheless, our investigation shows the potential of data-driven approaches to syndromic surveillance, due to their ability to process large amounts of data that cannot fully be understood and analyzed by humans.

In the future, we plan to investigate technical improvements to our algorithm that may help to prevent overfitting and allow for a more extensive, yet computationally efficient, exploration of promising combinations of indicators. In addition, valuable insights can possibly be obtained by applying our approach to different types of health-related data sources, as well as by the investigation of different correlation measures that can potentially be used to guide the search for meaningful syndromes.

DATA AVAILABILITY STATEMENT

Our source code is publicly available at <https://github.com/mrapp-ke/SyndromeLearner>. The reported numbers of infections used in this work are provided by the Robert-Koch-Institut for public access at <https://survstat.rki.de>. The emergency department data cannot be published due to privacy protection.

AUTHOR CONTRIBUTIONS

EL, MR, and MK developed the basic concept and methodology for the presented research. MR was responsible for the software implementation and carried out the experiments using real-word

data. MK conducted the empirical study using synthetic data. All authors have discussed the results and collaborated on the manuscript.

FUNDING

Parts of this work were funded by the German Innovation Committee of the Federal Joint Committee (G-BA) [ESEG project, grant number 01VSF17034]. We further acknowledge support by the Deutsche Forschungsgemeinschaft (DFG - German Research Foundation) and the Open Access Publishing Fund of the Technical University of Darmstadt.

REFERENCES

- Barro, R. J., Ursúa, J. F., and Weng, J. (2020). *The Coronavirus and the Great Influenza Pandemic: Lessons From the 'Spanish Flu' for the Coronavirus'S Potential Effects on Mortality and Economic Activity*. Technical Report, National Bureau of Economic Research, Cambridge, MA.
- Bjerring, L., and Frank, E. (2011). "Beyond trees: adopting MITI to learn rules and ensemble classifiers for multi-instance data," in *Proc. Australasian Joint Conference on Artificial Intelligence* (Perth, WA), 41–50.
- Boender, S. T., Cai, W., Schranz, M., Kocher, T., Wagner, B., Ullrich, A., et al. (2021). Using routine emergency department data for syndromic surveillance of acute respiratory illness before and during the COVID-19 pandemic in Germany, week 10-2017 and 10-2021. *medRxiv*.
- Bouchouar, E., Hetman, B. M., and Hanley, B. (2021). Development and validation of an automated emergency department-based syndromic surveillance system to enhance public health surveillance in Yukon: a lower-resourced and remote setting. *BMC Public Health* 21, 1–13. doi: 10.1186/s12889-021-11132-w
- Carbonneau, M.-A., Cheplygina, V., Granger, E., and Gagnon, G. (2018). Multiple instance learning: a survey of problem characteristics and applications. *Pattern Recognit.* 77, 329–353. doi: 10.1016/j.patcog.2017.10.009
- Chevalyere, Y., and Zucker, J.-D. (2001). "A framework for learning rules from multiple instance data," in *Proc. European Conference on Machine Learning (ECML)* (Freiburg), 49–60.
- Edge, V. L., Pollari, F., King, L., Michel, P., McEwen, S. A., Wilson, J. B., et al. (2006). Syndromic surveillance of norovirus using over the counter sales of medications related to gastrointestinal illness. *Can. J. Infectious Diseases Med. Microbiol.* 17, 235–241. doi: 10.1155/2006/958191
- Fürnkranz, J., Gamberger, D., and Lavrač, N. (2012). *Foundations of Rule Learning*. (Berlin; Heidelberg: Springer Science & Business Media).
- Goldstein, E., Cowling, B. J., Aiello, A. E., Takahashi, S., King, G., Lu, Y., et al. (2011). Estimating incidence curves of several infections using symptom surveillance data. *PLoS One* 6, 1–8. doi: 10.1371/journal.pone.0023380
- Gräff, I., Goldschmidt, B., Glien, P., Bogdanow, M., Fimmers, R., Hoef, A., et al. (2014). The german version of the Manchester triage system and its quality criteria – first assessment of validity and reliability. *PLoS One* 9:e88995. doi: 10.1371/journal.pone.0088995
- Hartnett, K. P., Kite-Powell, A., DeVies, J., Coletta, M. A., Boehmer, T. K., Adjemian, J., et al. (2020). Impact of the COVID-19 pandemic on emergency department visits United states, January 1, 2019–May 30, 2020. *Morb. Mortal. Week. Rep.* 69, 699–704. doi: 10.15585/mmwr.mm6923e1
- Heffernan, R., Mostashari, F., Das, D., Karpati, A., Kulldorff, M., and Weiss, D. (2004). Syndromic surveillance in public health practice, New York City. *Emerg. Infect. Dis.* 10, 858–864. doi: 10.3201/eid1005.030646
- Henning, K. J. (2004). What is syndromic surveillance? *Morbidity Mortality Week. Rep.* Druid Hills, GA: Centers for Disease Control & Prevention (CDC). 53, 7–11.
- Hopkins, R. S., Tong, C. C., Burkom, H. S., Akkina, J. E., Berezowski, J., Shigematsu, M., et al. (2017). A practitioner-driven research

