PASS: PHASED ATTENTIVE STATE SPACE MODELING OF DISEASE PROGRESSION TRAJECTORIES

Anonymous authors

Paper under double-blind review

Abstract

Disease progression models are instrumental in *predicting* individual-level health trajectories and understanding disease dynamics. Existing models are capable of providing either accurate predictions of patients' prognoses or clinically interpretable representations of disease pathophysiology, but not both. In this paper, we develop the *phased attentive state space* (PASS) model of disease progression, a deep probabilistic model that captures complex representations for disease progression while maintaining clinical interpretability. Unlike Markovian state space models which assume memoryless dynamics, PASS uses an attention mechanism to induce "memoryful" state transitions, whereby repeatedly updated attention weights are used to focus on past state realizations that best predict future states. This gives rise to complex, non-stationary state dynamics that remain interpretable through the generated attention weights, which designate the relationships between the realized state variables for *individual* patients. PASS uses *phased* LSTM units (with time gates controlled by parametrized oscillations) to generate the attention weights in continuous time, which enables handling irregularly-sampled and potentially missing medical observations. Experiments on data from a realworld cohort of patients show that PASS successfully balances the tradeoff between accuracy and interpretability: it demonstrates superior predictive accuracy and learns insightful individual-level representations of disease progression.

1 INTRODUCTION

Chronic diseases – such as cardiovascular disease, cancer and diabetes – progress slowly throughout a patient's lifetime, causing increasing burden to the patients, their carers, and the healthcare delivery system Sevick et al. (2007). Modern electronic health records (EHR) keep track of individual patients' disease progression trajectories through follow-up data sequences of the form $(X_1, ..., X_t)$, where X_t is a set of clinical observations collected for the patient at time t. The advent of EHRs¹ provides an opportunity for building models of disease progression that can fulfill two central goals of healthcare delivery systems:

- <u>Goal A</u>: *Predicting* individual-level disease trajectories.
- <u>Goal B</u>: *Understanding* disease progression mechanisms.

Goal A entails the *supervised* problem of predicting future clinical observations $(X_{t+1}, X_{t+2},...)$ on the basis of past observations $(X_1, ..., X_t)$. **Goal B** entails the *unsupervised* problem of discovering clinically-interpretable latent structures that explain the mechanisms underlying disease progression. Both goals **A** and **B** are entangled. This is because accurate predictions need to be transparent and interpretable in order to ensure their actionability, whereas interpretable representations explaining disease progression can only be trustworthy if they possess high predictive power.

Unfortunately, as a consequence of the inherent tension between model *accuracy* and *interpretability* Lipton (2016), most existing models of disease progression fulfill either **Goal A** or **Goal b**, but not both. State-of-the-art prediction performance is achieved by models based on *recurrent neural networks* (RNN) Lim & van der Schaar (2018); Lipton et al. (2016); Choi et al. (2016a). RNN-based models are often used for sequence prediction (or sequence labeling), where they are trained

¹https://www.healthit.gov/sites/default/files/briefs/



(d) Unrolled graphical depiction for an attentive state space

Figure 1: Depictions for different models of sequential data: (a) Graphical model for an RNN. \Diamond denotes a deterministic intermediate representation, (b) Graphical model for an HMM. \bigcirc denotes probabilistic states, (c) Graphical model for the proposed attentive state space model. (c) Unrolled instance-wise graphical depiction for an attentive state space. Thickness of the arrows reflect the attention weights.

to estimate the predictive distribution $P(X_t | X_{t-1}, ..., X_1)$ by propagating a sequence of hidden states $(Z_1, ..., Z_t)$ through intermediate deterministic mappings (Figure 1a). Unfortunately, an RNN is of a "black-box" nature since its hidden states $(Z_1, ..., Z_t)$ do not explicitly map to clinically meaningful states of disease progression. On the contrary, *state space* approaches based on Hidden Markov Models (HMM) provide a natural interpretation of a disease trajectory as a sequence of transitions between latent "progression stages" $(Z_1, ..., Z_t)$ (Figure 1b), each of which corresponds to a clinically distinguishable disease state Alaa & van der Schaar (2018); Liu et al. (2015); Wang et al. (2014). However, the interpretability of HMMs comes at a price. That is, while RNNs can in principle approximate *any* dynamical system, an HMM is limited to memoryless Markovian state dynamics, which greatly undermine its predictive performance.

Our Contribution In this paper, we develop a deep probabilistic model of disease progression that capitalizes on both the predictive power of RNNs and the interpretable nature of state space models to fulfill **Goals A** and **B**. Our model maintains the probabilistic structure of a state space representation, which decouples *emission* and *transition* distributions, but uses an RNN to model a flexible non-Markovian state transition dynamic $P(\mathbf{Z}_t | \mathbf{Z}_{t-1}, ..., \mathbf{Z}_1)$ that allows future states to depend on all past states. To model state transitions, we use an attention mechanism whereby the RNN generates a (repeatedly updated) set of attention weights $(\alpha^1, ..., \alpha^t)$ that designate the (relative) influence that past state realizations ($\mathbf{Z}_1, ..., \mathbf{Z}_t$) have on the transition probabilities to the future state \mathbf{Z}_{t+1} . Our model for state transitions at time step t can be summarized as follows:

$$(\boldsymbol{\alpha}_{1}^{t},...,\boldsymbol{\alpha}_{t}^{t}) = \mathbf{RNN}(\boldsymbol{X}_{1},...,\boldsymbol{X}_{t}),$$

 $P(\boldsymbol{Z}_{t+1} = z) = \sum_{i < t} \boldsymbol{\alpha}_{i}^{t} \times (A \text{ baseline transition probability from } \boldsymbol{Z}_{i} = z' \text{ to } z).$

The model described above, which we call an *attentive state space* model, uses a set of RNNgenerated attention weights to induce a time-varying Markov blanket for the state variable Z_t . This Markov blanket changes for every new state transition, putting more or less attention on previous state realizations depending on the patient's clinical history. If we restrict the attention weights to be binary ($\alpha_i^t \in \{0, 1\}, \forall i \leq t$), the attentive state space becomes a *variable-order* Markov model that decides the extent of memory involved in state transitions in every time step depending on the patient's current context. This allows for realistic non-stationary and time-inhomogeneous dynamics that are implicitly captured by an RNN but could not be possibly modeled with an HMM. The attentive state representation is clinically interpretable because the complex state dynamics that it captures are fully explicable through the attention weights, which indicates the extent to which past clinical events contribute to future state realizations. Figure 1c provides a formal graphical model for our attentive state space representation. Figure 1d shows an unrolled graphical depiction of the model for a particular exemplary patient, highlighting the time-varying nature of the attention weights and their straightforward interpretational benefit.

