Predicting Ankylosing Spondylitis Treatment Outcome Using Model-based Methods

Alexandre Almeida Abreu
alexandre.abreu@tecnico.ulisboa.pt
Instituto Superior Técnico, Universidade de Lisboa, Lisbon, Portugal

Ankylosing spondylitis (AS) is a chronic autoimmune inflammatory disease that primarily affects the axial skeleton. If inadequately treated, the disease causes severe pain and spinal deformity, leading to overall reduced quality of life. The treatment protocols are flexible, allowing for many different personalized adjustments. Therefore, there is great benefit in accurately predicting how a specific patient will respond to each therapy. To fulfil such task, model-based methods that rely on probabilistic graphical models are reviewed and applied to data from REUMA.PT database, responsible for collecting data from patients with rheumatic diseases in Portugal. The results show that it is sometimes possible to predict how patients will respond to biological therapy. Furthermore, it was discovered that specific features relative to the state of a patient could help in the choice of treatment. In addition, some explainable correlations between variables were found. Last but not least, it was uncovered that certain features present in REUMA.PT are not as relevant in this prediction as the rest.

***Index Terms*—Ankylosing spondylitis, time series, model-based, dynamic Bayesian network, dynamic Boltzmann machine.**

I. INTRODUCTION

ANKYLOSING SPONDYLITIS (AS) is a chronic autoimmune inflammatory disease with a strong predilection for the axial skeleton, but it can also affect the peripheral joints and organs namely, the eyes, the heart and the lungs [1]. If inadequately treated the disease causes severe pain and spinal deformity, and in extreme cases it can lead into a 'bamboo spine', where a total fusion of the axial skeleton occurs, resulting in immobility and restrictive lung function, leading to respiratory failure, and overall reduced quality of life.

A. Motivation

There is a benefit to both, the patient and the government, to adequately treat patients with AS, as there is a socio-economic burden since patients may struggle to work. Not only because of their symptoms but also because the workplace might not be adjusted for people with arthritis.

Since there is no cure for this disease, the main objectives of treatment are to reduce pain, avoid irreversible deformities, and preserve a good articular function in order to increase the patients' quality of life. Non-steroidal anti-inflammatory drugs (NSAID) have been widely used by patients with AS, they show effectiveness in reducing pain and rigidity, however, they do not allow for suitable control of the disease progression. More recently, biological therapies are being used, acting as TNF inhibitors, and have exhibited to intervene in the disease progression [2]. However, the treatment protocols are flexible, allowing for many different personalized adjustments. Therefore, there is great benefit in being able to accurately predict how a specific patient will respond to each therapy.

The data used in this study focuses on patients with AS under biological therapy and was extracted from the rheumatic diseases Portuguese register (RDPR) available at REUMA.PT [3]. This database has been actively collecting information on patients with rheumatic diseases in Portugal since June 2008. It contains over 2000 patients and more than 45000 appointments of patients with AS.

B. Objectives

In a general way, the main objective is to use model-based machine learning methods to analyse data from REUMA.PT of patients with AS. On a certain patient a biological therapy can be considered effective if there is no need to switch to a different one. Therefore, one objective is to be able to accurately predict if a given therapy will be effective, or in other words, predict which therapy would be the most beneficial to a certain patient. In addition, we would like to uncover what features have an impact on the effectiveness of the different biological therapies used to treat patients with AS.

To do so, model-based machine learning methods will be reviewed and applied. More precisely, dynamic models will be investigated and used, since this type of models are particularly adequate for time series data.

The outcome we want to accurately predict has to be created from the available data. As a result a ready-to-use preprocessed dataset will be created.

II. MODELS

Two dynamic models were selected for use with data from REUMA.PT. More specifically, the dynamic Bayesian network (DBN) and the dynamic Boltzmann machine (DyBM), which extends the Hebb rule to consider the concept of time, were used.

A. Dynamic Bayesian network

A dynamic Bayesian network is built on a Bayesian network (BN), allowing to model temporal processes [4]. Such

models have two different types of relations to learn, intra and inter-temporal. Intra-temporal relations express the dependencies between variables at a given time instance, while inter-temporal relations model how variables at times 0 to t influence the variables at time $t+1$.

A DBN is composed of a prior network, B^0, that describes the joint probability distribution over the initial state $\mathbf{X}[0]$, and a set of transition networks, B_t^{t+1}, describing the distribution over the variables at time t given the variables up to time $t-1$, for $0 \leq t < T$, that is,

$$B_t^{t+1} = P(\mathbf{X}[t] \| \mathbf{X}[: t-1]),$$

where $\mathbf{X}[t] = (X_1[t], \ldots, X_n[t])$ denotes the values of the variables at time t, and where $\mathbf{X}[: t-1]$ denotes the values of the variables up until time $t-1$. A common simplification of this model is to consider a mth-order Markov process, in which values at time t depend only on values from m timesteps prior to t,

$$B_t^{t+1} = P(\mathbf{X}[t] \| \mathbf{X}[t-m : t-1]),$$

where a first-order Markov process, $m = 1$, is often considered. Another simplification is to consider a stationary network, meaning that only one prior and transition network is used, where the transition network is unrolled over time. Meaning that B_t^{t+1} is the same for all $t \in [0, T-1]$.

In this work, the tree-augmented dynamic Bayesian network (tDBN) was used. The (tDBN) is a polynomial time complexity approach to a DBN [4]. This model finds both, optimal inter and intra-temporal relations in a transition network. To accomplish this, the model uses a tree network structure to model the intra-temporal relations where the number of parents from the preceding time-instances is restricted to at most p parents and only one parent from the current time-instance is allowed. This approach is polynomial in the number of attributes n, linear in the the number of observations N, and exponential in the number of parents p.

