Understanding of Comorbidities Using Modeling Techniques on EHR

Meenal Badki (✉ mhb180001@utdallas.edu)

The University of Texas at Dallas

Research Article

Keywords: Longitudinal data analysis, Disease dynamics, Comorbidity, Hidden Markov model, Bayesian analysis, Health informatics, Markov processes

Posted Date: November 16th, 2021

DOI: <https://doi.org/10.21203/rs.3.rs-1017334/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Understanding of comorbidities using modeling techniques on EHR

Meenal Badki

UT Dallas

mhb180001@utdallas.edu

ABSTRACT

Comorbidities refer to the existence of numerous, co-occurring diseases in medicine. The course of one comorbidity is typically extremely dependent on the course of the other condition due to their co-occurrence, and therapies can have major spill-over effects. Despite the high occurrence of comorbidities among patients, there is no complete statistical framework for modelling comorbidity longitudinal dynamics. We propose a probabilistic approach for studying comorbidity dynamics in patients over time in this paper. This approach is a non-homogenous transition technique/mechanism using Hidden Markov Model called as coupled-HMM. Clinical research influenced the design of our coupled-HMM: (1) It accounts for different disease stages (acute, stable) in disease progression by providing clinically meaningful latent phases. (2) It simulates a relationship between the trajectories of comorbidities and the dynamics of capturing co-evolution. (3) The transition mechanism takes into account between-patient heterogeneity (e.g., risk factors, treatments). Based on 675 health trajectories, we assessed our proposed Coupled-HMM, which investigates the concomitant evolution of diabetes mellitus and chronic liver disease. We find that our Coupled-HMM provides a superior fit when compared to competing models without coupling. We also assess the spill-over impact, or the amount to which diabetic therapies are linked to a shift in chronic liver disease from an acute to a stable condition. Immediate application in treatment planning and clinical research becomes possible as a result of our approach in context of comorbidities.

KEYWORDS

Longitudinal data analysis, Disease dynamics, Comorbidity, Hidden Markov model, Bayesian analysis, Health informatics, Markov processes

1 INTRODUCTION

Comorbidities are defined in medicine as the occurrence of many disorders that impact a patient's health at the same time [70]. According to recent estimates, 50 percent of hospitalizations in the United States involve individuals with eleven or more distinct disease codes [71]. Similarly, a high proportion of older people suffer from various diseases, especially chronic ones, at the same time [5, 78]. This necessitates a deeper knowledge of disease interactions, or the processes that underpin comorbidities.

Comorbidities are defined by a concurrent or simultaneous evolution, in which the progression of one disease is significantly dependent on the progression of the other [70]. Chronic pain and sadness are two examples. Both frequently coexist, and their evolution is inextricably linked. Chronic pain frequently leads to depression, and sadness, in turn, raises the perceived amount of pain [4, 27]. Diabetes and chronic liver diseases (e.g., non-alcoholic fatty liver disease) are another type of comorbidities [34, 37, 67]. The progression of both the diseases is known to be influenced by complex interactions and metabolic altercations. This is reflected in a recent suggestion to rename non-alcoholic fatty liver disease (NAFLD) to metabolic associated fatty liver disease (MAFLD) [21]. As a result, a greater understanding of the co-evolution could help with treatment planning, and is thus directly clinically relevant. However, this would necessitate models that take into account co-evolution, i.e., disease interactions throughout time are subjected to statistical modeling.

There are few statistical models that are optimized for comorbidities. A body of research has looked into which disorders are frequently encountered combined [32, 40, 73]. This inquiry, however, is about unsupervised learning (e.g., clustering) and does not look at the

diseases' longitudinal associations/interactions. Another body of research looks at the effect of comorbidities on patient health [e.g., 11, 27, 51, 59, 66] or readmission risk [71]. However, comorbidities are only employed as independent variables in these studies, not as dependent variables, or outcomes of interest. As a result, such models are unable to capture the co-evolution of comorbidities. Our research is motivated by the following question: *how does treating disease A impact the course (state) of disease B?* To solve this question, statistical inferences with interpretability are required. Diseases are typically defined in medicine by several phases of trajectory called "disease states" in which diseases progress through *acute* and *stable* states. These disease states underpin therapeutic decision-making, as treatment strategies and regimens alter depending on whether the condition is *acute* or *stable* (see the Corbin-Strauss trajectory framework [63]). The disease states, however, cannot be seen immediately [14].