Because EHR data comprises irregularly-sampled and asynchronous observations gathered only at the times when the patient visits a hospital, we use the *phased* LSTM units introduced in Neil et al. (2016), with time gates controlled by parametrized oscillations, in order to generate the attention weights at arbitrary time instances. Thus, we call our model a *phased attentive state space* (PASS) model. A detailed description of the construction of the PASS model is provided in Section 2. A detailed comparison between PASS and related models can be found in Section 4.

Indeed, state inference and parameter estimation of PASS is nontrivial since non-Markovianity hinders the application of conventional backward message passing algorithms Alaa & van der Schaar (2018); Dai et al. (2016). In Section 3, we show that PASS can be re-parameterized as a nonstationary dynamic Bayesian network, for which conventional forward message passing algorithms can be implemented with a complexity resembling that of the forward filtering algorithm used for HMMs. We conduct parameter estimation via a variant of the Expectation-Maximization algorithm. In Section 5, we conduct experiments on data from a real-world longitudinal cohort of more than 10,000 Cystic Fibrosis (CF) patients. Our experiments show that PASS successfully balances the tradeoff between accuracy and interpretability: it demonstrates superior predictive accuracy and learns insightful individual-level representations of disease progression. In particular, we show that PASS learns meaningful population-level CF progression stage, and that the attention weights can inform treatment decisions on the level of individual patients.

2 A PHASED ATTENTIVE STATE-SPACE MODEL OF DISEASE PROGRESSION

We model the progression of a target chronic disease using longitudinal EHR data for patients who have developed, or are at risk of developing such disease. We start by describing the model variables in Section 2.1, and then we develop the attentive state dynamics in Sections 2.2 and 2.3.

2.1 MODEL VARIABLES AND NOTATION

Structure of the EHR data A patient's EHR record, denoted as \mathcal{D} , is a collection of timestamped follow-up data gathered during repeated, irregularly-spaced hospital visits, in addition to static features (e.g., genetic variables). We represent a given patient's EHR record as follows:

$$\mathcal{D} = \{ \mathbf{Y} \} \cup \{ (\mathbf{X}_m, \mathbf{t}_m) \}_{m=1}^M, \tag{1}$$

Static features Visit times

where Y is the static features' vector, X_m is the follow-up data collected in the m^{th} hospital visit, t_m is the time of the m^{th} visit, and M is the total number of hospital visits. (The time-horizon t is taken to be the patient's chronological age.) The follow-up data X_m comprises information on biomarkers and clinical events, such as treatments and diagnoses of comorbidities. (Refer to the Appendix for a more elaborate discussion on the type of follow-up data collected in EHRs.) An EHR dataset $\{\mathcal{D}^{(i)}\}_{i=1}^{N}$ is an assembly of records for N independent patients.

Disease progression stages We assume that the target disease evolves through *D* different *progression stages*. Each stage corresponds to a distinct level of disease severity that manifests through the follow-up data. We model the evolution of progression stages via a (continuous-time) stochastic process Z(t) of the following form:

$$\mathbf{Z}(t) = \sum_{n \in \mathbb{N}_{+}} \tilde{Z}_{n} \cdot \mathbf{1}_{\{T_{n} < t \le T_{n+1}\}}, \ \tilde{Z}_{n} \in \{1, \dots, D\},$$
(2)

where $\{T_n\}_n$ is the sequence of *onsets* for the realized progression stages, and \tilde{Z}_n is the progression stage occupying the interval $(T_n, T_{n+1}]$. We assume that Z(0) = 1 (i.e., $\tilde{Z}_1 = 1$) almost surely, with stage 1 being the *asymptomatic* stage designating "healthy" patients. The sequence $\{Z_m\}_m$ is the *embedded* discrete-time process induced by Z(t) at the hospital visit times $\{t_m\}_m$, i.e. $Z_m = Z(t_m)$.

2.2 ATTENTIVE STATE SPACE REPRESENTATION

We adopt a state space representation for the disease progression process, with the state space being the set of all stages of progression $\{1, \ldots, D\}$. The states' sequence $\{Z_m\}_m$ is *hidden* whereas the EHR data \mathcal{D} is *observed*. We consider a graphical model that defines probabilistic dependencies between $\{Z_m\}_m$ and \mathcal{D} through the following factorization of *emission* and *transition* distributions:

$$P(\{\boldsymbol{Z}_m\}_m, \{\boldsymbol{X}_m\}_m | \boldsymbol{Y}, \{t_m\}_m) = \prod_{m'=1}^m \underbrace{P(\boldsymbol{X}_{m'} | \boldsymbol{Z}_{m'})}_{\text{Emission}} \cdot \underbrace{P(\boldsymbol{Z}_{m'} | \mathcal{F}_{t_{m'-1}})}_{\text{Transition}}, \quad (3)$$

 $\forall m \in \{1, \dots, M\}$, where the filtration $\mathcal{F}_{t_{m'}} = \{Y, (Z_1, X_1, t_1), \dots, (Z_{m'}, X_{m'}, t_{m'})\}$ conveys all the information available in the model up to time $t_{m'}$. We model the emission distribution in (3) as a Gaussian distribution with state-specific parameters as follows:

$$P(\boldsymbol{X}_m \mid \boldsymbol{Z}_m = z) = \mathcal{N}(\boldsymbol{\mu}_z, \boldsymbol{\Sigma}_z), \ z \in \{1, \dots, D\}.$$
(4)

