Factorial HMMs with Collapsed Gibbs Sampling for Optimizing Long-term HIV Therapy

Amit Gruber[1,†], **Chen Yanover**[1], **Tal El-Hay**[1], **Anders Sönnerborg**[2]
Vanni Borghi[3], **Francesca Incardona**[4], **Yaara Goldschmidt**[1]
[1]IBM Research Haifa Research Lab
[2]Karolinska Institute Karolinska University Hospital
[3]Modena University Hospital
[4]EuResist Network GEIE InformaPro S.r.l.

Abstract

Combined antiretroviral therapies (CART) can successfully suppress HIV in the serum and bring its viral load below detection rate. However, drug resistance remains a major challenge. As resistance patterns vary between patients, personalized therapy is required. Automatic systems for therapy personalization exist and were shown to better predict therapy outcome than HIV experts in some settings. However, these systems focus only on selecting the therapy most likely to suppress the virus for several weeks, a choice that may be suboptimal over the longer term due to evolution of drug resistance.

We present a novel generative model for HIV drug resistance evolution. This model is based on factorial HMMs, applying a novel collapsed Gibbs Sampling algorithm for approximate learning. Using the suggested model, we obtain better therapy outcome predictions than existing methods and recommend therapies that may be more effective over the long term. We demonstrate our results using simulated data and using real data from the EuResist dataset.

1 Introduction

Much progress has been made in recent years in treating HIV and modern Combined Antiretroviral Therapies (CART) can successfully suppress the virus and prolong patients' life. These therapies consist of multiple drugs, typically 2-3, that target different stages of the HIV reproduction mechanism. When taken properly, these drugs can suppress the virus and reduce its level in the blood below detection level. However, the virus is suppressed but not eradicated from the patient's body, and it still resides in viral reservoirs Pierson et al. [2000], also called latent reservoirs. These are cells infected by the virus but not actively producing it, and they are not affected by drugs targeting the virus replication process. These cells contain integrated HIV-1 DNA. When they become active, they replicate the virus, contributing to its persistence in the body despite it being suppressed in the blood. If the therapy is stopped, or if the virus develops resistance to the drug, its level in the blood will rise again.

When HIV replicates, new mutations are randomly created. Some of these mutations may show drug resistance, and prevail under therapy. HIV may acquire drug resistance through such mutations over time, and moreover, an initial infection by a drug resistant strain can also be transmitted from another person Ocfemia et al. [2015], Van de Vijver et al. [2006]. Due to the prevalence of drug resistant mutations, care must be taken when selecting a therapy. As resistance patterns vary between different patients, therapy selection need to be personalized. Such personalization is done by following guidelines AIDSInfo [2016], and recently also using genotypic resistance tests.

Due to the difficulty and importance of selecting an effective therapy, machine learning decision support systems for therapy selection have been proposed. Rosen-Zvi et al. [2008] and Prosperi et al. [2010] use the EuResist dataset[1] to train a system combining 3 prediction engines. Each of these engines applies logistic regression on a set of clinical and demographic features to predict the probability of success for a given treatment. A study by Zazzi et al. [2011] shows that this system is better at predicting therapy outcome than human HIV experts. After random forests were

Proceedings of the 21st International Conference on Artificial Intelligence and Statistics (AISTATS) 2018, Lanzarote, Spain. PMLR: Volume 84.

[†]Author is currently with Facebook Inc.

[1]http://www.euresist.org

shown to outperform SVM and artificial neural networks in Wang et al. [2009], Revell et al. [2014] used the HIV-RDI dataset[2] to train random forest using clinical features to predict therapy outcome. They showed that prediction accuracy is not harmed when genotype information is not available. A similar result was published by Prosperi et al. [2010].

While these systems are successful in optimizing therapy outcome in a range of several weeks (8-24 weeks), their performance significantly deteriorates when predicting therapy outcome in a longer range (48 weeks). Further, these systems do not consider the effect of the selected therapy on the evolution of future drug resistance and may recommend therapies that are suboptimal in a longer term. In a recent study, Prosperi et al. [2016] study therapy time to failure. Aiming to optimize treatments outcome in the long term, in this paper we suggest a novel generative model for drug resistance evolution.

With the increasing availability of electronic health records (EHR), graphical models have been suggested to model disease progression: Krishnan et el. [2016] model disease progression under therapy using an HMM (Figure 1a). The hidden states in this HMM correspond to the patient's health status and the observations are clinical measurements included in the EHR. Patient's health is affected by the known therapy the patient receives, whose effect is noticeable after some period of time. El-Hay et al. [2014] propose structured proportional jump process for non-homogenous data, in which the model factorizes into an element that depends on time and an element that depends on system configuration. They apply their model to diabetes and HIV data. Arora and Dixit [2009] use a HMM to model a certain aspect of HIV disease progression. Specifically, they model the evolution of drug resistance against a single specific drug.

We follow a similar approach and model drug resistance evolution in HIV patients using an HMM. Unlike Arora and Dixit [2009], we consider CART therapies consisting of multiple drugs. Using a standard HMM to model resistance against multiple drugs would require an exponentially large state space to represent all the combinations of possible virus resistance against each one of these drugs. Further, as therapy may be altered over time, such an HMM would need to represent the possible virus resistance against each one of the drugs in the arsenal. Instead of using a standard HMM, we use a factorial HMM Ghahramani et al. [1997] consisting of multiple chains interconnected through the observation (Figure 1b). In our case, each chain corresponds to a possible resistance with respect to a specific drug and its evolution over time, and the observation connecting the chains is the therapy outcome. Using this factorized representation, we avoid the exponential state space representation. The main contributions of this paper are: (1) A novel graphical model for modeling disease progression under combined drug therapy. Specifically, we apply this model to the evolution of drug resistance in HIV patients taking CART (therapy) and model patient's adherence to the therapy. (2) We use this graphical model to recommend a series of treatments that are effective over the long-term. (3) A novel collapsed Gibbs Sampling algorithm for Factorial HMM in general and for the suggested model specifically, different from the previously suggested approximate learning approaches including Gibbs sampling (non-collapsed).