ACKNOWLEDGMENTS

We thank our project partners, namely the *Health Protection Authority of Frankfurt*, the *Hesse State Health Office and Centre for Health Protection*, the *Hesse Ministry of Social Affairs and Integration*, the *Robert Koch-Institut*, the *Epias GmbH*, and the *Sana Klinikum Offenbach GmbH*, who provided insight and expertise that greatly assisted the research. Especially, we thank Theresa Kocher, Birte Wagner, and Sonia Boender from the *Robert-Koch-Institut* for providing the data used for evaluation and for their valuable feedback that substantially improved the manuscript.

- agenda for syndromic surveillance. *Public Health Rep.* 132, 116–126. doi: 10.1177/0033354917709784
- Ivanov, O., Wagner, M. M., Chapman, W. W., and Olszewski, R. T. (2002). "Accuracy of three classifiers of acute gastrointestinal syndrome for syndromic surveillance," in *Proc. AMIA Symposium* (San Antonio, TX), 345–349.
- Kalimeri, K., Delfino, M., Cattuto, C., Perrotta, D., Colizza, V., Guerrisi, C., et al. (2019). Unsupervised extraction of epidemic syndromes from participatory influenza surveillance self-reported symptoms. *PLOS Comput. Biol.* 15:e1006173. doi: 10.1371/journal.pcbi.1006173
- Lall, R., Abdelnabi, J., Ngai, S., Parton, H. B., Saunders, K., Sell, J., et al. (2017). Advancing the use of emergency department syndromic surveillance data, New York City, 2012–2016. *Public Health Rep.* 132, 23S–30S. doi: 10.1177/0033354917711183
- Mandl, K. D., Overhage, M. J., Wagner, M. M., Lober, W. B., Sebastiani, P., Mostashari, F., et al. (2004). Implementing syndromic surveillance: a practical guide informed by the early experience. *J. Amer. Med. Informat. Assoc.* 11, 141–150. doi: 10.1197/jamia.M1356
- Muchaal, P. K., Parker, S., Meganath, K., Landry, L., and Aramini, J. (2015). Big data: evaluation of a national pharmacy-based syndromic surveillance system. *Canada Commun. Disease Rep.* 41, 203. doi: 10.14745/ccdr.v41i09a01
- Müller, M. (2007). "Dynamic time warping," in *Information Retrieval for Music and Motion* (Berlin; Heidelberg: Springer), 69–84.
- Nolan, M. L., Kunins, H. V., Lall, R., and Paone, D. (2017). Developing syndromic surveillance to monitor and respond to adverse health events related to psychoactive substance use: methods and applications. *Public Health Rep.* 132, 65S–72S. doi: 10.1177/0033354917718074
- Rapp, M., Loza Mencía, E., Fürnkranz, J., Nguyen, V.-L., and Hüllermeier, E. (2020). "Learning gradient boosted multi-label classification rules," in *Proc. European Conference on Machine Learning and Knowledge Discovery in Databases (ECML-PKDD)* (Ghent), 124–140.
- Ray, S., and Page, D. (2001). "Multiple instance regression," in *Proc. International Conference on Machine Learning (ICML)* (Williamstown, MA), 425–432.
- Suyama, J., Sztajnkrzyer, M., Lindsell, C., Otten, E. J., Daniels, J. M., and Kressel, A. B. (2003). Surveillance of infectious disease occurrences in the community: An analysis of symptom presentation in the emergency department. *Acad. Emergency Med.* 10, 753–763. doi: 10.1111/j.1553-2712.2003.tb00070.x
- Trott, P. A. (1977). International classification of diseases for oncology. *J. Clin. Pathol.* 30, 782.
- Velardi, P., Stilo, G., Tozzi, A. E., and Gesualdo, F. (2014). Twitter mining for fine-grained syndromic surveillance. *Artif. Intell. Med.* 61, 153–163. doi: 10.1016/j.artmed.2014.01.002

Conflict of Interest: The authors MR, MK, and EL have received funding by the German Innovation Committee of the Federal Joint Committee (G-BA).

The remaining author declares that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may

be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Copyright © 2022 Rapp, Kulesa, Loza Mencía and Fürnkranz. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.