Binary variables are modeled with a Bernoulli state-specific distributions. The transition probability factor in (3) assumes that the realized state at time t_m depends on the entire process history $\mathcal{F}_{t_{m-1}}$. To model $P(\mathbf{Z}_{m'} | \mathcal{F}_{t_{m'-1}})$, we first define a $D \times D$ baseline Markov generator matrix Λ as follows:

$$\mathbf{\Lambda} = \begin{bmatrix} -\lambda_{12} & \lambda_{12} & 0 & \dots & 0\\ 0 & -\lambda_{23} & \lambda_{23} & \dots & 0\\ \vdots & \vdots & \ddots & \vdots & \vdots\\ 0 & \dots & 0 & -\lambda_{D^{-1,D}} & \lambda_{D^{-1,D}}\\ 0 & \dots & 0 & 0 & 0 \end{bmatrix}, \ \lambda_{ij} \ge 0, \ \forall i, j \in \{1, \dots, D\},$$
(5)

where λ_{ij} is a Markovian transition rate from state *i* to state *j*. Λ is the transition rate matrix of a continuous-time Markov chain model on the state space $\{1, \ldots, D\}$; its bidiagonal structure forces transitions to be permissible only between adjacent states (in an ascending order), with the last state (state *D*) being an absorbing state. This enforces the states in the set $\{1, \ldots, D\}$ to map properly to the disease progression stages (state 1 is the least severe stage and state *D* is the terminal stage of illness). We model the state transition probability $P(\mathbf{Z}_m = z | \mathcal{F}_{t_{m-1}})$ by creating a "memoryful" version of the Markov chain model in (5) through the following parametrization:

$$P(\mathbf{Z}_{m} = z \mid \mathcal{F}_{t_{m-1}}) = P(\mathbf{Z}_{m} = z \mid \underbrace{\{(\mathbf{Z}_{k} = z_{k}, \boldsymbol{\alpha}_{k}^{m}, \Delta_{k})\}_{k=1}^{m-1}}_{\mathsf{Sufficient statistics}}) = \sum_{k=1}^{m-1} \underbrace{\boldsymbol{\alpha}_{k}^{m}}_{\downarrow} (e^{\Delta_{k} \Lambda})_{z, z_{k}}, \quad (6)$$

where $\Delta_k = t_m - t_k$ is the time interval between the k^{th} and the m^{th} hospital visits, $\alpha_k^m \in [0, 1]$ is an *attention weight* assigned to the k^{th} visit, with $\sum_{k=1}^{m-1} \alpha_k^m = 1$, and $(e^{\Delta_k \Lambda})_{z,z_k}$ is entry (z, z_k) of the exponentiation of matrix Λ . The attention weights in (6) are generated via an attention function φ with parameter Θ as follows:

$$(\boldsymbol{\alpha}_1^m,\ldots,\boldsymbol{\alpha}_{m-1}^m) = \boldsymbol{\varphi}(\boldsymbol{Y},t_m,(\boldsymbol{X}_1,t_1),\ldots,(\boldsymbol{X}_{m-1},t_{m-1});\boldsymbol{\Theta}).$$
(7)

We call the representation in (6) an *attentive state space* representation. As shown in (6), the attentive representation starts with a baseline Markov chain, and creates memory in state transitions by weighting the baseline Markovian transition probabilities from all previous states, i.e. $\{e^{\Delta_k Q}\}_{k=1}^{m-1}$, using a set of attention weights $\{\alpha_k^m\}_{k=1}^{m-1}$. (The attention weights are generated on the basis of a patient's static features and follow-up data.) That is, instead of the memoryless Markovian dynamics in which a new state realization Z_{m+1} is fully determined by the current realization Z_m , the attentive state dynamics pay attention to all previous realizations in proportion to their attention weights. This gives rise to "memoryful", non-stationary, and time-inhomogeneous state transitions, whereby the time-varying Markov blanket (sufficient statistic) $\{(Z_k, \alpha_k^m, \Delta_k)\}_{k=1}^{m-1}$ of every new state realization Z_m determines which state realizations in the past matter most for the future.

Similar to an HMM model, the factorization in (3) decouples the transition and emission distributions by assuming that Z_m d-separates X_m from all other variables. This decoupling ensures that the clinical interpretability of the hidden states is maintained since each state is associated with a distinct emission distribution for the observed follow-up data. Moreover, regardless of the choice of the attention function φ , the attentive state transition matrix $\sum_{k=1}^{m-1} \alpha_k^m e^{\Delta_k \Lambda}$ will always be interpretable because the influence that a previous progression stage has on the future progression trajectory is encoded in its corresponding attention weight.



Figure 2: Illustration of phased attention mechanism. Left: Architecture of the phased attention network. The follow-up data (augmented with the static features) are fed in a reversed order into the phased LSTM, with reversed and shifted timestamps for the hospital visits. **Right:** Illustration for the operation of the phased LSTM. The time gate g_t of 4 neurons are depicted; each has different oscillatory parameters. The contents of the cell state c_t decay as we go backwards in time, implying smaller attention weights for older hospital visits.

2.3 THE PHASED ATTENTION MECHANISM

To implement the attentive state dynamics, the attention weights in (7) must be repeatedly updated (after each hospital visit) using variable-length sequences of data. Hence, we model the attention function φ as an RNN that maps a patient's history to attention weights as follows:

$$\boldsymbol{h}_{m-1}, \dots, \boldsymbol{h}_1 = \textbf{pLSTM}((\boldsymbol{X}_{m-1}, t_m - t_{m-1}), \dots, (\boldsymbol{X}_1, t_m - t_1); \boldsymbol{\Theta}),$$
$$\boldsymbol{e}_j = \boldsymbol{w}^T \boldsymbol{h}_j + \boldsymbol{b}, \, \forall j \in \{1, \dots, m-1\},$$
$$(\boldsymbol{\alpha}_1^m, \dots, \boldsymbol{\alpha}_{m-1}^m) = \text{Softmax}(\boldsymbol{e}_1, \dots, \boldsymbol{e}_{m-1}),$$
(8)

where **pLSTM** is a *phased* LSTM network Neil et al. (2016), w and b are the output layer parameters, h_j is the hidden layer, and $\tilde{X}_j = [Y, X_j, t_j]$ is the input at time step j. Unlike traditional RNNs, a phased LSTM takes a timestamped sequence as an input, and performs updates at arbitrary points of time. Phased LSTMs can also handle asynchronously-sampled sequences, which is particularly important for EHR data as not all of the components of X_m are necessarily measured in each hospital visits. Through phased LSTMs, we can update the attention weights with whatever follow-up data available at arbitrary time instances without the need for explicitly imputing missing observations. We call the attention mechanism in (8) the *phased attention* mechanism. The PASS model is an attentive state space model that uses the phased attention mechanism.

In the m^{th} time step, phased attention operates by feeding the phased LSTM with the sequence $\{\tilde{X}_j\}_{j=1}^{m-1}$ in reversed order, with timestamps reversed and shifted by t_m as shown in Figure 2 (left). This allows all attention weights allocated to all previous state realizations (or equivalently, hospital visits) to be dynamically updated at every time step while preserving the relative time spacing between hospital visits. The phased attention mechanism in (8) can be thought of as a continuous-time analogue of the reverse-time attention mechanism in Choi et al. (2016b).