B. Dynamic Boltzmann machine

The Hebb rule used to define a Boltzmann machine (BM) has a limitation, the concept of time is missing [5], [6], therefore, such networks are not ideal for time series data. To overcome his limitation, the Hebb rule is extended by spike-timing dependent plasticity (STDP) [7], stating that a synaptic weight is strengthened when a post-synaptic neuron fires shortly after a pre-synaptic neuron (*i.e.*, long term potentiation (LTP)), and is decreased if the order of firing is reversed (*i.e.*, long term depression (LTD)).

The dynamic Boltzmann machine (DyBM) [5], [6] is trained with a learning rule that has some of the properties of STDP. The DyBM consists of a BM having multiple layers each with N units or neurons having memory units and conduction delays, where each layer corresponds to a different moment in time. The structure of a connection between a pre-synaptic neuron and a post-synaptic neuron is illustrated in Fig. 1.

Let $x_j^{[t]}$ denote the value of the j-th neuron at time t for $j \in [1, N]$. A pre-synaptic neuron, $i \in [1, N]$ may be connected to a post-synaptic neuron, $j \in [1, N]$ with a first in, first out (FIFO) queue of length $d_{i,j} - 1$, where $d_{i,j}$ is the conduction delay from the pre-synaptic neuron i to the post-synaptic neuron j. This way, a spike[1] in a pre-synaptic neuron, i, reaches a post-synaptic neuron, j, though a synapse after a conduction delay, $d_{i,j}$. A neuron can be connected to itself via a FIFO queue.

The memory units of a neuron store the neural eligibility traces and the synaptic eligibility traces. Specifically, each neuron stores a fixed amount, L, of neural eligibility traces, which summarize the neuron's activities in the past. Let $\gamma_{j,l}^{[t-1]}$ denote the l-th neural eligibility traces of the j-th neuron just before time t, for $l \in [1, L]$,

$$\gamma_{j,l}^{[t-1]} \equiv \sum_{s=-\infty}^{t-1} \mu_l^{t-s} x_j^{[s]},$$

where $\mu_l \in (0, 1)$ is the decay rate for the l-th neural eligibility trace. So, the neural eligibility trace is the weighted sum of the past values of the neuron, in which recent values have greater weights than older ones.

The number of synaptic eligibility traces depends on the number of neurons that are connected to that neuron. For each pre-synaptic neuron i that connects to a given post-synaptic neuron j, neuron j stores a fixed amount, K, of synaptic eligibility traces. Let $\alpha_{i,j,k}^{[t-1]}$ denote the k-th synaptic eligibility trace of neuron j for pre-synaptic neuron i just before time t, for $k \in [1, K]$,

$$\alpha_{i,j,k}^{[t-1]} \equiv \sum_{s=-\infty}^{t-d_{i,j}} \lambda_k^{t-s-d_{i,j}} x_i^{[s]},$$

where $\lambda_k \in (0, 1)$ is the decay rate for the k-th synaptic eligibility trace. Therefore, the synaptic eligibility trace is the weighted sum of the values that reached neuron j from a pre-synaptic neuron, i, after the conduction delay $d_{i,j}$.

Apart from these values, the DyBM has learnable parameters, bias and weight. Every neuron j has an associated bias, b_j, and each connected pair or neurons (synapse) is associated with weights of LTP and weights of LTD. Let $u_{i,j,k}$ denote the k-th LTP weight from a pre-synaptic neuron, i, to a post-synaptic neuron, j, corresponding to the k-th synaptic eligibility trace for $k \in [1, K]$. The LTD weight from a pre-synaptic neuron i to a post-synaptic neuron j is described by L parameters and is denoted by $v_{i,j,l}$ for $l \in [0, L]$, corresponding to the l-th neural synaptic trace.

Unlike a BM, the energy of a given pattern at a moment of a DyBM depends on the previously generated patterns. Let $\mathbf{x}^{[t]} = \left(x_j^{[t]}\right)_{j \in [1,N]}$ denote the vector of the values of the neurons at time t and $\mathbf{x}^{[:t-1]} = \left(\mathbf{x}^{[s]}\right)_{s<t}$ denote the values of the neurons up to time $t-1$.

The energy of the DyBM at time t is denoted by $E_\theta(\mathbf{x}^{[t]} \,|\, \mathbf{x}^{[:t-1]})$ and can be decomposed into the energies of the neurons at time t,

$$E_\theta(\mathbf{x}^{[t]} \,|\, \mathbf{x}^{[:t-1]}) = \sum_{j=1}^{N} E_{\theta,j}(x_j^{[t]} \,|\, \mathbf{x}^{[:t-1]}),$$

[1]When a neuron takes the value 1 we say that it spiked.

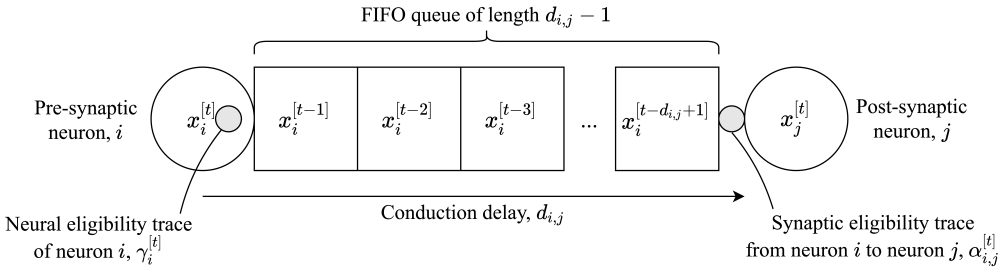

Fig. 1. Representation of the connection between two neurons in a DyBM. A pre-synaptic neuron is connected to a post-synaptic neuron with a FIFO queue. Resulting in spikes from the pre-synaptic neuron reaching the post-synaptic neuron after a constant conduction delay, $d_{i,j}$. Each neuron has memory units that store neural and synaptic eligibility traces. Adapted from [5].

where the energy of neuron j at time t is given by,

$$E_{\theta,j}(x_j^{[t]} \mid \mathbf{x}^{[:t-1]}) = -b_j x_j^{[t]} - \sum_{i=1}^{N}\sum_{k=1}^{K} u_{i,j,k}\alpha_{i,j,k}^{[t-1]} x_j^{[t]} + \sum_{i=1}^{N}\sum_{l=1}^{L} v_{i,j,l}\beta_{i,j,l}^{[t-1]} x_j^{[t]} + \sum_{i=1}^{N}\sum_{l=1}^{L} v_{j,i,l}\gamma_{i,l}^{[t-1]} x_j^{[t]}, \tag{1}$$

where $\beta_{i,j,l}^{[t-1]}$ is another eligibility trace,

$$\beta_{i,j,l}^{[t-1]} \equiv \sum_{s=t-d_{i,j}+1}^{t-1} \mu_l^{s-t} x_i^{[s]},$$

that is, the weighted sum of the values that are going to reach neuron j from a pre-synaptic neuron i within the period of conduction delay, $d_{i,j}$.