Diseases are commonly classified in medicine by several phases of trajectory called "disease states" in which diseases progress through acute and stable states. These disease states underpin therapeutic decision-making, as treatment strategies and regimens alter depending on whether the condition is acute or stable (see the Corbin-Strauss trajectory framework [63]). The disease states, however, cannot be seen immediately [14]. Lab measurements, in particular, only give noisy realizations of the disease state, and so both measurements (e.g., glucose levels or even other types of symptoms) and disease states are only related stochastically. To reflect this, a frequent technique is to model health trajectories using latent states to formalize disease states and thereby depict hidden Markov models (HMMs) [e.g., 6, 50, 54, 60, 62, 65]. We follow the so-called trajectory framework [14, 63] and characterize disease evolution using disease states, but we extend the HMM-based framework from a single-disease to a multi-disease context in this paper.

Data and data extraction: In this manuscript, we utilized MIMIC-III database which was a large US-based critical care database, containing several records of patients admitted in intensive care unit. Data including vital signs, medications, laboratory measurements, fluid balance, procedure codes, imaging reports, diagnostic codes, observations and notes charted by care providers, hospital length of stay, and survival data were comprehensively recorded. Structure query language was used for data extraction from the database. General information including age, sex, BMI, time-since-prediabetes were collected. The clinical outcomes/observations are blood glucose level and alanine aminotransferase levels associated with conditions that are diabetes and chronic liver disease and their interactions.

Model: In this paper, we develop a probabilistic model to analyse the longitudinal course/progression of comorbidities in patients. In particular, we propose a custom coupled Hidden Markov Model with a personalised, non-homogeneous mechanism containing transition called coupled HMM. Based on clinical practice, we have specified coupled HMM as follows: (1) We introduce clinically meaningful latent states that is accountable for distinct phases (ie, stable or acute) in the course of the disease progression by introducing latent conditions that are of clinical importance. (2) We model a coupling between the trajectories of the comorbidities in order to capture the coevolution dynamics. (3) We consider heterogeneity for between-patients (like treatments, risk factors) in the transition mechanism. Through Markov chain Monte Carlo, we also give a completely Bayesian estimate methodology (MCMC). We forecast and evaluate a spill-over impact based on our model, which measures the indirect effect of therapies due to disease co-evolution. This enables statistical judgments with high interpretability.

Findings: A large longitudinal dataset is used to test our HMM model. Specifically, we use electronic health records from 675 individuals collected over a 10 year period. The concomitant advancement of (A) diabetes and (B) chronic liver disease is documented in the electronic health records. Since it is one of the most common diseases in the world, patients at risk of developing diabetes mellitus type 2 were chosen and its presence is frequently associated with a chronic liver disease as a related/concomitant comorbidity [34, 37, 67]. We come up with three major conclusions. First, modelling disease relationships using means coupling yields a better model fit. Second, an **acute** condition in one of the two diseases is linked to a higher likelihood of transitioning to an **acute** state in the other. Third, a diabetic treatment has significant spillover effects: it is associated with a decrease in the risk of an acute disease state in diabetes as well as chronic liver disease. Overall, the findings illustrate the necessity of modelling comorbidity co-evolution empirically.

Contributions: The work mentioned in this research advances existing research on longitudinal disease dynamics modeling in following ways:

- (1) The approach mentioned in this research is the first statistical model designed particularly to capture longitudinal comorbidities' dynamics.
- (2) We build a coupled hidden Markov model with a non-homogeneous transition mechanism for longitudinal data and provide a Bayesian estimation approach by deriving the likelihood.
- (3) The spill-over impact, which assesses the indirect effect of treatments through disease interactions, is formalized and estimated. It quantifies empirically how treating disease A impacts the course (disease state) of disease B, allowing for direct interpretability. Thereby, new knowledge for clinical practice (e. g., treatment planning) and research is generated.