2 FResist: A Generative Model for Drug Resistance Evolution

HIV is typically treated with combined antiretroviral therapy (CART): a combination of multiple (usually 2-3) compounds given in a single pill that needs to be taken daily by the patient. If taken properly, and if the virus is sensitive to at least one of the compounds in the CART, the therapy may succeed and prevent the virus from multiplying (but not eradicate it from reservoirs). If the virus is resistant to all the compounds in the CART, the therapy is expected to fail.

We suggest a novel generative model for the evolution of HIV drug resistance under CART. This model, which we call FResist, is a special case of factorial HMM where each of its K chains represents the sensitivity or resistance of the virus to a specific compound k and its possible evolution. We first describe the components of a single chain and the relation between them. Then we describe the interchain relation to the treatment outcome and provide a full generative process. In Section 2.1 we extend the model to account for patient adherence.

HIV acquires drug resistance through mutations. When the virus replicates, mutations are randomly created. When no treatment is taken by the patient, the wild type, which is advantageous to other mutations, outgrows and dominates other mutants. But in the presence of drugs, the wild type is suppressed and drug resistant mutations prevail Clavel and Hance [2004]. Once drug resistant mutants are present in a significant amount in the blood stream, they may form latent reservoirs - collections of immune cells infected by the virus. These latent reservoirs are not affected by anti-HIV drugs as these only target the virus replication process Pierson et al. [2000]. Once drug resistant mutants are present in the reservoirs, this may be viewed as acquisition of permanent resistance. We begin with a detailed

[2]https://www.hivrdi.org

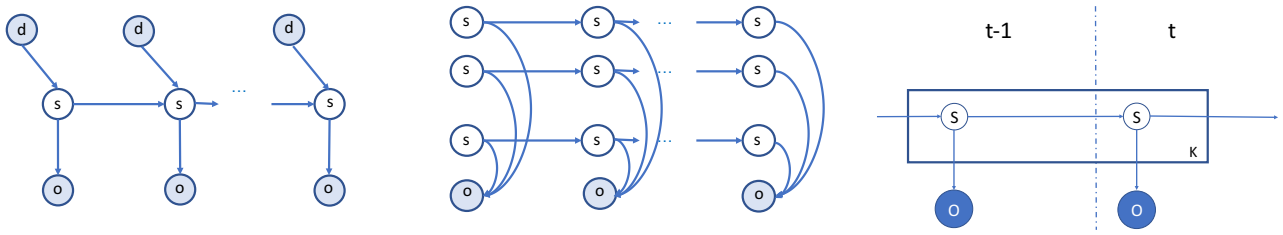

(a) (b) (c)

Figure 1: **a.** Markov model of a patient under therapy. The therapy affects the patient's health after a period of time; this effect may be observed through clinical measurements. **b.** Factorial HMM Ghahramani et al. [1997]. **c.** Factorial HMM in plates notation.

description of this process for a single patient, omitting the patient specific index for simplicity of notation. We describe how mutations may be created, how resistance may be acquired, the relation to the therapy outcome and the transition to time $t+1$. Then we describe the entire generative process.

Focusing on a single chain k, let $d_{t,k}$ be a binary variable denoting whether drug k was taken at time t and let O_t be the multi-drug treatment outcome ($O_t = 1$ for a successful treatment). Let $m_{t,k}$ be a binary variable representing the existence of mutations resistant to drug k in the serum, and let $R_{t,k}$ be a binary variable representing whether permanent resistance to drug k already exists (through viral reservoirs) at time t. If permanent resistance was already acquired at time t (i.e. $R_{t,k} = 1$), then the drug resistant strain may also be found in the serum: $m_{t,k} = 1$. Otherwise, when $R_{t,k} = 0$, when drug k is taken ($d_{t,k} = 1$), new drug resistant mutations may appear with probability p_k^M. If the drug is not taken at time t, no such mutations will prevail. Equation 1 summarizes this relation:

$$\Pr(m_{t,k} = 1 | R_{t,k}, d_{t,k}) = \begin{cases} 1 & \text{if } R_{t,k} = 1 \\ p_k^M & \text{if } R_{t,k} = 0, d_{t,k} = 1 \\ 0 & \text{otherwise} \end{cases} \tag{1}$$

When drug resistant mutations appear in the serum ($m_t = 1$) and are not suppressed ($O_t = 0$), these may form reservoirs, and the virus will consequently acquire permanent resistance to drug k ($R_{t+1,k} = 1$). We denote the probability for this event by p^C. There are two conditions for this situation: (1) Mutations resistant to drug k have emerged, and (2) The treatment fails, i.e., no other drug taken at time t is successful in suppressing the virus. It is also possible, that viral reservoirs resistant to drug k had already existed at time t and will be preserved at time $t+1$. Finally, there is also a low probability for a new infection by a strain resistant to drug k between time t and $t+1$ leading to the creation of a new viral reservoir. We denote the probability for this event by p^I. In practice, we assume $p^C = 1 - \epsilon$ and $p^I = \epsilon$ for a small fixed ϵ.

$$\Pr(R_{t+1,k} = 1 | R_{t,k}, m_{t,k}, O_t) = \begin{cases} 1-\epsilon & \text{if } R_{t,k} = 1 \\ 1-\epsilon & \text{if } R_{t,k} = 0, \\ & \quad m_{t,k} = 1, \\ & \quad O_t = 0 \\ \epsilon & \text{otherwise} \end{cases} \tag{2}$$

We place a prior $p_k^{R_0} = \Pr(R_{1,k} = 1)$ on resistance that had already been acquired by the patient before the first recorded treatment. Such resistance may be acquired due to past treatments missing from our data before $t = 1$, or due to an initial infection by a drug resistant strain Ocfemia et al. [2015], Van de Vijver et al. [2006].