The first term in the right side of (1) translates that a neuron j having larger positive bias, b_j, is more likely to spike at time t. The second term corresponds to LTP. A post-synaptic neuron, j, is more likely to fire at time t, if the spikes from a pre-synaptic neuron, i, have recently and frequently arrived (*i.e.* $\alpha_{i,j,k}^{[t-1]}$ is large). The third term corresponds to LTD. Neuron j is less likely to fire when spikes will soon and frequently arrive from neuron i (*i.e.*, $\beta_{i,j,l}^{[t-1]}$ is large) and the corresponding LTD weight, $v_{i,j,l}$, is large. The fourth and last term also corresponds to LTD. While the third term only considers the spikes that are going to reach the post-synaptic neuron within the period of conduction delay, the last term considers the spikes that will arrive after the conduction delay. In this term neurons i and j switch roles. The pre-synaptic neuron, j, is less likely to spike at time t, if the post-synaptic neuron, i, has recently and frequently spiked (*i.e.*, $\gamma_{i,l}^{[t-1]}$ is large).

The probability that neuron j takes a specific value at time t, $x_j^{[t]} \in \{0,1\}$, depends on the values of the DyBM before time t, $\mathbf{x}^{[:t-1]}$, and is given by,

$$P_{\theta,j}(x_j^{[t]} \mid \mathbf{x}^{[:t-1]}) = \frac{\exp(-E_{\theta,j}(x_j^{[t]} \mid \mathbf{x}^{[:t-1]}))}{\sum_{\tilde{x}\in\{0,1\}} \exp(-E_{\theta,j}(\tilde{x} \mid \mathbf{x}^{[:t-1]}))}.$$

In a DyBM, the values of the neurons at time t are conditionally independent of each other, given the values before time t, $\mathbf{x}^{[:\mathbf{t-1}]}$, therefore, the probability of the values of the neurons at time t is given by,

$$P_\theta(\mathbf{x}^{[t]} \mid \mathbf{x}^{[:t-1]}) = \prod_{j=1}^{N} P_{\theta,j}(x_j^{[t]} \mid \mathbf{x}^{[:t-1]}).$$

Some simplifications are performed to this model. For simplicity, a constant conduction delay, d, and a single neural and synaptic eligibility trace (*i.e.* $L = 1$ and $K = 1$) are considered [8].

These simplifications allow to rewrite the energy of a pattern in a matrix form by,

$$E_\theta(\mathbf{x}^{[t]} \mid \mathbf{x}^{[:t-1]}) = -\left(\mathbf{b}^\mathsf{T} + \sum_{\delta=1}^{d-1}(\mathbf{x}^{[t-\delta]})^\mathsf{T}\mathbf{W}^{[\delta]} + \sum_{z=1}^{Z}(\boldsymbol{\gamma}_{\mu_z}^{[t-1]})^\mathsf{T}\mathbf{V}^{[z]}\right)\mathbf{x}^{[t]}, \tag{2}$$

where $\mathbf{b}$ is a vector of bias, $\mathbf{W}^{[\delta]}$ for $\delta \in [1, d-1]$ and $\mathbf{V}^{[z]}$ for $z \in [1, Z]$ are $N \times N$ matrices of weights, and $\boldsymbol{\gamma}_{\mu_z}^{[t-1]}$ for $z \in [1, Z]$ are column vectors,

The dependency on the decay rate, μ_z, in $\boldsymbol{\gamma}_{\mu_z}^{[t-1]}$, was made explicit by adding the respective subscript.

III. Rheumatological Dataset

The data used throughout the rest of this thesis was extracted from Reuma.pt, where the information is in regard to patients with AS under treatment by biological therapies. The dataset extracted contains a total of 2293 patients with a total of 45331 appointments.

The extracted dataset is subject to four preprocessing steps. To start, since there is a high amount of missing values, imputation is performed following two steps. First, for each feature, the last valid observation is propagated forward until the next valid observation. Next, the still missing values are filled with the value of the next valid observation. Note that, if a given patient does not have a single valid observation for a certain feature (only missing values), imputation is not possible. Therefore, missing values can and are still present.

Over the follow-up period, some patients switched biological drug at least once. Such switch in therapy means that the first biological drug used was not effective in treating the patient, or that other adverse event occurred. For this reason, the next step is to add a new feature to this dataset to express this event. So, every patient that has only been under

treatment with one biological drug is considered to have an effective response to that treatment. Meanwhile, for the other patients, treatment is considered to not be effective and further appointments after switching medication are ignored.

The third preprocessing step is to drop and truncate the patients data. Meaning that patients with less than 5 appointments are dropped to avoid having patients with a reduced number of appointments, and patients with more than 9 appointments are truncated to the first 9 appointments.

The last step is to perform discretization of the variables following TABLE I. In the end, the preprocessed datasets has 36 features, as present in TABLE I, and has 1290 patients and 10758 appointments in total. It is worth noting that there are still missing values present where imputation was not possible. This preprocessed dataset is ready for use in any application by further researchers.