2 RELATED WORK

2.1 Comorbidities

The term "comorbidity" is used in medicine to describe diseases that occur together (see [70]). Comorbidities are sometimes divided into main and secondary disorders in the literature, but this classification is not consistently applicable [23, 69]. As a result, we no longer distinguish between the two diseases and simply refer to them as disease A and disease B. Both diseases could co-occur by chance, but they could also have a causal relationship or be caused by correlated or shared risk factors [69]. This indicates that diseases interact with one another. In particular, treating one condition frequently results in an improved prognosis for co-occurring diseases. In clinical practice, comorbidities are common.

Patients in hospitals, in particular, frequently have several comorbidities (e.g., every second hospitalization contains eleven or more separate disease codes [71]). Comorbidities are also of direct relevance for clinical decision-making due to their prevalence. Because the existence of comorbidities has been shown to affect health outcomes, clinical practice stresses that their treatment be given specific consideration [30].

There are ample examples of comorbidities, in clinical practice. For example, depression and chronic pain, frequently co-occur, and their severity is frequently dependent [4, 27]. Similarly, up to 40% of cancer patients experience psychological distress, such as anxiety or depression [81]. Furthermore, diabetes patients are commonly diagnosed with comorbidities [41], such as hypertension, cardiovascular disease, or chronic liver illnesses (e.g., nonalcoholic fatty liver disease) [34, 37, 67], as explored in this research.

2.2 Statistical Inferences using Comorbidity Data

Despite the importance of comorbidities in medicine, only a few research have taken their structure into account in their models. We provide a summary of key references above. While they use comorbidity data as well, their objective is different from the work described in this paper.

Statistical models have been employed in the literature to detect co-occurrence patterns among diseases [e.g., 32, 40, 73]. This allows for the identification of diseases that co-occur and should thus be considered comorbidities. These studies, on the other hand, address the topic of which comorbidities occur frequently together but do not predict their progression. Various models can be utilized for this. Comorbidities, for instance, can be shown as networks with nodes representing symptoms and edges representing potentially causal relationships [16]. The start of comorbidities is then indicated by shared or overlapping symptoms of different conditions, as determined by the network. This hypothesis has also been tested using dynamic structural equation models [9, 31] or deep diffusion processes [56]. This answers the question of whether diseases co-occur, but not how they co-occur, i.e., without elucidating longitudinal disease interactions. Comorbidities and their temporal connections have also been studied using Bayesian networks [22, 46, 47]. However, these studies provide only a limited understanding of disease progression, such as underlying disease conditions.

Other studies [e.g., 11, 27, 59, 66, 83], including readmission risk [71], employ data on comorbidities as input when generating inferences about patient outcomes [e.g., 11, 27, 59, 66, 83]. This helps researchers to account for the fact that the co-occurrence of diseases frequently leads to nonlinear relationships with patient outcomes. When a patient has a number of co-occurring diseases, for example, patient survival may be reduced. Different indices, such as the Charlson comorbidity index [13], reflect this in clinical practice. However, in the aforementioned studies, comorbidities are treated as independent variables rather than dependent variables.

2.3 Hidden Markov Models

Hidden Markov models (HMMs) are a flexible statistical tool for modelling time series that can be used in a variety of situations. HMMs represent a set of latent states that follow a Markov chain and are used to generate observations [29]. There are widespread application of HMMs in modeling human behavior like hiring decisions [45].

Previously, HMMs were used to analyze data on patient health [e.g., 2, 42, 43, 54, 60, 65]. HMMs are further altered in these works to reflect patient trajectories via follow-up specification. First, different disease states (i.e., phases) along the trajectory frequently correspond to latent states [see 14]. Second, in order to account for patient variability and other sources of unobserved heterogeneity, the transition mechanism incorporates covariates as well as random effects [2]. The aforementioned HMMs, on the other hand, describe the evolution of single-disease outcomes but not the emergence of comorbidities in multi-disease contexts.

HMMs come in a variety of forms [e.g., 64, pp. 76–77]. Coupled HMMs, for example, are a type of HMM that can be used to simulate the evolution of numerous interacting processes. The latent states of distinct but presumably linked processes can be conditionally dependent using these coupled HMMs [8]. A Cartesian product can be used to link the latent states of distinct sequences, forming the basis of a so-called Cartesian product coupled HMM. When

the underlying processes are dependent, a Cartesian product coupled HMM outperforms conventional and multivariate HMMs in a simulation study [55].