The main difference between the phased LSTM model and the conventional LSTM model is the addition of a *time gate*, g_t , which controls the updates to the LSTM cell state c_t (and consequently the hidden layer h_t). The opening and closing of the time gate g_t for every neuron is controlled through an independent (continuous-time) rhythmic oscillation specified by 3 parameters (an oscillation period, a phase shift, and the ratio of the duration of the "open" phase to the full period) that can be learned from the data Neil et al. (2016). With every neuron having its own oscillatory parameters, the phased LSTM generates a continuum of possible updates that can be probed at arbitrary time instances as illustrated in Figure 3 (right). The updated equations of the phased LSTM can be found in Neil et al. (2016).

Algorithm 1 EM algorithm for learning the PASS model parameters

Input: EHR data \mathcal{D} , initial guess $\hat{\Gamma}^{(0)}$, state space $\{1, ..., D\}$ and number of iterations R. Output: Parameter estimate $\hat{\Gamma} = \{\hat{\Lambda}, (\hat{\mu}, \hat{\Sigma}), \hat{\Theta}\}$ while $i \leq R$ do $\mathbf{1} \in \{\hat{z}_m\}_m \leftarrow \text{Forward-Backward}(\mathcal{D} \mid \hat{\Gamma}^{(i)})$ $\mathbf{2} \in (\hat{\alpha}_1^m, ..., \hat{\alpha}_{m-1}^m) \leftarrow \varphi((\tilde{X}_{m-1}, t_m - t_{m-1}), ..., (\tilde{X}_1, t_m - t_1); \hat{\Theta}^{(i)}), m \in \{1, ..., M\}$ $\mathbf{3} \in \hat{p}_m \leftarrow \sum_{k=1}^{m-1} \hat{\alpha}_k^m (e^{\Delta_k} \hat{\Lambda}^{(i)})_{\hat{z}, \hat{z}_k} m \in \{1, ..., M\}$ $\mathbf{4} \in \hat{\Theta}^{(i+1)} \leftarrow \arg\min(-\sum_m \hat{z}_m \cdot \log(\hat{p}_m))$ $\mathbf{5} \in \hat{\Lambda}^{(i+1)} \leftarrow \arg\max\log(P(\mathcal{D} \mid \{\hat{z}_m\}_{m=1}^M, \Lambda)$ $\mathbf{6} \in (\hat{\mu}^{(i+1)}, \hat{\Sigma}^{(i+1)}) \leftarrow \arg\max\log(P(\{X_m\}_m \mid \{\hat{z}_m\}_{m=1}^M))$ $\mathbf{7} \cdot i \leftarrow i + 1$ end while

3 LEARNING AND INFERENCE

In this Section, we present the parameter learning and state inference algorithms for the PASS model. Throughout this Section, we continue with a single patient EHR record for the ease of notation.

Parameter learning Let Γ be the set of all PASS model parameters, i.e. $\Gamma = \{\Lambda, (\mu, \Sigma), \Theta\}$, where $\mu = \{\mu_z\}_{z=1}^{D}$ and $\Sigma = \{\Sigma_z\}_{z=1}^{D}$ are the emission parameters. The *complete data* log-likelihood of an EHR record \mathcal{D} and a state sequence realization $\{Z_m = z_m\}_m$ is given by:

$$\log(P(\boldsymbol{\mathcal{D}}, \{z_m\}_{m=1}^M \mid \boldsymbol{\Gamma})) = \sum_{m=1}^M \log\left(\sum_{k=1}^{m-1} \boldsymbol{\alpha}_k^m \cdot \exp(\boldsymbol{\Lambda} \left(t_m - t_k\right))_{z_k, z_m}\right).$$
(9)

Because the state sequence $\{Z_m\}_m$ is hidden, the complete data likelihood in (9) is inaccessible, and hence we resort to the Expectation-Maximization (EM) algorithm. The EM algorithm operates iteratively to update its guess of Γ , where the *i*th iteration implements 2 steps:

E-step:
$$Q(\Gamma \mid \hat{\Gamma}^{(i)}) = \mathbb{E}_{\{Z_m\}_m \mid \mathcal{D}, \hat{\Gamma}^{(i)}} [\log (P(\mathcal{D}, \{Z_m\}_m))],$$

M-step: $\hat{\Gamma}^{(i+1)} = \underset{\Gamma}{\operatorname{arg\,max}} Q(\Gamma \mid \hat{\Gamma}^{(i)}).$ (10)

Algorithm 1 lists the steps involved in implementing the EM algorithm in (10). In Step 1, we first infer the hidden states via a message passing algorithm (described later) using the current guess of Γ . Next, in Steps 2 and 3, we compute the attention weights associated with all hospital visit times, and then compute the transition probabilities using the formula in (6). In Step 4, the phased LSTM parameters are updated by optimizing the cross-entropy loss of the estimated transition probabilities and the inferred states. Maximum-likelihood is used to update the emission and transition parameters using the complete data likelihood obtained by plugging the inferred states into the expression in (9). To updated the Markov generator matrix Λ , we use the *Expm* method in Liu et al. (2015).



Figure 3: Rearranged super states for the attentive model in Figure 1c with truncated attention and Q = 1. The resulting graphical model (right) corresponds to a standard Hidden Markov model.

State inference One key advantage of the attentive state construction in (9) is that the attention weights explicitly quantify the importance of each past state to any given future state. Thus, efficient

inference can be conducted by limiting the Markov blanket for every state variable to "important" past states with attention weights exceeding a certain threshold. Since attention weights already decline as we go back in time, we approximate the Markov blanket for every state Z_m by only considering the Q most recent states $(Z_{m-1}, \ldots, Z_{m-Q})$. The resulting graphical model can be rearrange by lumping together every Q consecutive states into one "super state" (as shown in Figure 3), we retrieve a first-order Markovian dynamic Bayesian network (or equivalently, a higher-order Markov model Murphy & Russell (2002)), for which standard forward and backward message passing apply.

Finally, we note that in (6), we specified a finite-dimensional distribution over the embedded discrete process $\{(\mathbf{Z}_m, t_m)\}_m$ rather than directly modeling the continuous-time process $\mathbf{Z}(t)$. This is because non-Markovian modeling of $\mathbf{Z}(t)$ would require solving an intractable system of integral equations in order to compute the probability distribution of the discrete sequence $\{\mathbf{Z}_m\}_m$ Alaa & van der Schaar (2018). It is easy to show, using the *Kolmogorov extension Theorem*, that the model in (6) induces a consistent probability measure on $\mathbf{Z}(t)$ Kolmogorov (2018).