IV. Methodology

Before applying this data to the selected models, it is necessary to prepare the data. Since the models do not work with missing values, and since feature VAS MEDIC has a high amount of such values, this feature is removed followed by patients with missing values.

Next, the dataset is separated by therapy used, resulting in 6 different datasets, where in each one, only patients under the same treatment are present. This separation is performed because each dataset is considered in a separate model. As seen in TABLE II, datasets $\mathtt{D}_{\mathtt{Certolizumab}}$ and $\mathtt{D}_{\mathtt{Secucinumab}}$ have a reduced number of patients. For this reason, these two datasets were ignored.

Lastly, each of the resulting datasets are split into two datasets. For example, $\mathtt{D}_{\mathtt{Adalimumab}}$ is separated into $\mathtt{D}^1_{\mathtt{Adalimumab}}$ and $\mathtt{D}^2_{\mathtt{Adalimumab}}$. On the first one, features relative to the BASDAI and BASFI questions are not considered, that is, features BASDAI Q1-6 and BASFI Q1-10 are removed. While for the second one, features BASDAI and BASFI are dropped but the respective questions are kept. The set of datasets following the first approach are collectively denoted by DATASET 1 while the rest are denoted by DATASET 2. The objective with DATASET 2 is to explore the impact of the different questions and possibly identify some that are less relevant in the prediction.

Relative to the DyBM, since this model only accept binary features, non-binary features have to be converted. To do so, each feature is encoded using a one-hot ('one-of-K' or 'dummy') scheme.

Every dataset used is separated into a training set and a test set. This separation is performed such that, 80% of the dataset is used to train the model and the remaining is left for testing. Furthermore, the separation is balanced in such a way that, the ratio of patients that stopped therapy due to inefficiency in the resulting subsets (train and test sets) is equal to the initial ratio of the dataset.

A. Training procedure

Given that the datasets are highly class imbalanced as shown in Table II, balanced datasets are also used to train the models. Two different methods are used to perform such task, where, in both situations, the objective is to have the same amount of samples in each class. That is, the number of patients that stopped therapy due to inefficiency (positive outcome) is the same as the number of patients that did not (negative outcome).

One method is by performing undersampling, in other words, by randomly removing patients from the majority class. Another is by oversampling using a technique called synthetic minority oversampling technique (SMOTE) [9]. The SMOTE implementation used is available on GitHub[2] [10].

In short, SMOTE generates synthetic samples of the minority class. To do this, the method takes a sample from the minority class and finds its k nearest neighbors. Then, to create a new synthetic sample, the vector between one of the neighbors and the sample is multiplied by a random number in the range of 0 to 1.

For the DyBM, the model is provided with data from the patients and it is considered as outcome whether the patients stopped therapy due to inefficiency. In other words, the input to the model is the patients' temporal data until switching drug or until the end of the follow-up period, whichever occurs first, and feature INEFFICIENCY is the output.

This training procedure of the DyBM follows closely the procedure carried out in [5], where the authors used this model to learn two different sequences. As a result, the data available for each patient is considered to be a separate sequence that is supplied to the network. In each iteration of the learning process, every sequence corresponding to each patient is provided to the model a specified number of times (epochs). Before each sequence is provided to the model, the values of the eligibility traces and the FIFO queues are initialised to zero.

Given the restriction in processing power, all DyBMs were trained with a total of 100 iterations, and the sequence of each patient was supplied to the model for 200 epochs in each iteration.

Relative to the tDBN, for each considered dataset, two models are trained. One model is trained with data from patients of the negative class and another is trained with data from patients of the positive class. Classification is performed by calculating the probability that a given pattern was generated by each model.

B. Evaluation procedure

With respect to the DyBM, evaluation is performed by giving a pattern to the model and allowing the model to return a prediction. This prediction corresponds to the probability of the given patient eventually stopping therapy due to inefficiency.

Once again, before presenting the model with the first pattern of a given patient, the eligibility traces and FIFO queues are initialised with zeros. Then, the model's prediction is obtained by first allowing the model to update its state, that is, the eligibility traces and FIFO queues, and then retrieving the model's prediction given the model parameters and state.

[2]https://github.com/scikit-learn-contrib/imbalanced-learn

TABLE I
FEATURES PRESENT IN THE PREPROCESSED DATASET WITH RESPECTIVE DESCRIPTION AND DISCRETIZATION PERFORMED.

Feature	Description	Discretization
Static Variables		
Adalimumab	Whether the patient is under treatment by Adalimumab	**0**: No; **1**: Yes
Certolizumab	Whether the patient is under treatment by Certolizumab	**0**: No; **1**: Yes
Etanercept	Whether the patient is under treatment by Etanercept	**0**: No; **1**: Yes
Golimumab	Whether the patient is under treatment by Golimumab	**0**: No; **1**: Yes
Infliximab	Whether the patient is under treatment by Infliximab	**0**: No; **1**: Yes
Secucinumab	Whether the patient is under treatment by Secucinumab	**0**: No; **1**: Yes
Dynamic Variables		
ASDAS	ASDAS calculated with CRP	**0**: $[0; 1.3[$; **1**: $[1.3; 2.1[$; **2**: $[2.1; 3.5]$; **3**: $]3.5; +\infty[$
ASDAS ESR	ASDAS calculated with ESR	**0**: $[0; 1.3[$; **1**: $[1.3; 2.1[$; **2**: $[2.1; 3.5]$; **3**: $]3.5; +\infty[$
BASDAI	Index used to assess disease activity	**0**: $[0; 4[$; **1**: $[4; 7[$; **2**: $[7; 10]$
BASDAI Q1-6	Questions 1 through 6 of the BASDAI index	**0**: $[0; 4[$; **1**: $[4; 7[$; **2**: $[7; 10]$
BASFI	Index used to assess the degree of functional limitation	**0**: $[0; 3]$; **1**: $]3; 7]$; **2**: $]7; 10]$
BASFI Q1-10	Questions 1 through 10 of the BASFI index	**0**: $[0; 3]$; **1**: $]3; 7]$; **2**: $]7; 10]$
Cortic	Whether a patient is taking corticosteroid	**0**: No; **1**: Yes
CRP	C-reactive protein (mg/L)	**0**: $[0; 5]$; **1**: $]5; 10]$; **2**: $]10; 100]$; **3**: $]100; +\infty[$
ESR	Erythrocyte Sedimentation Rate (mm/hr)	**0**: $[0; 6]$; **1**: $]6; 17]$; **2**: $]17; +\infty[$
Painful joints	Number of painful joints	**0**: $[0; 5[$; **1**: $[5; 75]$
Swollen joints	Number of swollen joints	**0**: 0; **1**: $[1; 75]$
VAS Medic	VAS measured from the medic's perspective	**0**: $[0; 1[$; **1**: $[1; 3[$; **2**: $[3; 10]$
VAS Nocturnal	VAS measured during the nighttime	**0**: $[0; 2[$; **1**: $[2; 5[$; **2**: $[5; 10]$
VAS Patient	VAS measured from the patient's perspective	**0**: $[0; 2[$; **1**: $[2; 5[$; **2**: $[5; 10]$
VAS Total Column	VAS measured with respect to the column	**0**: $[0; 2[$; **1**: $[2; 5[$; **2**: $[5; 10]$
Outcome		
Inefficiency	Whether a patient stopped therapy due to inefficiency	**0**: No; **1**: Yes