Speech recognition [53], epileptic seizure detection [15], and sleep staging [58] are only a few of the applications of coupled HMMs. They've also been used to simulate the spread of infectious diseases [19, 68]. [55] describes the development of the Cartesian product coupled HMM for modelling the combined advancement of lab measures in a single-disease environment from intensive care units. A coupled HMM was used in ecology to look at the dependency of voles from different woodland zones [61]. It was especially investigated whether the presence of one disease makes the spread of another disease more likely. A multi-chain Markov switching model, similar to a coupled HMM, was also used to see if there was any volatility spillover between financial markets [26].

Research gap: Although comorbidities are widely recognized as important in clinical practice, statistical frameworks for predicting their co-evolution across time have received little attention. Understanding the underlying disease patterns, on the other hand, could be extremely useful for therapy planning. To do this, we created a Coupled-HMM that is specifically designed to capture comorbidity dynamics. On the basis of longitudinal data, we simulate disease relationships among comorbidities in patients.

3 PROPOSED MODEL: Coupled-HMM Model

3.1 Problem Statement

The goal of this study is to simulate the co-evolution of comorbidities. We do this by analyzing a longitudinal dataset containing patient health data. We also take into account two disorders, referred to as disease A and sickness B. The evolution of both is clearly assumed to be characterized by disease interactions in this case. This encourages researchers to learn more about how the two diseases interact. We want to know if there are any potential spillover effects while treating disease A: *how does treating disease A alter the status of disease B?* For instance, how does Metformin as a diabetes treatment affect the chance of a co-occurring liver disease becoming acute? Formally, we want to see how one sickness affects the state of the other. To allow for such statistical conclusions, we design a model that ensures interpretability.

Our model specification must include additional qualities, such as (1) disease stages, (2) coupling, and (3) between-patient heterogeneity, in order to be clinically relevant.

(1) Disease states. Many diseases (especially chronic or otherwise long-term disorders) progress through several phases throughout time, with the condition being classified as "**acute**" or "**stable**" [14]. The phases are commonly referred to as "disease states" in medicine. These states can't be readily detected; instead, they have to be deduced from data [14]. To do so, we refer to previous research [e.g., 6, 50, 54, 60, 62, 65] and model disease states using latent states. To highlight their clinical significance, we later term the latent states "**acute**" and "**stable**". We build on the HMM-based framework [57] for simulating the history of diseases A and B, which is consistent with previous literature [e.g., 6, 50, 54, 60, 62, 65]. Both are intuitively represented by two different HMMs.

(2) Coupling. We expect disease interactions among comorbidities and, as a result, we explicitly assume that the states of disease A and B are interdependent. When disease A transitions from a stable to an acute condition, for example, it should raise the likelihood of disease B moving from a stable to an acute state as well. As a result, we introduce a coupling between the HMMs from disease A and B later on.

(3) Between-patient heterogeneity. The dynamics of disease are known to differ greatly between patients [e.g., 52, 75]. In medicine, risk factors (e.g., age, gender) and the existence of therapies (e.g., 36) are used to define this. As a result, our approach accounts for patient heterogeneity as follows. On the one hand, we include covariates at the patient level that represent risk factors and therapies. These are incorporated into the HMMs' transition mechanism, allowing risk factors (as well as therapies) to influence disease dynamics: this allows risk factors (as well as treatments) to enhance the likelihood of acute vs. stable disease states.¹ Unobserved heterogeneity, on the other hand, is taken into account. We follow previous research [7, 17, 20, 33, 38, 76] and account for within-treatment effects in the treatment variable (see Sec. 4.2 for details).

We'll build a model called Coupled-HMM in the following section. The aforesaid properties are accommodated by our Coupled-HMM (1)–(3). We define a spill-over effect based on the Coupled-HMM, which measures the impact of a treatment for disease A on the status of disease B. Here, we follow clinical standards that prioritize disease states over measurements or symptoms in therapy planning [14, 48, 63]. As a result, the spill-over impact is defined in terms of disease states rather than measurements or symptoms. This enables us to draw statistical conclusions about how a treatment for disease A affects the chance of disease B becoming acute vs. stable.