In combined antiretroviral therapy, the therapy is expected to succeed if the virus is sensitive to at least one of the compounds in it, for at least one k, $m_{t,k} = 0$. To account for deviations from this model and since medical records are prone to errors, we model the observed outcome as a noisy version of this expected outcome: Let O_t^E be the expected therapy outcome and let O_t be the therapy outcome observed in the EHR data, then

$$O_t^E = \bigvee_{\{k: d_{t,k}=1\}} \neg m_{t,k}$$

$$\Pr(O_t = O_t^E | \vec{m}_t, \vec{d}_t) = 1 - p^N \tag{3}$$

At any given time t, the observed treatment outcome, O_t, depends only on a small number of variables $m_{t,k}$ (typically $2-3$) associated with the drugs in the CART.

The generative process is as follows:

1. For all drug compounds $k \in 1, \ldots, K$
   (a) draw mutation probability $p_k^M \sim Beta(\beta)$
   (b) draw prior resistance probability $p_k^{R_0} \sim Beta(\gamma)$
2. draw outcome noise probability $p^N \sim Beta(\eta)$

3. For all patients i=1,...,N
   (a) for t=1,...,T
      i. for all drugs k=1,..,K
         A. if $t = 1$, draw $R_{i,1,k} \sim p_k^{R_0}$, else draw $R_{i,t,k} \sim \Pr(R_{i,t,k}|R_{i,t-1,k}, O_{i,t-1}, m_{i,t-1,k})$
         B. draw $m_{i,t,k}$ conditioned on $R_{i,t,k}$, $d_{i,t,k}$
      ii. draw treatment outcome O_t conditioned on $\vec{m}_{i,t}, \vec{d}_{i,t}$

where N is the number of patients and β, γ, and η are Beta priors that we treat as hyper-parameters of the model. This model is depicted in Figure 2(b).

2.1 Patient Adherence

The generative process in Section 2 assumes full knowledge of the drugs taken by the patient, i.e. that patients take drugs as prescribed. However, in real life, patients often fail to adhere to the prescribed drug regimen Kim et al. [2014], Glass et al. [2006]. Therefore, when a therapy fails, there is now uncertainty if this is due to drug resistance or due to lack of adherence.

We extend the model described in the previous section to account for patient adherence: Let $a_{i,t}$ be a binary variable denoting whether patient i has taken the prescribed drugs at time t, and let $p_i^A = Pr(a_{i,t} = 1)$ be the patient specific adherence probability. Since modern CART therapies typically use a single pill that contains all therapy drugs, all drugs can either be taken or not at a specific time t, and a single adherence variable $a_{i,t}$ for all drugs is sufficient. We extend the generative process to include adherence by adding the following steps as a prior process before the generative process described in Section 2:

1. For all patients $i \in 1, \ldots, N$
   (a) draw $p_i^A \sim Beta(\alpha)$
   (b) for t=1,...,T
      i. draw $a_{i,t} \sim p_i^A$
      ii. set $d_{i,t,k} = a_{i,t} \wedge s_{i,t,k}$ for all $k = 1 \ldots K$

where $s_{i,t,k}$ is a binary variable indicating whether drug k was prescribed to patient i under the therapy at time t and α is a Beta hyper-parameter. Note that these additional steps only set $d_{i,t,k}$ to the drugs taken in practice by patient i at time t which are now unobserved. Once these drugs have been set, the generative process proceeds as before. This extended model is depicted in Figure 2(c). This extension also modifies the factorial structure of the chain by adding a new interchain dependency at each time point, as the adherence $a_{i,t}$ affects all $d_{i,t,k}$ for all therapy drugs (for which $s_{i,t,k} = 1$).

2.2 Related Models

FResist is closely related to Factorial HMM Ghahramani et al. [1997], with the addition of observed treatment, and an effect of the observed therapy outcome at time t on the hidden states at time $t + 1$. FResist specifies a structure on the hidden states of each one of the chains induced by the relations between R and m. Further, since CARTs consist of only 2-3 drugs, the observation in FResist at each time (therapy outcome) depends on only 2-3 chains. The two models are depicted in figures 1(c) and 2(b). Modeling patient adherence adds a dependency between the chains, this time through a latent variable (Figure 2(c)).

Coupled HMMs Brand et al. [1997] are a related class of models where the chains interact during their progression. However, each chain has its own observations: $s_{t,i}$, the state of chain i at time t depends on the states of other chains, but the observation $o_{t,i}$ depends only on $s_{t,i}$. Pan et al. [2012] model influence by human interaction in a social network using coupled HMMs and employ variational EM algorithm for approximate learning. Dong et al [2012] and Fan et al. [2015] model disease infection in a social network using Graph-coupled HMM, a special case of coupled HHM. They take advantage of a sparse network structure to derive a Gibbs Sampling algorithm for a disease spread model. Dong et al [2016] developed variational inference methods for a class of coupled HMM models for spread of an epidemic.