4 RELATED WORKS

Previous works related to PASS fall into three areas: state space models of disease progression, RNN-based predictive models for healthcare applications, and (general-purpose) deep probabilistic models. In what follows, we discuss previous works in these three areas, and then conclude the Section by demonstrating the generality of the PASS model.

State space models of disease progression Almost all existing models of disease progression are based on variants of the HMM model Wang et al. (2014); Liu et al. (2015); Alaa et al. (2017). Disease dynamics in such models are very easily interpretable as they can be perfectly summarized through a single matrix of probabilities that describes the transition rates among the different disease states. Markovian dynamics also greatly simplify inference because the model likelihood factorizes in a way that makes efficient forward and backward message passing possible Murphy & Russell (2002). However, memoryless Markov models assume that a patient's current state *d*-separates her future trajectory from her clinical history. This renders HMM-based models incapable of properly explaining the heterogeneity in the patients' progression trajectories, which often results from their varying clinical histories or the chronologies (timing and order) of their experienced clinical events Valderas et al. (2009). This limitation is particularly crucial in complex chronic diseases that are accompanied with multiple morbidities. As discussed earlier, PASS addresses this limitation by creating memoryful state transitions that depend on the patient's entire clinical history. (In the Appendix material, we provide a detailed discussion on how PASS can better explain patient heterogeneity compared to Markov models.)

RNN-based predictive modeling for healthcare Various RNN-based predictive models have been recently developed for healthcare settings; examples of such models include **Doctor AI** Choi et al. (2016a), **L2D** Lipton et al. (2016), and **Disease-Atlas** Lim & van der Schaar (2018). All those methods do not attempt to model a disease progression trajectory, but rather predict target clinical events on the basis of (discrete) sequential observations. Because of their black-box nature, none of these models can help understand the mechanisms underlying disease progression.

There have been various attempts to create interpretable RNN-based predictive models using attention. The models in Choi et al. (2016b) and Ma et al. (2017) use the *reverse-time attention* mechanism to learn visit-level and variable-level attention weights that explain the prediction of a target label through measures of variable importance. The phased attention mechanism proposed in Section 2.3 is a generalization of the reverse-time attention mechanism in Choi et al. (2016b) that can operate in continuous-time, and update the attention weights at irregularly-spaced and potentially incomplete observations. The main difference between the way attention is used in PASS and the way it is used in models like **RETAIN** Choi et al. (2016b) can be summarized as follows. PASS applies attention to the latent *state space*, whereas **RETAIN** applies attention to the observable *sample space*. Hence, the attention mechanism interprets the hidden disease dynamics, and hence it provides an explanation for the mechanisms underlying disease progression. On the contrary, **RETAIN** uses attention to measure feature importance, and hence it only explains predictions, but does not explain the disease progression mechanisms. **Deep probabilistic models** Most existing works on deep probabilistic models have focused on developed structured inference algorithms for deep Markov models and their variants Krishnan et al. (2017); Dai et al. (2016); Karl et al. (2016); Johnson et al. (2016). All such models use neural networks to model the transition and emission distributions, but are limited to Markovian dynamics. Other works develop stochastic versions of RNNs for the sake of generative modeling; examples include variational RNNs Chung et al. (2015), SRNN Fraccaro et al. (2016), and STORN Bayer & Osendorfer (2014). All such models augment stochastic layers to an RNN in order to enrich its output distribution. However, the transition and emission distributions in all these models cannot be decoupled, and hence their latent state representations would not lead to clinically meaningful identification of disease states. To the best of our knowledge, PASS is the first deep probabilistic model that provides both a clinically interpretable latent representation, and interpretable non-Markovian state dynamics.

Generality of the attentive state space representation For particular choices of the attention function in (7), the attentive state space representation in (6) reduces to various classical models of sequential data. For instance, if φ always sets α_{m-1}^m to 1 and all other weights to 0, then we retrieve an HMM. If the attention weights are binarized, then we retrieve a variable-order HMM Willems et al. (1995); Begleiter et al. (2004). Furthermore, if the attention weights are fixed, then we recover an auto-regressive model. This is a powerful feature of our model as it implies that by learning the attention function φ , we are effectively testing the assumptions of various commonly-used time series models in a data-driven fashion.

5 EXPERIMENTS

To validate the PASS model, we conducted a set of experiments using retrospective data for a longitudinal cohort of cystic fibrosis (CF) patients. CF is a life-shortening chronic condition that causes severe lung dysfunction, and is the most common genetic disease in Caucasian populations Szczesniak et al. (2017). All experimental details are listed hereunder.

Recall that, as stated in Section 1, the main purpose of the PASS model is to simultaneously fulfill **Goal A** (predicting individualized disease trajectories) and **Goal B** (understanding disease progression mechanisms). Thus, in Sections 5.1 and 5.2, we evaluate our model with respect to both goals.

Data description. The dataset involved in the experiments was extracted from the UK CF registry, a database maintained by the UK CF trust². Data was gathered from hospitals all over the UK, with 99% of patients consenting to their data being submitted, and hence the cohort is representative of the UK CF population. The dataset comprises longitudinal follow-ups for 10,263 patients over the period spanning between 2008 and 2015, with a total of 60,218 hospital visits. Each patient is associated with 90 variables, including the intake of 36 possible treatments, diagnoses for 31 possible comorbidities and 16 possible infections, FEV1 biomarkers, gender, and CF genetic mutations.