3.2 Model specification

The five components of our Coupled-HMM are: (i) latent states, which represent disease states; (ii) observations, which represent measurements or symptoms; (iii) an emission component, which connects states and observations; (iv) a transition mechanism; and (v) a coupling. We use it to create an HMM-based framework that is also subject to coupling to describe the evolution of two diseases, A and B. The following is a list of the components (i)–(v):

(i) Observations: For two diseases **A** and **B**, observations denote measurements (e.g., from laboratories or data on symptoms). Let $i = 1, \dots, N$ represents/denotes different patients. Further, let $t = 1, \dots, T$ denote the different time steps. Then observations of diseases **A** and **B** are given by $Y_{it}^{(A)}$ and $Y_{it}^{(B)}$, respectively.

(ii) Latent states: The different disease states are represented by the latent states (i. e., **acute** vs. **stable**). For patient i at time step t , we introduce latent states $S_{it}^{(A)} \in S^{(A)}$ and $S_{it}^{(B)} \in S^{(B)}$, where the latent states relate to disease **A** and **B**, respectively. As a result, each disease has its own latent condition and might progress to a different stage. We assume two different states "1" and "2" for each disease, i.e., $S^{(A)} = S^{(B)} = 1, 2$. State 1 is referred to as "**stable**" whereas state 2 is referred to as "**acute**." This enables for a total of $|S^{(A)} \times S^{(B)}|$ distinct combinations. Based on the Cartesian product of the two independent/separate latent states, we now establish/define a global latent state. $S_{it}^{(G)} = (S_{it}^{(A)}, S_{it}^{(B)})$ is defined in formal terms. This will be required later when we detail the coupling as part of component (v). The underlying mapping is listed in Table 1. For instance, a global state $S_{it}^{(G)} = (2, 2)$ means that both the

Table 1: Mapping of latent states

Disease states	Latent states	Global states
----------------	---------------	---------------

Disease A	Disease B	$(S_{it}^{(A)}, S_{it}^{(B)})$	$S_{it}^{(G)}$
stable	stable	(1,1)	1
stable	acute	(1,2)	2
acute	stable	(2,1)	3
acute	acute	(2,2)	4

¹ Clinical research determined that risk factors (and treatments) should be included in the transition mechanism rather than the emission mechanism. Variables in the emission only affect observable measures (e.g., pain resistance) because they moderate how acute vs. stable illness states relate to measurements; nevertheless, they have no effect on the course of disease development. Our numerical experiments later back up our decision.

diseases are in an **acute** state. As a short form, we use the notation $S_{it}^{(G)} \in \mathcal{S}^{(G)} = \{1, \dots, 4\}$.

(iii) Emission component: The probability of an observation conditional on the latent state is defined by the emission component. It defines the likelihood of noticing a measurement or symptom based on the disease state in our situation. Both observations of the two disorders, $Y_{it}^{(A)}$ and $Y_{it}^{(B)}$ are assumed to depend on the latent states $S_{it}^{(A)}$ and $S_{it}^{(B)}$, respectively. We represent a latent state s as being conditionally normally distributed via

$$Y_{it}^{(A)} \mid S_{it}^{(A)} = s \sim \mathcal{N}(\mu_s^{(A)}, \sigma_{(A)}^2), \quad (1)$$

$$Y_{it}^{(B)} \mid S_{it}^{(B)} = s \sim \mathcal{N}(\mu_s^{(B)}, \sigma_{(B)}^2)$$

with mean $\mu_s^{(A)}, \mu_s^{(B)}$ and variance $\sigma_{(A)}^2, \sigma_{(B)}^2$. The data is modelled using an individual mean for each state and condition. As a result, an acute condition caused by disease **A** may have, on average, higher observable measurements or symptoms. The variation is disease-specific, but not state-specific. We utilized a normal distribution to represent the data in the above specification since the measurements (e.g., blood glucose levels) in our research follow a normal distribution.

(iv) Transition mechanism: The probability of transitioning from one latent disease state to another is defined by the transition mechanism. We assume the latent states follow a Markov process, as is characteristic of HMMs. However, we take into account other covariates (such as risk factors and treatments), resulting in a non-homogeneous transition mechanism that is tailored to individual patient profiles.