TABLE II
NUMBER OF PATIENTS AND NUMBER OF PATIENTS THAT STOPPED THERAPY DUE TO INEFFICIENCY IN EACH DATASET.

Dataset	Number of patients	Number of patients that stopped therapy
$\mathtt{D}_{\mathtt{Adalimumab}}$	337	59 (17.5%)
$\mathtt{D}_{\mathtt{Certolizumab}}$	19	1 (5.3%)
$\mathtt{D}_{\mathtt{Etanercept}}$	329	62 (18.8%)
$\mathtt{D}_{\mathtt{Golimumab}}$	183	33 (18.0%)
$\mathtt{D}_{\mathtt{Infliximab}}$	261	82 (31.4%)
$\mathtt{D}_{\mathtt{Secucinumab}}$	12	2 (16.7%)

As furthers patterns are supplied to the model, predictions are obtained in the same way.

Relative to the tDBN, for each model, the probability that a given pattern was generated by the model is calculated. Then, classification is made according to the model that is more likely to have generated the pattern (*i.e.*, higher probability).

V. RESULTS

The tDBN method was applied to the datasets, where a stationary network was considered, as well as a first-order Markov process, $m = 1$. The minimum description length (MDL) score was used as the scoring function and, for DATASET 1 p was set to 3, while for DATASET 2 this value was set to 2 because of the higher number of features. The the implementation of this model is available on GitHub[3].

Relative to the DyBM, as it was mentioned a binary model was used. Specifically, the `BinaryDyBM` model available on GitHub[4] was used. For this model, the conduction delay parameter was varied between 1, 2 and 3. For the rest of the parameters their default values were used.

A. Performance

The performance is relative to the ability of the models in predicting the outcome (*i.e.*, whether the patient switched therapy).

As an example, relative to the performance of the models on dataset $\mathtt{D}^1_{\mathtt{Etanercept}}$, TABLE III shows the results obtained in all timesteps.

In a general sense, the results present in TABLE III, show that the DyBM when trained with the imbalanced dataset obtained high accuracy scores but low F_1 scores implying that the model has overfitted. Further, using SMOTE as an oversampling technique to balance the datasets revealed that this method is not suitable for this data, or in a more extensive way, appropriate for time series, given that models trained with such data did not improve in performance compared to the models trained with the imbalanced dataset. Moreover, using undersampling to balance the dataset helped improve the performance of the DyBM.

Regarding the remaining datasets, similar conclusion can be drawn.

[3]https://github.com/SSamDav/learnDBN

[4]https://github.com/ibm-research-tokyo/dybm

TABLE III
ACCURACY, PRECISION, RECALL AND F_1 SCORES IN PERCENTAGE OBTAINED IN EACH TIMESTEP FOR MODELS TRAINED WITH DATASET $\mathtt{D}^1_{\mathtt{Etanercept}}$. MODELS WERE TRAINED WITH AN UNCHANGED DATASET, AND BALANCED DATASETS, BY UNDERSAMPLING AND ANOTHER BY APPLYING SMOTE. *DIVISION BY ZERO.

	Timestep																	
	1		2		3		4		5		6		7		8		9	
	DyBM	tDBN	DyBM	tDBN	DyBM	tDBN	DyBM	tDBN	DyBM	tDBN	DyBM	tDBN	DyBM	tDBN	DyBM	tDBN	DyBM	tDBN
Unchanged dataset																		
Accuracy	81.82	31.82	83.33	50.00	83.33	54.55	83.33	59.09	84.85	63.64	85.25	65.57	84.48	67.24	85.71	66.07	85.19	66.67
Precision	*0.00	16.33	100.0	21.62	60.00	20.00	60.00	22.22	100.0	25.00	100.0	21.05	100.0	23.53	100.0	26.32	100.0	27.78
Recall	0.00	66.67	8.33	66.67	25.00	50.00	25.00	50.00	16.67	50.00	10.00	40.00	10.00	40.00	20.00	50.00	20.00	50.00
F_1	0.00	26.23	15.38	32.65	35.29	28.57	35.29	30.77	28.57	33.33	18.18	27.59	18.18	29.63	33.33	34.48	33.33	35.71
Undersampled dataset																		
Accuracy	53.03	—	56.06	—	60.61	—	60.61	—	71.21	—	70.49	—	65.52	—	71.43	—	66.67	—
Precision	22.86	—	25.71	—	25.00	—	26.67	—	34.78	—	27.78	—	25.00	—	33.33	—	30.00	—
Recall	66.67	—	75.00	—	58.33	—	66.67	—	66.67	—	50.00	—	50.00	—	60.00	—	60.00	—
F_1	34.04	—	38.30	—	35.00	—	38.10	—	45.71	—	35.71	—	33.33	—	42.86	—	40.00	—
SMOTE dataset																		
Accuracy	66.67	63.64	63.64	66.67	69.70	68.18	66.67	69.70	63.64	66.67	63.93	68.85	74.14	72.41	66.07	75.00	68.52	75.93
Precision	14.29	7.14	12.50	18.75	25.00	15.38	14.29	21.43	7.14	18.75	7.14	20.00	27.27	28.57	20.00	35.71	23.08	38.46
Recall	16.67	8.33	16.67	25.00	33.33	16.67	16.67	25.00	8.33	25.00	10.00	30.00	30.00	40.00	30.00	50.00	30.00	50.00
F_1	15.38	7.69	14.29	21.43	28.57	16.00	15.38	23.08	7.69	21.43	8.33	24.00	28.57	33.33	24.00	41.67	26.09	43.48