Our work differs from those mentioned above in the structure of the model, the algorithm we use for approximate learning and the application.

3 Approximate Learning

Similar to Factorial HMM and coupled HMM, exact learning in FResist is intractable. Various approaches have been suggested for approximate learning in these models, including EM and variational EM Ghahramani et al. [1997], Pan et al. [2012], Dong [2016], EM with Gibbs sampling Ghahramani et al. [1997], Fan et al. [2015] and Gibbs Sampling for graphs with sparse connectivity between the chains Dong et al. [2012].

In this paper we suggest a novel *Collapsed Gibbs Sampling* algorithm for factorial HMMs in general and for FResist in particular. It differs from Gibbs sampling as in Dong et al. [2012] by integrating out the continuous variables rather than sampling them. This approach is commonly applied in topic models such as LDA Griffiths and Steyvers [2004] and in the context of a single HMM Griffiths et al. [2004]. We derived the sampling equations for the general case of factorial HMM and we provide them in Appendix A. Here we describe the

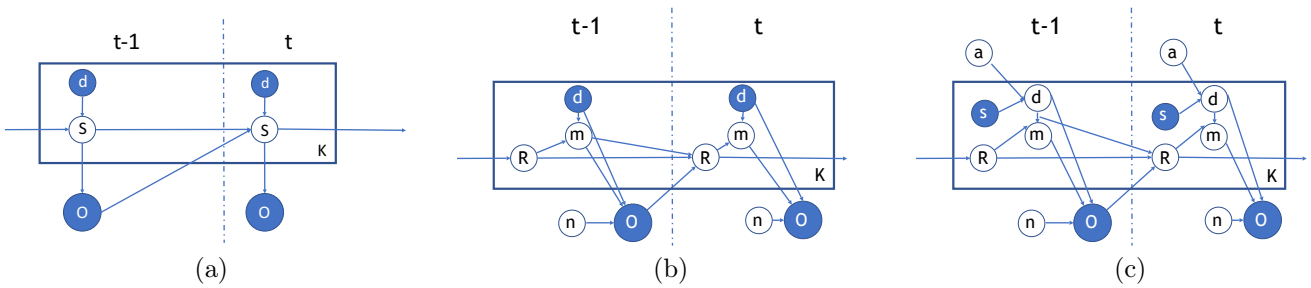

Figure 2: **a.** Multi-drug therapy model. Notice the addition of drugs (d) and the effect of treatment outcome on the patients health. **b.** The suggested model, FResist, is a special case of the model in (a). **c.** FResist with patient adherence (a).

sampling equations for the special case of FResist.

3.1 Collapsed Gibbs Sampling for FResist

In collapsed Gibbs Sampling Liu [1994], the discrete variables are sampled while the continuous variables are integrated out. In the case of FResist, the discrete variables need to be sampled are m and R, and in the adherence model of Section 2.1 also a [3]. The continuous variables that integrated out are p^M, p^{R0}, p^N and in the case of adherence also p^a. We assume that the Beta priors are known and set $\epsilon = 0.01$[4].

We iterate over all patients and their treatments t, and at each iteration sample $\vec{R}_t, \vec{m}_t, a_t$ of a specific patient conditioned on all other variables (omitting again the patient specific subscripts for ease of notation). Due to the determinstic relations between $\vec{R}_t, \vec{m}_t, a_t$ in some cases, (e.g. $m_{t,k} = 1$ and $R_{t,k} = 0$ implies $a_{t,k} = 1$) we sample these $2K + 1$ variables together as a block. We use the factorization of the joint posterior distribution given in equation 4 to sample them efficiently (in time linear, rather than exponential, in the number of prescribed drugs). $\vec{R}_t, \vec{m}_t, a_t$ are sampled conditioned on all other variables, observations and hyper parameters. To simplify and shorten the equations below, we only explicitly write the variables on which $\vec{R}_t, \vec{m}_t, a_t$ depend (and omit all others. We also omit the hyper parameters and observed prescriptions on which we always condition).

$$\begin{aligned} &\Pr(\vec{R}_t, \vec{m}_t, a_t | \vec{R}_{t-1}, \vec{R}_{t+1}, O_{t-1}, O_t) \qquad (4)\\ &= \Pr(\vec{R}_t | \vec{m}_t, a_t, \vec{R}_{t-1}, \vec{R}_{t+1}, O_{t-1}) \\ &\Pr(\vec{m}_t | \vec{R}_{t-1}, \vec{R}_{t+1}, a_t, O_{t-1}, O_t,) \\ &\Pr(a_t | O_{t-1}, O_t, \vec{R}_{t-1}, \vec{R}_{t+1}) \end{aligned}$$

[3] $d_{t,k}$ is simply computed from a_t and the observed $s_{t,d}$; formally a_t and $\vec{d}_t$ are block sampled together.

[4] ϵ (or p^C and p^I) could be integrated out exactly like the other probabilities in the model. In practice, setting $\epsilon = 0.01$ proved sufficient and simplifies the presentation of the sampling algorithm.

The factorization given by equation 4 implies that we first sample a_t averaging over the $2K$ variables $\vec{R}_t, \vec{a}_t$, then sample $\vec{m}_t$ conditioned on the already sampled a_t and average over the $\vec{R}_t$, and finally sample $\vec{R}_t$ conditioned on $a_t, \vec{m}_t$.