Formally, we use a multinomial logit link function to model the probability of global states [2, 49]. As a result, the transition probability $\gamma_{it}^{(j \rightarrow k)}$ denotes the probability of patient i transitioning from global state j at time t to global state k at time $t + 1$. We specify the transition probability via

$$\gamma_{it}^{(j \rightarrow k)} = P(S_{i,t+1}^{(G)} = k \mid S_{it}^{(G)} = j)$$

$$= \frac{\exp(\eta_{jk})}{\sum_{l=1}^{|S^{(G)}|} \exp(\eta_{jl})} \quad (2)$$

where

$$\eta_{jk} = \begin{cases} \alpha_{jk} + x_{it}^T \beta_{jk}, & \text{if } j \neq k \\ 0, & \text{if } j = k \end{cases} \quad (3)$$

Additional variables are used in the above logit: the variable α_{jk} in Eq.3 corresponds to the intercept of the transition from global state j to k . All else equal, it captures how likely a certain disease state transition is. The covariates x_{it} describe the between-patient heterogeneity (e.g., risk factors, treatment). These represent an increased likelihood of acute states (above stable states) as a result of risk factors. Treatment covariates are further processed to account for within-variable effects in the treatment variable (see Sec. 4.2). To ensure identifiability, the coefficients for $j = k$ are set to zero. As a result, recurring transitions (transitions in which the same latent illness state is maintained) serve as the reference category; see Table 6.

The initial probability of latent state s is given by π_s (i. e., the probability of latent states in time step $t = 1$). The initial probabilities were estimated independently of each patient and any risk factors. This approach follows previous works [77].

(v) Coupling: We simulate the interaction between the two diseases by linking their latent states. We build on the concept of a Cartesian product connected HMM [55] in particular. We use the Cartesian product of the underlying latent states from above for this. Recall that $S_{it}^{(G)} = (S_{it}^{(A)}, S_{it}^{(B)})$, while $S_{it}^{(G)}$ is referred to as global state. Then we assume that the global states follow a Markov chain (Fig.1).

Figure 1: Illustrative scheme of coupling inside the HMM model on comorbidities

The coupling is mathematically expressed/captured as part of the transition process by introducing the global states. The reason for this is that the transition mechanism is now dependent on the global state, which means that the likelihood of $S_{it}^{(A)}$ depends not only on the prior latent state $S_{i,t-1}^{(A)}$ but also on the other disease $S_{i,t-1}^{(B)}$. When disease A and B are exchanged, the same analogy applies. As a result, the transition mechanism accommodates for probable co-movements from both disease states.

An example of the coupling among the latent states is shown below. Remember that the transition mechanism is influenced by other covariates as well as an intercept (i. e., risk factors). Now we'll look at how the transitions of one disease differ from those of other diseases. For example, given that disease **A** is in an **acute** condition, the likelihood of a shift from a **stable** to an **acute** state for disease **B** may be increased. For example, if the equivalent intercept α_{34} [for transition (2=acute,1=stable) $\rightarrow$ (2=acute,2=acute)] is significantly bigger than the intercept α_{12} [for transition (1=stable,1=stable) $\rightarrow$ (1=stable,2=acute)], a case like this could develop. This indicates that when disease A is in an acute condition, there's a higher chance that disease B will be in an acute state as well (rather than a stable state). As a result, comparing the estimated intercepts of such transitions can reveal information about the disease interactions that lie beneath. Hence, we publish the findings of our estimations for α_{jk} later because this reflects the strength of the underlying coupling, i.e., the underlying disease interaction.

3.3 Spill-Over Effect

Medication and other treatments can also be included as covariates. If such a treatment is targeted especially at one comorbidity, it will have a direct impact on the disease's state transitions. A medicine developed to treat disease **A**, for example, could prevent the shift from state (1, 1) to state (2, 1) (i. e., prevent A from becoming **acute**). However, due to their connection, the same medicine may also have an influence on disease B's state transitions. As a result, the former will be referred to as direct treatment effects, while the latter will be referred to as spill-over effects.

The following is an example of a spill-over effect. Consider the state transitions (2,2). It depicts a situation in which a therapy causes the condition to worsen. A change from an acute to a steady state (i. e., direct effect). In the ensuing time frame, disease B also transforms to a stable condition (i. e., as an additional indirect effect).