B. DyBM weights

In this section the analysis focuses on the weights of the DyBM. By taking an approach where only the weights that affect the first pattern are considered, it is possible to investigate the effect of each feature on the outcome when considering only the first pattern.

Recollect that lower energy values translate into higher probabilities. In the case at hand, a lower energy value signifies a higher probability of stopping therapy. Therefore, following (2), a positive weight causes a reduction of energy, or equivalently a higher probability of inefficiency. Inversely, a negative weight induces an increase in energy, resulting in a lower probability of inefficiency.

With the help of figures like Fig. 2 with respect to all datasets from DATASET 1 and DATASET 2, it was possible to draw some conclusions.

Given that feature CORTIC has a negative weight on models trained with dataset $\mathtt{D}_{\mathtt{Etanercept}}$ but shows a positive weight in almost all other models, etarnercept can be the better choice for patients that are taking corticosteroids. Similarly, patients with low VAS values can benefit from the use of adalimumab. It is worth noting that a CRP value of 3 is, is almost all models, correlated with inefficiency. Indicating that patients with such high CRP values tend to need to switch to a different biological drug. In addition, patients with a BASFI value of 0 benefit from being treated with etanercept as given by the negative values this feature presents on models trained with dataset $\mathtt{D}^1_{\mathtt{Etanercept}}$, meanwhile, patients with higher values do not. In such cases, patients should instead be treated with adalimumab. On a last note, patients with a low ASDAS value of 0 or 1 can benefit more from adalimumab, as indicated by the low weights obtained with such models.

Relative to the particular BASDAI and BASFI questions, some were found to not be as relevant. However, these questions were not the same in all therapies. Nevertheless, note that BASDAI question 3 and BASFI questions 2, 6 and 7 are not as relevant in two of the four datasets.

C. DBN relations

In general, the trained DBN models found correlations between variables that make sense. Fig. 3 shows two diagrams with such interesting correlations.

To start, both network diagrams in Fig. 3 found a correlation between features CORTIC and SWOLLEN JOINTS. Such correlation makes sense since usually corticosteroids are administrated to patient with the objective of reducing inflammation. Moreover, both networks also show a correlation between VAS PATIENT and BASFI. Logically, the pain intensity (measured by VAS) affects the patient's ability to perform certain tasks (measured by BASFI). On the same note, the network present in Fig. 3b, also found a correlation between BASFI and PAINFUL JOINTS which can be explained by the same reasoning. Relative to ESR, two interesting relations are present. On the network present in Fig. 3b there is a relation between ASDAS ESR and ESR. This relation makes sense since the ESR value is used to calculate ASDAS ESR. Lastly, Fig. 3a shows a correlation between SWOLLEN JOINTS and ESR. This correlation is understandable since an increase in ESR can be directly caused by inflammation.

D. DBN relative influence

In this section, the relative influence of each variable in the transition network for each tDBN trained is investigated. To do so, on a transition network the number of children of each variable is counted. Intuitively, a variable x_1 with a higher number of children than a variable x_2 is considered to be more relevant. Then, these counts are normalized, representing the relative influence of each variable in the transition network. Allowing, with the help of a stacked bar plot, like the one present in Fig. 4, to easily identify features that are less relevant.

Across all models, feature VAS NOCTURNAL obtained a low relative influence. Moreover, feature CORTIC has a higher relative influence in models trained with data from the positive class (class 1). In fact, in models trained with data from the negative class, this feature shows to not be relevant at all.

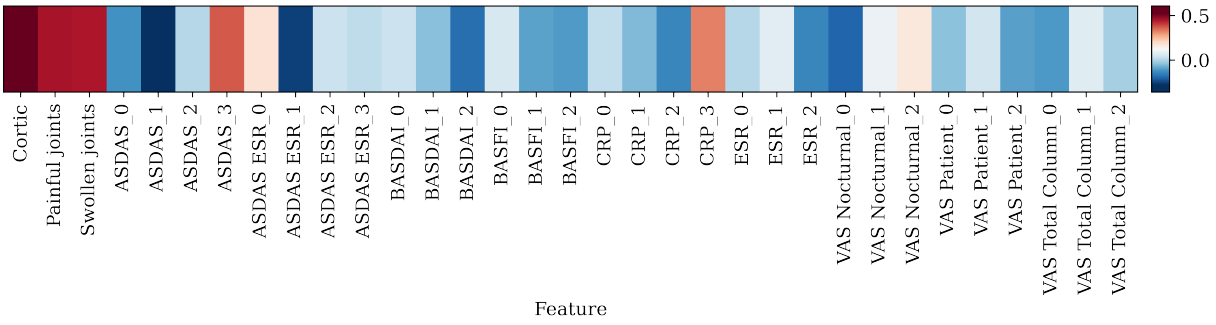

Fig. 2. Heatmap of weights of models trained with $\mathtt{D}^1_{\mathtt{Adalimumab}}$. Weights that contribute to the energy of the first pattern are summed together.