We begin with sampling $\vec{R}_t$: The individual elements $R_{t,k}$ of $\vec{R}_t$ are conditionally independent given $a_t, m_t, R_{t+1}, R_{t-1}, O_{t-1}, O_t$ and we can sample them individually:

$$\begin{aligned} &\Pr(R_{t,k} | R_{t-1,k}, R_{t+1,k}, m_t, a_t, O_{t-1}, O_t) \qquad (5)\\ &\propto \Pr(R_{t+1} | R_{t,k}, O_t, m_{t,k}) \\ &\Pr(m_{t,k} | R_{t,k}, a_t) \\ &\Pr(R_{t,k} | R_{t-1,k}, O_{t-1}) \end{aligned}$$

where the terms in equation 5 can be computed using equations 2, and 11.

For sampling $\vec{m}_t$, we use the factorization $\Pr(\vec{m}_t | \ldots) = \prod_k \Pr(m_{t,k} | m_{t,k'<k}, \ldots)$ to individually sample $m_{t,k}$ conditioned on $m_{t,k'}$ for all $k' < k$ and averaged over $m_{t,k'}$ for all $k' > k$ as described in equations 6 - 7.

$$\begin{aligned} &\Pr(m_{t,k} | \vec{R}_{t-1}, \vec{R}_{t+1}, a_t, O_t, O_{t-1}, m_{t,k'<k}) \qquad (6)\\ \propto &\sum_{R_{t,k}} \Pr(R_{t+1,k} | m_t, O_t, R_{t,k}) \\ &\Pr(O_t | m_{t,k'\le k}, \vec{R}_{t-1,k'>k}, \vec{R}_{t+1,k'>k}) \\ &\Pr(m_{t,k} | a_t, R_{t,k}) \Pr(R_{t,k} | R_{t-1,k}, O_{t-1}) \end{aligned}$$

Computing $\Pr(O_t | m_{t,k}, \vec{R}_{t-1}, \vec{R}_{t+1}, m_{t,k'<k})$ requires averaging over all configurations of $m_{t,k'>k}$ (an exponentially large number). We show in equation 7 how to compute it in linear time. The other factors from equation 6 are computed using equations 2 and 11.

$$\begin{aligned} &\Pr(O_t | m_{t,k'\le k}, O_{t-1}, \vec{R}_{t-1,k'>k}, \vec{R}_{t+1,k'>k}) \qquad (7)\\ \propto &\sum_{o=0,1} \sum_{m \in M_o} \Big[\Pr(O_t | O^E(\vec{m}_t = o)) \\ &\Pr(\vec{m}_{t,k'>k} | a_t, R_{t-1,k'>k}, O_{t-1}) \Pr(R_{t+1,k'>k} | O_t, m_{t,k'}) \Big] \end{aligned}$$

where $M_1 = \{m : O^E(m) = 1\}$ and $M_0 = \{m : O^E(m) = 0\}$ are the sets of all configurations of $\vec{m}$ that yield the specific expected outcome O^E. The sum over M_0 includes the single term where all $m_k = 1$ (for drugs prescribed at time t). The other sum, over M_1 includes all other configurations. We compute it in time linear in the number of drugs in the therapy using the factoriaztion:

$$\Pr(\vec{m}_t|\vec{R}_{t-1}, a_t)\Pr(\vec{R}_{t+1}|O_t, \vec{m}_t, a_t) = \quad (8)$$

$$\prod_k \Big[\sum_{R_k} \Pr(m_{t,k}|R_{t,k}, a_t)\Pr(R_{t+1,k}|O_t, m_{t,k}, R_{t,k})$$

$$\Pr(R_{t,k}|R_{t-1,k}, O_{t-1})\Big]$$

Similarly, a_t is sampled while averaging over $\vec{m}_t, \vec{R}_t$ in linear time:

$$\Pr(a_t|O_{t-1}, O_t, \vec{R}_{t-1}, \vec{R}_{t+1}) \quad (9)$$

$$\propto \Pr(a_t)\Pr(O_t|a_t, \vec{R}_{t-1}, \vec{R}_{t+1}) \quad (10)$$

where $\Pr(O_t|a_t, \vec{R}_{t-1}, \vec{R}_{t+1})$ is computed using equation 7 setting $k' = 0$ (i.e. averaging over all m_k, R_k).

The other probabilities required for sampling according to equations 5-7 are computed from counts of sampled variables:

$$\Pr(m_{t,k} = m|R_{t,k} = 0, a_t = 1) = \frac{C^M_{k,m} + \beta}{\sum_{m'} C^M_{k,m'} + 2\beta} \quad (11)$$

$$\Pr(O_t|\vec{m}_t) = \frac{C_{O_t = O^E_t(m)} + \eta}{\sum_n C_n + 2\eta} \quad (12)$$

$$\Pr(R_{1,k} = r) = \frac{C^{R_0}_{k,r} + \gamma}{\sum_{r'} C^{R_0}_{k,r'} + 2\gamma} \quad (13)$$

$$\Pr(a_{i,t} = a) = \frac{C^A_{i,a} + \alpha}{\sum_{a'} C^A_{i,a'} + 2\alpha} \quad (14)$$

Where $C^M_{k,m}, C^{R_0}_{k,r}, C^A_{i,a}$ are the counts of the sampled variables $m_{\cdot,k}, n_{i,t}, R_{0,k}, a_i$ respectively excluding $\vec{m}_{i,t}, \vec{R}_{0,i,k}, a_{i,t}$, and C_n counts the number of times O_t is equal or different from O^E_t (computed from the sampled m's).