5.1 Goal A: PREDICTING INDIVIDUAL-LEVEL CF PROGRESSION TRAJECTORIES

Baselines. We compared the predictive accuracy of PASS to the following models:

- MLP: A multi-layer perceptron (MLP) classifier that is trained to sequentially predict the clinical events in 1 hospital visit given the observations in the prior visits. The MLP is trained on a static dataset that is created by unrolling the longitudinal follow-up data for all patients and treating every hospital visit as a separate data point.
- HMM: A standard continuous-time HMM model Wang et al. (2014); Liu et al. (2015) trained with the Baum-Welch EM algorithm. The number of HMM states was set via the Akaike information criterion (AIC). Similar to the PASS model, the observations (X_1, \ldots, X_t) are modeled as Gaussian emission variables with state-specific mean and variance parameters.
- **RNN:** A standard LSTM network with 2 hidden layers of size 200. The follow-up data (X_1, \ldots, X_t) was used as an input, and the output was defined as a set of (binary) labels

²https://www.cysticfibrosis.org.uk/the-work-we-do/uk-cf-registry/

	ABPA	Diabetes	I. Obstruction	K. Pneumoniae	E. Coli	Aspergillus
Model	AUC-ROC	AUC-ROC	AUC-ROC	AUC-ROC	AUC-ROC	AUC-ROC
PASS	$\textbf{0.687} \pm \textbf{0.022}$	$\textbf{0.771} \pm \textbf{0.012}$	$\textbf{0.577} \pm \textbf{0.018}$	$\textbf{0.718} \pm \textbf{0.026}$	$\textbf{0.701} \pm \textbf{0.019}$	$\textbf{0.640} \pm \textbf{0.011}$
RETAIN	0.685 ± 0.026	0.764 ± 0.014	0.578 ± 0.014	0.715 ± 0.031	0.697 ± 0.015	0.641 ± 0.010
RNN	0.681 ± 0.016	0.762 ± 0.021	0.577 ± 0.010	0.719 ± 0.036	0.696 ± 0.014	0.641 ± 0.012
НММ	0.666 ± 0.021	0.755 ± 0.031	0.551 ± 0.014	0.689 ± 0.021	0.665 ± 0.013	0.620 ± 0.009
MLP	0.657 ± 0.036	0.751 ± 0.056	0.553 ± 0.024	0.685 ± 0.052	0.656 ± 0.018	0.601 ± 0.012

Table 1: Performance of the different competing models for the 6 prognostic tasks under consideration.

designating the prediction targets at every hospital visit. A sigmoid transformation was applied to the top hidden layer.

RETAIN: An RNN-based reverse-time attention model proposed in Choi et al. (2016b). To ensure a fair comparison with PASS and the standard RNN benchmark, we implemented the attention layer of **RETAIN** via an LSTM with 2 hidden layers of size 200, and restricted its architecture to generate only visit-level attention. (This is equivalent to the **RNN**- α_M benchmark in Choi et al. (2016b).)

The baseline algorithms above are selected so as to highlight the added value of every modeling component in PASS. That is, an MLP only uses current information to predict the future clinical outcomes, whereas an HMM only looks 1 step back but provides a fully-fledged probabilistic model for disease progression. On the other hand an RNN capture more flexible dynamics (and is more memoryful) than an HMM but lacks interpretability, whereas **RETAIN** can provide explanations for its predictions, but does not explain the actual mechanisms of disease progression.

Implementation of PASS. We implemented the phased LSTM in Section 3 with 2 hidden layers of size 200 in order to match the model complexity of **RETAIN** and the standard RNN baseline. Hospital visits after which a death event happens within 3 years were explicitly labeled as the absorbing state D in our model. The number of states D was tuned via cross-validation to optimize the accuracy of predicting mortality events. PASS was implemented in **Tensorflow** Abadi et al. (2016), and the phased attention layer was implemented via tf.contrib.rnn.PhasedLSTMCell.

Prediction tasks and evaluation metric. All models were used to sequentially predict the 1-year risk for 6 prognostic tasks of predicting 3 comorbidities and 3 lung infections that are common in the CF population. The comorbidities are Allergic bronchopulmonary aspergillosis (ABPA), diabetes and intestinal obstruction. The lung infections are Klebsiella Pneumoniae, E. coli and Aspergillus. We used the area under the ROC curve (AUC-ROC) with 5-fold cross-validation for performance evaluation, where error counts are taken over all patients and all hospital visits.

Results. The AUC-ROC performance of all models on the 6 prognostic tasks under consideration is provided in Table 1. We note that CF is a very complex disease, for which patients encounter various possible comorbidities and are prescribed a wide variety of possible treatments. This leads to each patient having a very rich clinical history that influence their outcomes. Because HMMs fail to properly integrate the patients' rich clinical histories into the state dynamics, they displayed modest predictive performance on the 6 prognostic tasks. As we can see in Table 1, the HMM model did not provide any significant improvement over the static MLP model in any of the prognostic tasks. On the contrary, RNN-based model provided significant improvements over the static MLP model on all of the 6 tasks. The results in Table 1 show that the predictive accuracy of PASS is comparable to that of the RNN-based models. Note that the standard RNN model issue predictions without modeling disease progression, and hence it does not offer any interpretation benefit, whereas **RETAIN** provides the CF physiology (in terms of its latent progression stages), and hence it ensures the interpretational and modeling benefits of an HMM while maintaining the predictive accuracy of an RNN-based predictive model.

5.2 **Goal B**: UNDERSTANDING DISEASE PROGRESSION MECHANISMS

In this Section, we show that the PASS model successfully extracts meaningful clinical knowledge about the CF progression mechanisms. We also show how the PASS model parameters can inform clinical practitioners about the progression mechanisms for individual patients.

CF progression stages. Unlike many other chronic diseases, current clinical guidelines do not provide classifications for the progression stages of CF Szczesniak et al. (2017). In a completely unsupervised fashion, the PASS model successfully learned D = 3 progression stages (Stages I, II and III) that corresponded to clinically distinguishable levels of CF severity. The learned baseline Markov generator matrix for the transition rates among the 3 stages is given by:

$$\hat{\mathbf{\Lambda}} = \begin{bmatrix} -0.0578 & 0.0578 & 0\\ 0 & -0.0691 & 0.0691\\ 0 & 0 & 0 \end{bmatrix}.$$
(11)

From the baseline transition rates in (11), it follows that the average occupancy of a patient in Stage I is $\frac{1}{0.0578} = 17.30$ years, and the occupancy in Stage II is around $\frac{1}{0.0691} = 14.47$ years. That is, a typical patient who is born with CF progresses to Stage II by adulthood, before reaching the terminal stage by the age of 31. These figures match the survival rates in CF populations, where the median lifetime is known to be as low as 40 years of age McCarthy et al. (2013). Note that the baseline generator matrix in (11) only describes population-level rates of progression: individual variability among patients are captured via the patient-specific attention weights.



Figure 4: Depiction for the distributions of clinical observations and generated attention in the 3 learned progression stages.

The FEV1 % biomarker is the main spirometric measure of lung function that is currently used to guide clinical and therapeutic decisions. In order to check that the learned progression stages correspond to different levels of disease severity, we plot the estimated emission distribution for the FEV1 % biomarker in Stages I, II and III in Figure 4. As we can see from the emission distributions in Figure 4, the mean values of the FEV1 biomarker in each stage were 87%, 65% and 36%, respectively. This coincided with the current practice guidelines for referring critically-ill patients to a lung transplant, which recommends a transplant for patients with FEV1 < 30%, monitoring for a transplant for patients with FEV1 ranging from 30% to 80%, and no transplant for patients with FEV1 above 80% Braun & Merlo (2011). Thus, the learned progression stages can be translated into actionable information for clinical decision-making.