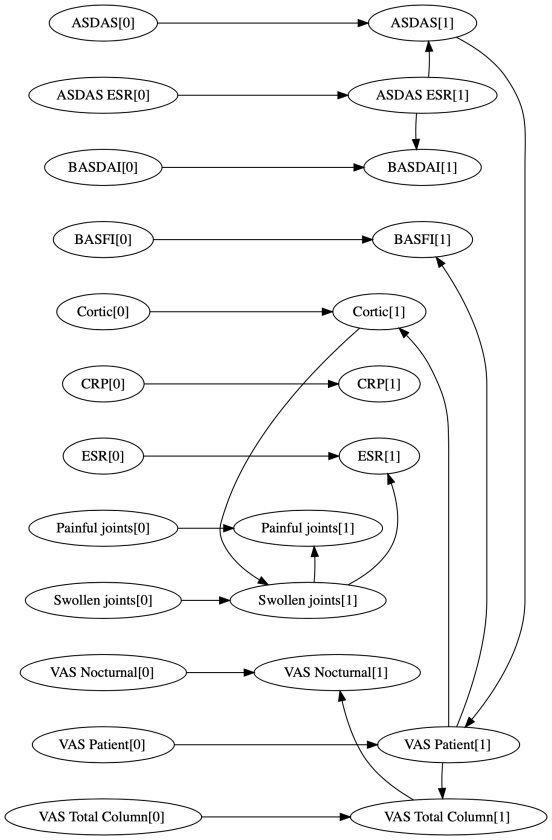

(a) Diagram of the DBN trained with patients from the positive class of dataset $\mathtt{D}^1_{\mathtt{Etanercept}}$.

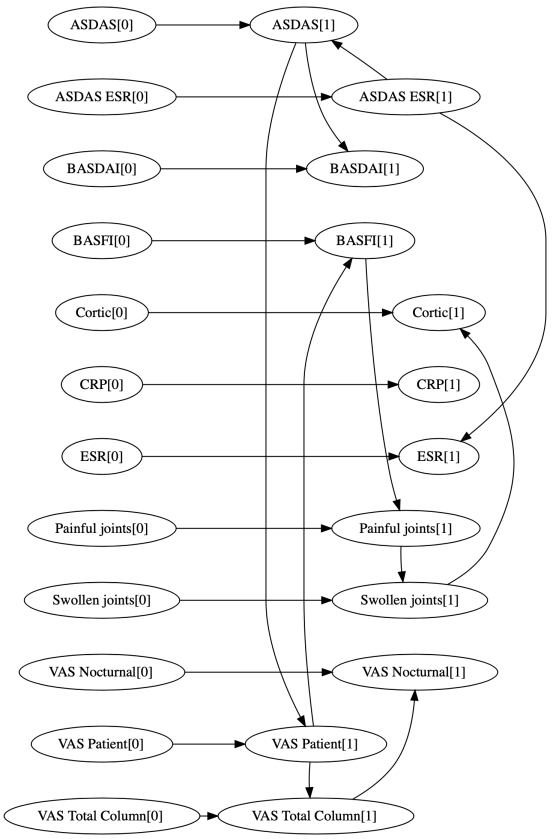

(b) Diagram of the DBN trained with patients from the negative class of dataset $\mathtt{D}^1_{\mathtt{Infliximab}}$.

Fig. 3. Diagrams of a couple of DBN transition networks.

Relative to individual BASDAI and BASFI questions, some questions also show a low relative influence. Where, analogously to the DyBM, these questions vary between therapy. However, unlike with the DyBM, BASDAI question 6 has a low relative influence across the four therapies. Additionally, BASFI question 8 has a low relative influence in all therapies with the exception of golimumab, where the model related with the positive class has a slightly higher relative influence. On a last note, BASDAI question 1 together with BASFI questions 1 and 6 have revealed a low relative influence in two of the four therapies.

VI. Conclusion

This work focused on using model-based machine learning methods to analyse time series data from patients with AS, a chronic autoimmune inflammatory disease which causes a significant decrease in the subject's quality of life. To achieve this, two dynamic models were review and applied. Furthermore, the data extracted from Reuma.pt had to be preprocessed. As a result, a ready-to-use dataset by other researchers was created.

In general, the DyBM trained with imbalanced datasets obtained high accuracy scores, however such models obtained low F_1 scores, indicating that overfitting has occurred. In

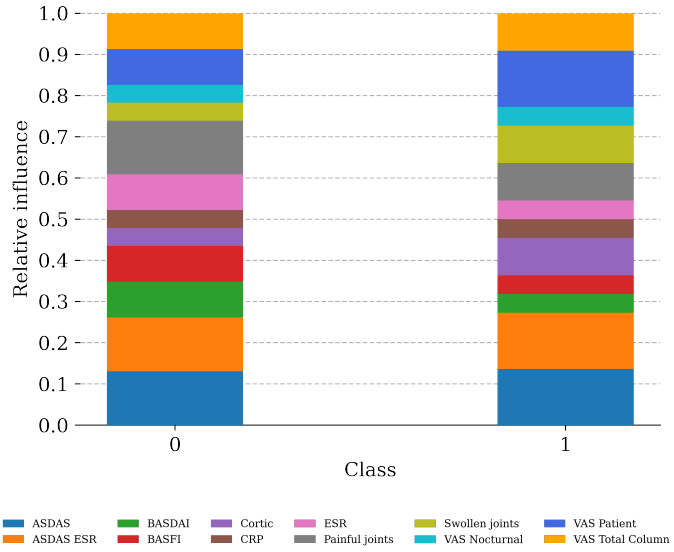

Fig. 4. Relative influence of each variable in the transition network for each tDBN trained with an unchanged (imbalanced) $D^1_{\texttt{Adalimumab}}$. Where, class 0 corresponds to the tDBN model trained with data from class 0 (negative sample). Likewise, class 1 corresponds to the model trained with data from the positive class.

addition, with the exception of some cases, the use of SMOTE did not improve the performance for both models relative to the imbalanced datasets. Suggesting that this oversampling method is not appropriate for this data or, in a more general sense, not adequate for time series. Lastly, the DyBM trained with the balanced, by undersampling, datasets obtained better results than both, imbalanced and SMOTE datasets, across most timesteps.