To summarize, despite the coupling of the chains through patient adherence and therapy outcome, we have derived a sampling algorithm linear in the number of variables to be sampled. This is in contrast to sampling in a general factorial HMM (see appendix A in the supplementary material for details).

4 Planning Long-Term Therapies

Having modeled drug resistance evolution using FResist and having learned the model dynamics, we proceed to using the model for therapy design. Our goal is to plan an entire series of treatments $a_1, \ldots, a_T$, possibly different from each other, that would optimize patient's health over the entire time period (in contrast to optimizing only for the first treatments). Due to the Markovian structure of FResist, optimizing such a sequence of treatments is a finite Markov Decision Process (MDP) that can be solved using dynamic programming to find a policy that would maximize the total cumulative reward[5]. Let $q_t = (m, R, a)_t$ the state at time t, $a_t \in A$ the action (a combination of drugs) we take at time t and $O_t(q_t, a_t)$ the outcome of the treatment a_t if the patient is at state q_t at time t. We define the immediate reward of a treatment a_t to be 1 if the treatment is successful and 0 otherwise, and the expected reward of a state q_t and action a_t is $E(O_{t'}|q_t, a_t, \ldots, a_T)$. We define the value of state q_t at time t to be the expected outcome of the best treatment (optimized over all possible actions): Our goal is to find the sequence of actions (therapies) yielding an overall maximal reward (maximal cumulative patient health along the entire time period).

$$V_t(q_t) = \max_{at,\ldots,a_T} \sum_{t'=t}^{T} E(O_{t'}|q_t, a_t, \ldots, a_T) \quad (15)$$

With this definition, it follows that:

$$V_T(q_T) = \max_{a \in A} E(O_T(q_T, a)) \quad (16)$$

$$V_{t-1}(q_{t-1}) = \max_a E(O_{t-1}(q_{t-1}, a) + \sum_{qt} p(q_t|q_{t-1}, a)V(q_t)) \quad (17)$$

Equations 16-17 can be solved by dynamic programming with a single backward pass.

5 Experiments

We empirically evaluate and analyse the performance of FResist using synthetic and real data sets. We use predictive log likelihood on a previously unseen test set as our means of validating the model and comparing to other methods. We compare it to logistic regression and random forest as these are the methods currently being used in state of the art decision support systems for selecting a therapy for HIV patients Revell et al. [2016, 2014], Rosen-Zvi et al. [2008]. In a previous work, Wang et al. [2009] compared random forest to artificial neural networks and SVM with RBF kernel and concluded that random forest outperformed the other methods. The systems in Revell et al. [2016,

[5]For optimizing the total cumulative reward indefinitely, one can define a decay factor and use value iteration to solve the optimization problem.

2014], Rosen-Zvi et al. [2008] use a diverse range of features including demographic features and have shown them to be very informative for predicting treatment outcome. However, in the comparisons we include in this paper we use only clinical data in the form of therapy and its outcome, as the focus of this paper is on the drug resistance evolution. As a part of decision support system, the method suggested in this paper can be used as an additional feature to logistic regression or random forest similar to the use of a Bayesian network in Rosen-Zvi et al. [2008]. Following Rosen-Zvi et al. [2008], the clinical features that we used for both logistic regression and random forest were: drugs of the current therapy, drugs previously taken by the patient, number of previous treatments, outcome of last treatment.

In the first experiment, we generated a family of datasets with varying outcome noise levels (p^N from equation 3) in the range $0 - 0.2$. There were 1000 patients in all the datasets, all of them with full adherence ($a_i = 1$). Each patient received two randomly selected combination therapies, each of which was a series of 10 continuous treatments, to a total of 20 treatments per patient. Each therapy was a random combination of 3 antiretroviral compounds out of a total of 5 available compounds to choose from (yielding $\binom{5}{3}$ possible therapies). We split each of these datasets to a train set with 500 patients and a test set with the remaining 500. Figure 3(a) shows predictive log likelihood of the three methods as a function of outcome noise. As expected, FResist outperforms the other methods on data generated according to the model. It is interesting to see that as we add noise to the data, the performance of all methods dropped drastically, well beyond the decrease in likelihood expected only due to introducing noisy outcomes in the test set (where random values need to be predicted). The reason for this decrease is that observing a noisy therapy failure, even if just one, is a strong indication that persistent resistance has evolved.

In the second experiment, we tested the effect of patient adherence on predicting therapy outcome. Similar to the first experiment, we generated a family of datasets with the same characteristics, but without any outcome noise ($p^N = 0$) and we varied the adherence level a_i between 1 and 0.8. Figure 3(b) shows predictive log likelihood of the three methods as a function of patient's adherence. Again, FResist outperforms the other two methods (the performance of FResist decreased compared to the previous experiment as adherence was sampled rather than set). Once again, we see a significant drop in performance for all methods as patient adherence decreases. The reason is the increased uncertainty when observing a treatment failure regarding its cause - whether it is due to drug resistance or because the patient simply did not take the prescribed drugs.