In Figure 4, we also obtain the maximum a posterior inferences for the progression stages of every individual patient as described in Section 3, and plot the average attention weights assigned to the patients' last 6 hospital visits in every progression stage. We found that the attentive dynamics tend to be less relevant for patients in Stage I, where most of the attention is allocated to the most recent visit. Memory starts getting more important in Stages II and III, where the attention weights allocated to older hospital visits gets higher. This can be explained by the fact that patients in Stages II and III are more likely to have been diagnosed with more comorbidities in the past, and hence more segments of their clinical history matters for predicting their outcomes.

Significance of attention weights for individual patients. How can clinicians interpret and make use of the generated attention weights for the patient at hand? One important way to utilize the attention weights is to reason about the effect of different treatment decisions and how their outcomes are impacted by the patient's history. For the sake of illustrating this point, in Figure 5 we pick an out-of-sample patient who has repeatedly visited the hospital over the years 2012, 2013 and 2014. We see the attention weights generated by the PASS model can inform the clinician about the potential efficacy of the Ivacaftor treatment (a gene targeted therapy Wainwright et al. (2015)) prescribed for this particular patient in the year 2014 by predicting the risk of progression to Stage III of lung function severity by the year 2015.

The inferred progression stage for the patient was Stage II for all of the 3 hospital visits. We toggle the patient's followup data vector X_3 (in year 2014) to let the variable indicating the prescription of the Ivacaftor drug be once set to 0 and once set to 1, and compute the probability of progressing to Stage III by the year 2015 in each case. We found that assigning Ivacaftor treatment to this patients is actually associated with an elevated risk of progressing to the severe Stage III within 1 year. By inspecting at the attention weights, we found that most attention is assigned to the most recent visit when Ivacaftor treatment is not prescribed, but the highest attention is paid to the follow-up data in 2012 when Ivacaftor is prescribed.



Figure 5: Usage of the PASS attention weights to reason about treatment decisions.

The patient's follow-up data in year 2012 included a diagnosis with Liver disease. Hence, the elevated risk of progression upon taking the Ivacaftor treatment may be linked to its side effects concerning liver function complication, which may exacerbate the patient's liver disease. The PASS model altered the state dynamics to take into account the 2-year old follow-up data that is only important in determining the state dynamics conditional on prescription of an Ivacaftor treatment.

APPENDIX

DETAILED STRUCTURE OF THE EHR DATA

A patient's EHR record, denoted as \mathcal{E} , is a collection of timestamped follow-up data gathered during repeated, irregularly-spaced hospital visits, in addition to static features of the patient (e.g., genetic variables). We represent a given patient's EHR record as follows:



```
Static features Visit times
```

Treatments Anchors Observations

where Y is the static features' vector, X_m is the follow-up data collected in the m^{th} hospital visit, t_m is the time of the m^{th} visit, and M is the total number of hospital visits. (The time-horizon t is taken to be the patient's chronological age.) The follow-up data vector X_m comprises three components:

- <u>Treatments indicator</u> $(u_m \in \{0,1\}^U)$: A binary vector indicating the prescription of (a subset of) U possible treatments to the patient during the m^{th} hospital visit.
- <u>Anchor findings</u> $(C_m \in \{0,1\}^K)$: A binary vector indicating the presence of concrete diagnoses (i.e., ICD or HCPCS codes Blumenthal & Tavenner (2010)) for K distinct comorbidities that may co-occur with the target disease.
- <u>Clinical observations</u> $(O_m \in \mathbb{R}^O)$: A set of laboratory-measured biomarkers that reflect the severity of the target disease.

An EHR dataset \mathcal{D} is an assembly of records for N patients, i.e., $\mathcal{D} = \{\mathcal{E}^{(i)}\}_{i=1}^{N}$. Figure 6 provides an illustration for the structure of the EHR data.



Figure 6: Illustration for the structure of the EHR data.

COMPARISON BETWEEN MARKOVIAN AND ATTENTIVE STATE DYNAMICS

Because Markovianity simplifies inference, most existing models of disease progression are based on HMMs (e.g., Alaa et al. (2017), Liu et al. (2015), and Wang et al. (2014)). However, memoryless Markovian dynamics hinder a model's capacity for "explaining" individual-level progression trajectories. This is because under a Markov model, all patients at the same stage of progression would have the same expected future trajectory, irrespective of their potentially different individual clinical histories (i.e., the timing and order of treatments and comorbidities). That is, a Markov model captures the *population-level* transition rates among progression stages, but explains away *individual-level* variations in progression trajectories through the randomness of the transition probability $P(\mathbf{Z}_m | \mathbf{Z}_{m-1}, t_m - t_{m-1})$. This can render Markov models highly misleading since clinical actions are taken on an individual basis.

Interpreting the attentive state dynamics Now we use an illustrative example to show how a clinician can interpret the attentive state dynamics for individual patients, and highlight the information that is missed by Markovian dynamics but can be captured via attentive dynamics.

In Figure 7, we display exemplary progression trajectories for 2 chronic kidney disease (CKD) patients through the unrolled graphical model of the DBN in (3). With a slight abuse of graphical model notation, we let the thickness of the arrows connecting states be proportional to the attention weights generated for predicting the state transition in the third hospital visit. Patients A and B have identical trajectories at the first 2 visits (both are diagnosed with hypertension and are in the same progression stages), with one exception being that Patient A is administered a medication for hypertension (ACE inhibitors) in the first visit. Patient B transits from stage 2 to stage 3 CKD because of hypertensive renal complications, whereas patient A stays in stage 2 thanks to the medication. The attentive model can capture the difference between the 2 trajectories by paying attention to the first visit for Patient A (when the medication was prescribed), and little attention to the same visit for Patient B.



Figure 7: Depiction for the attentive state dynamics.

By visually inspecting the attention weights assigned to past states of each individual patient, clinicians can interpret the decreased risk for Patient A (compared to Patient B) to be a result of the clinical events that Patient A encountered in the first visit (i.e., administration of ACE inhibitors). On the contrary, a memoryless Markov model would not be able to distinguish the different trajectories that Patients A and B exhibit as both patients are in Stage 2 CKD during the second visit.