The learned DyBMs have showed that if a patient is taking corticosteroids then etanercept can be a better choice of biological drug. On a similar note, patients with low VAS values can benefit from the use of adalimumab. In addition, the model has found that patients with high CRP (greater than $100mg/L$) usually don't respond as well to treatment and will, most likely, need to switch to a different biological drug. Moreover, patients with BASFI values lower than or equal to 3 are very good candidates for treatment with etanercept, while patients with higher values are not, and should instead be treated with adalimumab.

The model also allowed to identify some individual BASDAI and BASFI questions that show to not be as relevant, however, such questions vary between the different biological drugs. Nevertheless, it was noted that the third BASDAI question together with BASFI questions 2, 6 and 7 were found to not be relevant in two of the four therapies.

In the context of the DBN, the model found interesting and justifiable correlations between features. To start, the model found a correlation between the use of corticosteroids and the presence of swollen joints. On top of that, a correlation between the VAS measured from the patient's perspective and BASFI was also discovered, as well as a correlation between BASFI and the presence of painful joints. Lastly, two correlations involving the ESR were found. One between this value and ASDAS-ESR and another between the existence of swollen joints and ESR.

In regard to the relative influence of the features, VAS NOCTURNAL showed to not be a relevant feature. It was also noted that feature CORTIC shows an increased importance in models trained with the positive class. In fact, on models trained with the negative class this feature shows to not be relevant. Still on this topic, the model allowed to identify particular BASDAI and BASFI questions that exhibited a low relative influence. Like with the DyBM the questions differ between therapy. However the sixth BASDAI question was found to not be relevant across all therapies. In addition, the eight BASFI question was discovered to not be relevant in all therapies with the exception of one. Lastly, BASDAI question 1 and BASFI questions 1 and 6 were found to not be relevant in two of the four therapies.

A. Future work

As was stated, the RDPR database available at REUMA.PT is responsible for collecting data relative to patients with multiple rheumatic diseases. Therefore, a similar research process can be conducted with data from other rheumatic diseases. In future work one could also expand the research to more machine learning methods.

In this work, the analysis focused on one event, whether a patient stopped therapy due to inefficiency. Patients that stopped therapy end up switching to a different one. On that note, other events concerning subsequent therapies could be explored.

The amount of iterations and epochs used to train the DyBM could be increased. Increasing these parameters might increase the performance of this model, however, with the cost of an increased time to train. Having said that, the time this model takes to train could be reduced by taking advantage of a graphics processing unit (GPU) that supports the CUDA API.

With respect to the DBN a non-stationary network could instead be learned. However to fully take advantage of such network, it may be necessary to obtain more data. More precisely, the number of patients may need to be larger, given that a non-stationary network learns a transition network for each transition in the data.

ACKNOWLEDGMENT

The author would like to thank the Portuguese Foundation for Science and Technology (Fundação para a Ciência e a Tecnologia-FCT) for providing a Research Grant under the PREDICT project (PTDC/CCI-CIF/29877/2017).

REFERENCES

[1] F. Mahmood and P. Helliwell, "Ankylosing spondylitis: A review," *European Medical Journal*, vol. 2.4, p. 6, 2017.

[2] A. R. Cravo, V. Tavares, and J. C. D. Silva, "Terapêutica anti-TNF Alfa na espondilite anquilosante," *Acta Médica Portuguesa*, vol. 19, no. 2, pp. 141–150, 2006.

[3] H. Canhão, A. Faustino, F. Martins, and J. Fonseca, "Reuma. PT - the rheumatic diseases portuguese register," *Acta reumatológica portuguesa*, vol. 36, pp. 45–56, 2011.

[4] J. L. Monteiro, S. Vinga, and A. M. Carvalho, "Polynomial-Time Algorithm for Learning Optimal Tree-Augmented Dynamic Bayesian Networks," in *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence*. Amsterdam, Netherlands: AUAI Press, 2015, pp. 622–631.

[5] T. Osogami and M. Otsuka, "Seven neurons memorizing sequences of alphabetical images via spike-timing dependent plasticity," *Scientific Reports*, vol. 5, no. 1, p. 14149, Sep. 2015. [Online]. Available: https://doi.org/10.1038/srep14149

[6] ——, "Learning dynamic Boltzmann machines with spike-timing dependent plasticity," *arXiv:1509.08634*, Sep. 2015. [Online]. Available: http://arxiv.org/abs/1509.08634

[7] N. Caporale and Y. Dan, "Spike Timing–Dependent Plasticity: A Hebbian Learning Rule," *Annual Review of Neuroscience*, vol. 31, no. 1, pp. 25–46, Jul. 2008. [Online]. Available: http://www.annualreviews.org/doi/10.1146/annurev.neuro.31.060407.125639

[8] T. Osogami, "Learning binary or real-valued time-series via spike-timing dependent plasticity," *arXiv:1612.04897 [cs, stat]*, Dec. 2016, arXiv: 1612.04897. [Online]. Available: http://arxiv.org/abs/1612.04897

[9] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, "SMOTE: Synthetic Minority Over-sampling Technique," *Journal of Artificial Intelligence Research*, vol. 16, pp. 321–357, Jun. 2002, arXiv: 1106.1813. [Online]. Available: http://arxiv.org/abs/1106.1813

[10] G. Lemaître, F. Nogueira, and C. K. Aridas, "Imbalanced-learn: A Python Toolbox to Tackle the Curse of Imbalanced Datasets in Machine Learning," *Journal of Machine Learning Research*, vol. 18, no. 17, pp. 1–5, 2017. [Online]. Available: http://jmlr.org/papers/v18/16-365.html