The EuResist dataset[6] is an integrated database containing clinical and demographic data of more than 65000 patients in Europe. From this dataset, we extracted all patients with at least 30 viral load measurements. From these, we selected a random subset of 1000 patients and split them to train and test sets of 500 each. The dataset that we extracted included clinical samples taken between 1997 and 2013. Similar to Rosen-Zvi et al. [2008] and others, we associated a clinical sample with a therapy if the sample was taken at least 8 weeks after the therapy started and before the therapy was stopped. We used viral load as the therapy outcome. A treatment was considered successful when the viral load was below 500 copies/mL and failure otherwise (see additional discussion below). Figure 4 compares the predictive log likelihood of FResist, logistic regression and random forest on the 500 patients from the EuResist test set. FResist outperforms the other two methods. To better understand FResist's performance we inspected the samples drawn by it from the posterior distribution. For example, estimating an average patient adherence p^A by averaging all variables $a_{p,t}$, the value learned by FResist is 0.87, which suggests that modeling patient adherence is important. To verify that, we compared the performance of FResist with a range of pre-set values for patient adherence, and we indeed see its importance as shown in Figure 6 (in this experiment, each adherence value was set to all patients, which performs worse than an individually learned value on the train set, but generalizes better to the test set in this setting where adherence is not individually fit to the patients in the test set). Another interesting observation is the importance of modeling the prior resistance (Figure 5): In our dataset (train and test combined), 473 out of 1000 of the first treatments failed.

It is important to note some limitations of the above empirical validation: The EuResist data is retrospective (in contrast to data from clinical trials): Drugs were selected by physicians who expected them to be effective for a specific patient. This means that a therapy prescribed for patient i is more likely to succeed for i than for others. However, the above methods ignore this bias and assume that apriori (before observing any treatment outcome) there is no difference between the patients with respect to therapy chance of success. We did not adjust for this bias in our study and this is a subject for future work.

Another limitation of the above analysis is the significant progress made in HIV research and treatment during the years in which the EuResist data has been col-

[6] https://www.euresist.org

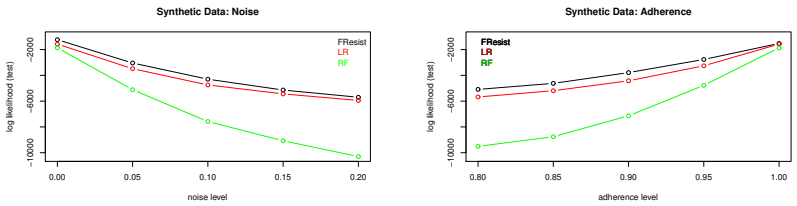

(a) (b)

Figure 3: **a.** Test log likelihood of FResist (black), logistic regression (LR, red) and random forest (RF, green) as a function of the noise added to synthetic data. **b.** Test log likelihood as a function of patient adherence

.

Figure 4: Log likelihood (test) of FResist, logistic regression and random forest on the EuResist dataset.

Figure 5: EuResist: Modeling prior drug resistance greatly outperforms assuming full initial drug sensitivity.

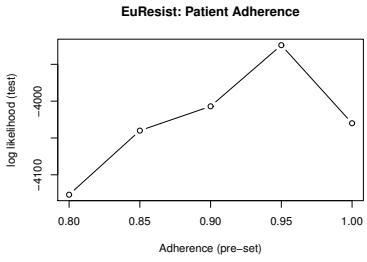

Figure 6: Patient Adherence: Optimal value is smaller than 1.

lected: availability of drugs and guidelines has changed with the introduction of new drugs and new drug classes. Moreover, the ability to detect HIV in the blood has improved significantly. While in the past viral load below 500 copies/mL could not be detected, it is now possible to detect the virus even when its level in the blood is below 50 copies/mL. To avoid an inconsistent definition of treatment success in different years, we chose 500 as our threshold for therapy success (for untreated patients, viral load is in the range of tens - hundreds of thousands).

6 Summary

In this paper we presented FResist, a novel probabilistic graphical model for the evolution of drug resistance in HIV. The model derives from Factorial HMM and specifies relations between therapy drugs, mutations, persistent resistance and therapy outcome. It further extends Factorial HMMs by modeling patient adherence. Approximate learning in FResist is done using a novel Collapsed Gibbs Sampling algorithm that we developed, new also in the context of Factorial HMM.

Compared to state of the art decision support systems for HIV therapy selection, FResist more accurately predicts therapy outcome, and gains insight regarding patient adherence to the therapy. Further, it plans a series of treatments optimizing patient's health in the long term, while existing systems provide greedy therapy recommendations optimized for the success of a single treatment at a time.

While motivated by drug resistance in HIV, it would be interesting to apply FResist to other diseases with emerging drug resistance and multi-drug therapies such as certain types of cancer. Other directions for future research include adjusting for bias in the clinical data used for training, taking possible side effects into account when planning a series of treatment or considering drug cross resistance.

References

AIDSInfo. Guidelines for the use of antiretroviral agents in hiv-1-infected adults and adolescents. *Panel on Antiretroviral Guidelines for Adults and Adolescents. Department of Health and Human Services.*, 2016. URL `http://www.aidsinfo.nih.gov/ContentFiles/AdultandAdolescentGL.pdf`.

P. Arora and N. M. Dixit. Timing the emergence of resistance to anti-hiv drugs with large genetic barriers. *PLoS Comput Biol*, 5(3):e1000305, 2009.

M. Brand, N. Oliver, and A. Pentland. Coupled hidden markov models for complex action recognition. In *Computer vision and pattern recognition, 1997. proceedings., 1997 ieee computer society conference on*, pages 994–999. IEEE, 1997.

F. Clavel and A. J. Hance. Hiv drug resistance. *New England Journal of Medicine*, 350(10):1023–1035, 2004.

W. Dong. Variational inference with agent-based models. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, pages 854–863. International Foundation for Autonomous Agents and Multiagent Systems, 2016.