REFERENCES

- Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. Tensorflow: a system for large-scale machine learning. In *OSDI*, volume 16, pp. 265–283, 2016.
- Ahmed M Alaa and Mihaela van der Schaar. A hidden absorbing semi-markov model for informatively censored temporal data: Learning and inference. *Journal of Machine Learning Research*, 2018.
- Ahmed M Alaa, Scott Hu, and Mihaela van der Schaar. Learning from clinical judgments: Semimarkov-modulated marked hawkes processes for risk prognosis. *International Conference on Machine Learning*, 2017.
- Justin Bayer and Christian Osendorfer. Learning stochastic recurrent networks. arXiv preprint arXiv:1411.7610, 2014.

- Ron Begleiter, Ran El-Yaniv, and Golan Yona. On prediction using variable order markov models. *Journal of Artificial Intelligence Research*, 22:385–421, 2004.
- David Blumenthal and Marilyn Tavenner. The meaningful use regulation for electronic health records. *New England Journal of Medicine*, 363(6):501–504, 2010.
- Andrew T Braun and Christian A Merlo. Cystic fibrosis lung transplantation. Current opinion in pulmonary medicine, 17(6):467–472, 2011.
- Edward Choi, Mohammad Taha Bahadori, Andy Schuetz, Walter F Stewart, and Jimeng Sun. Doctor ai: Predicting clinical events via recurrent neural networks. In *Machine Learning for Healthcare Conference*, pp. 301–318, 2016a.
- Edward Choi, Mohammad Taha Bahadori, Jimeng Sun, Joshua Kulas, Andy Schuetz, and Walter Stewart. Retain: An interpretable predictive model for healthcare using reverse time attention mechanism. In *Advances in Neural Information Processing Systems*, pp. 3504–3512, 2016b.
- Junyoung Chung, Kyle Kastner, Laurent Dinh, Kratarth Goel, Aaron C Courville, and Yoshua Bengio. A recurrent latent variable model for sequential data. In Advances in neural information processing systems, pp. 2980–2988, 2015.
- Hanjun Dai, Bo Dai, Yan-Ming Zhang, Shuang Li, and Le Song. Recurrent hidden semi-markov model. *International Conference on Learning Representations*, 2016.
- Marco Fraccaro, Søren Kaae Sønderby, Ulrich Paquet, and Ole Winther. Sequential neural models with stochastic layers. In Advances in neural information processing systems, pp. 2199–2207, 2016.
- Matthew Johnson, David K Duvenaud, Alex Wiltschko, Ryan P Adams, and Sandeep R Datta. Composing graphical models with neural networks for structured representations and fast inference. In *Advances in neural information processing systems*, pp. 2946–2954, 2016.
- Maximilian Karl, Maximilian Soelch, Justin Bayer, and Patrick van der Smagt. Deep variational bayes filters: Unsupervised learning of state space models from raw data. *arXiv preprint arXiv:1605.06432*, 2016.
- Andrei Nikolaevich Kolmogorov. *Foundations of the Theory of Probability: Second English Edition*. Courier Dover Publications, 2018.
- Rahul G Krishnan, Uri Shalit, and David Sontag. Structured inference networks for nonlinear state space models. In *AAAI*, pp. 2101–2109, 2017.
- Bryan Lim and Mihaela van der Schaar. Disease-atlas: Navigating disease trajectories with deep learning. *Machine Learning for Healthcare Conference (MLHC)*, 2018.
- Zachary C Lipton. The mythos of model interpretability. *ICML Workshop on Human Interpretability* of Machine Learning, 2016.
- Zachary C Lipton, David C Kale, Charles Elkan, and Randall Wetzel. Learning to diagnose with lstm recurrent neural networks. *International Conference on Learning Representations*, 2016.
- Yu-Ying Liu, Shuang Li, Fuxin Li, Le Song, and James M Rehg. Efficient learning of continuoustime hidden markov models for disease progression. In Advances in neural information processing systems, pp. 3600–3608, 2015.
- Fenglong Ma, Radha Chitta, Jing Zhou, Quanzeng You, Tong Sun, and Jing Gao. Dipole: Diagnosis prediction in healthcare via attention-based bidirectional recurrent neural networks. In Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 1903–1911. ACM, 2017.
- Cormac McCarthy, Borislav D Dimitrov, Imran J Meurling, Cedric Gunaratnam, and Noel G McElvaney. The cf-able score: a novel clinical prediction rule for prognosis in patients with cystic fibrosis. *Chest*, 143(5):1358–1364, 2013.

- Kevin Patrick Murphy and Stuart Russell. Dynamic bayesian networks: representation, inference and learning. 2002.
- Daniel Neil, Michael Pfeiffer, and Shih-Chii Liu. Phased lstm: Accelerating recurrent network training for long or event-based sequences. In Advances in Neural Information Processing Systems, pp. 3882–3890, 2016.
- Mary Ann Sevick, Jeanette M Trauth, Bruce S Ling, Roger T Anderson, Gretchen A Piatt, Amy M Kilbourne, and Robert M Goodman. Patients with complex chronic diseases: perspectives on supporting self-management. *Journal of general internal medicine*, 22(3):438–444, 2007.
- Rhonda D Szczesniak, Dan Li, Weiji Su, Cole Brokamp, John Pestian, Michael Seid, and John P Clancy. Phenotypes of rapid cystic fibrosis lung disease progression during adolescence and young adulthood. *American journal of respiratory and critical care medicine*, 196(4):471–478, 2017.
- Jose M Valderas, Barbara Starfield, Bonnie Sibbald, Chris Salisbury, and Martin Roland. Defining comorbidity: implications for understanding health and health services. *The Annals of Family Medicine*, 7(4):357–363, 2009.
- Claire E Wainwright, J Stuart Elborn, Bonnie W Ramsey, Gautham Marigowda, Xiaohong Huang, Marco Cipolli, Carla Colombo, Jane C Davies, Kris De Boeck, Patrick A Flume, et al. Lumacaftor–ivacaftor in patients with cystic fibrosis homozygous for phe508del cftr. *New England Journal of Medicine*, 373(3):220–231, 2015.
- Xiang Wang, David Sontag, and Fei Wang. Unsupervised learning of disease progression models. In Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining, pp. 85–94. ACM, 2014.
- Frans MJ Willems, Yuri M Shtarkov, and Tjalling J Tjalkens. The context-tree weighting method: basic properties. *IEEE Transactions on Information Theory*, 41(3):653–664, 1995.