W. Dong, K. Heller, and A. Pentland. Graph-coupled hmms for modeling the spread of infection. In *Uncertainty in Artificial Intelligence*, 2012.

T. El-Hay, O. Weissbrod, E. Eban, M. Zazzi, and F. Incardona. Structured proportional jump processes. In *Proceedings of the Thirtieth Conference on Uncertainty in Artificial Intelligence*, pages 172–181. AUAI Press, 2014.

K. Fan, M. Eisenberg, A. Walsh, A. Aiello, and K. Heller. Hierarchical graph-coupled hmms for heterogeneous personalized health data. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 239–248. ACM, 2015.

Z. Ghahramani, M. I. Jordan, and P. Smyth. Factorial hidden markov models. *Machine learning*, 29(2-3): 245–273, 1997.

T. R. Glass, S. De Geest, R. Weber, P. L. Vernazza, M. Rickenbach, H. Furrer, E. Bernasconi, M. Cavassini, B. Hirschel, M. Battegay, et al. Correlates of self-reported nonadherence to antiretroviral therapy in hiv-infected patients: the swiss hiv cohort study. *JAIDS Journal of Acquired Immune Deficiency Syndromes*, 41(3):385–392, 2006.

T. L. Griffiths and M. Steyvers. Finding scientific topics. *Proceedings of the National academy of Sciences*, 101 (suppl 1):5228–5235, 2004.

T. L. Griffiths, M. Steyvers, D. M. Blei, and J. B. Tenenbaum. Integrating topics and syntax. In *NIPS*, volume 4, pages 537–544, 2004.

S.-H. Kim, S. M. Gerver, S. Fidler, and H. Ward. Adherence to antiretroviral therapy in adolescents living with hiv: systematic review and meta-analysis. *Aids*, 28(13):1945–1956, 2014.

R. G. Krishnan, U. Shalit, and D. Sontag. Structured inference networks for nonlinear state space models. *arXiv preprint arXiv:1609.09869*, 2016.

J. S. Liu. The collapsed gibbs sampler in bayesian computations with applications to a gene regulation problem. *Journal of the American Statistical Association*, 89(427):958–966, 1994.

M. Ocfemia, A. Hernandez, N. Saduvala, A. Oster, and H. Hall. Epidemiology of hiv-1 transmitted drug resistance among men who have sex with men in the united states. In *XXIV International HIV Drug Resistance Workshop*, 2015.

W. Pan, W. Dong, M. Cebrian, T. Kim, J. H. Fowler, and A. S. Pentland. Modeling dynamical influence in human interaction: Using data to make better inferences about influence within social systems. *IEEE Signal Processing Magazine*, 29(2):77–86, 2012.

T. Pierson, J. McArthur, and R. F. Siliciano. Reservoirs for hiv-1: mechanisms for viral persistence in the presence of antiviral immune responses and antiretroviral therapy. *Annual review of immunology*, 18(1):665–708, 2000.

M. Prosperi, A. Pironti, F. Incardona, G. Tradigo, and M. Zazzi. Predicting human-immunodeficiency virus rebound after therapy initiation/switch using genetic, laboratory, and clinical data. In *Proceedings of the 7th ACM International Conference on Bioinformatics, Computational Biology, and Health Informatics*, pages 611–615. ACM, 2016.

M. C. Prosperi, M. Rosen-Zvi, A. Altmann, M. Zazzi, S. Di Giambenedetto, R. Kaiser, E. Schülter, D. Struck, P. Sloot, D. A. Van De Vijver, et al. Antiretroviral therapy optimisation without genotype resistance testing: a perspective on treatment history based models. *PloS one*, 5(10):e13753, 2010.

A. Revell, M. Boyd, D. Wang, S. Emery, B. Gazzard, P. Reiss, A. Sighem, J. Montaner, H. Lane, and B. Larder. A comparison of computational models with and without genotyping for prediction of response to second-line hiv therapy. *HIV medicine*, 15 (7):442–448, 2014.

A. D. Revell, D. Wang, R. Wood, C. Morrow, H. Tempelman, R. L. Hamers, P. Reiss, A. I. van Sighem, M. Nelson, J. S. Montaner, et al. An update to the hiv-treps system: the development and evaluation

of new global and local computational models to predict hiv treatment outcomes, with or without a genotype. *Journal of Antimicrobial Chemotherapy*, page dkw217, 2016.

M. Rosen-Zvi, A. Altmann, M. Prosperi, E. Aharoni, H. Neuvirth, A. Sönnerborg, E. Schülter, D. Struck, Y. Peres, F. Incardona, et al. Selecting anti-hiv therapies based on a variety of genomic and clinical factors. *Bioinformatics*, 24(13):i399–i406, 2008.

D. Van de Vijver, A. M. Wensing, C. A. Boucher, et al. The epidemiology of transmission of drug resistant hiv-1. *HIV sequence compendium*, 2007:17–36, 2006.

D. Wang, B. Larder, A. Revell, J. Montaner, R. Harrigan, F. De Wolf, J. Lange, S. Wegner, L. Ruiz, M. J. Pérez-Elías, et al. A comparison of three computational modelling methods for the prediction of virological response to combination hiv therapy. *Artificial intelligence in medicine*, 47(1):63–74, 2009.

M. Zazzi, R. Kaiser, A. Sönnerborg, D. Struck, A. Altmann, M. Prosperi, M. Rosen-Zvi, A. Petroczi, Y. Peres, E. Schülter, et al. Prediction of response to antiretroviral therapy by human experts and by the euresist data-driven expert system (the eve study). *HIV medicine*, 12(4):211–218, 2011.