The above spill-over effect is formalized as follows. Let z be a disease-specific treatment that impacts transition probabilities via the covariates x_{it} , as shown in Equation 3. Then, during treatment z , the probability of the aforementioned state changes of patient i is given by

$$\begin{aligned}\xi_{it}(z) &= \gamma_{it}^{(4 \rightarrow 2)} \times \gamma_{i,t+1}^{(4 \rightarrow 2)} \\ &= P\left(S_{i,t+1}^{(G)} = 2 \mid S_{it}^{(G)} = 4, Z = z\right) \\ &= P\left(S_{i,t+2}^{(G)} = 1 \mid S_{it}^{(G)} = 2, Z = z\right)\end{aligned}\tag{4}$$

Analogously, the probability under no treatment denoted with z' is given by

$$\begin{aligned}\xi_{it}(z) &= \gamma_{it}^{(4 \rightarrow 2)'} \times \gamma_{i,t+1}^{(2 \rightarrow 1)'} \\ &= P\left(S_{i,t+1}^{(G)} = 2 \mid S_{it}^{(G)} = 4, Z = z'\right) \\ &= P\left(S_{i,t+2}^{(G)} = 1 \mid S_{it}^{(G)} = 2, Z = z\right)\end{aligned}\tag{5}$$

Thus, if the therapy z increases the probability $\xi_{it}(z)$, it is specifically created for condition A, but it also has a favorable effect on disease B. Hence, we refer to

$$\text{SpillOver}_{(2,2) \rightarrow (1,2) \rightarrow (1,1)} := \xi_{it}(z) - \xi_{it}(z') \quad (6)$$

as spill-over effect for $(2,2) \rightarrow (1,2) \rightarrow (1,1)$. If $\xi_{it}(z) - \xi_{it}(z') > 0$, there exists a positive spill-over effect. Without loss of generality, potential spill-over effects could also arise in different state transitions such as $(1,1) \rightarrow (2,1) \rightarrow (2,2)$ and, hence, the above formulation can be calculated for them analogously.

3.4 Model Performance

Model fit: The expected log pointwise predictive density (elpd), which is a standard performance metric for Bayesian modelling [28], is used to assess the fit of our HMM model. Using Bayesian leave-one-out cross-validation and Pareto smoothed importance sampling, the elpd can be determined quickly [72]. In our scenario, we assess model performance at the patient level for the hold-out sample (rather than at observation level). The loo package, which is part of the statistical software R, was used for this. In addition, we present a widely applicable information criterion (WAIC).

Model variants: Alternative model variants are compared to our suggested approach of coupled HMM. The following are the model variations: (1) A straightforward version in which disease **A** has only one latent state, i.e. $|S^{(A)}| = 1$. (2) A simpler variation with only one latent state for disease **B**, i.e. $|S^{(B)}| = 1$. These are simplified disease dynamics with a smaller number of possible states. It's worth noting that both model variations have the same data as our HMM model approach. As we'll see later, the proposed HMM with $|S^{(A)}| = |S^{(B)}| = 2$ is superior, thus the necessity of accommodating different disease states is implied here.

3.5 Estimation Procedure

Sampling: In the coupling HMM, all model parameters are derived via completely Bayesian estimates [28], especially using Markov chain Monte Carlo (MCMC) sampling. We utilized the probabilistic language Stan [12] for this. We used the Hamiltonian Markov chain technique and the No-U-Turn sampler to get posterior estimates based on it [39]. To acquire posterior estimates, we performed four chains with 1500 iterations each (1500 further iterations were discarded as part of a warm-up). To do so, we derived the likelihood $\mathcal{L}$ of the HMM model via

$$\begin{aligned} \mathcal{L} = & \prod_{i=1}^N P(y_{it}^{(A)}, \dots, y_{iT}^{(A)}, y_{it}^{(B)}, \dots, y_{iT}^{(B)}) = \\ & \prod_{i=1}^N \sum_{s_1=1}^{|S^{(G)}|} \sum_{s_2=1}^{|S^{(G)}|} \dots \sum_{s_T=1}^{|S^{(G)}|} \left[\pi_{s_1} \times \prod_{t=2}^T \gamma_{i,t-1}^{(s_{t-1} \rightarrow s_t)} \times \right. \\ & \left. \prod_{t=1}^T \mathcal{N}(y_{it}^{(A)}; \mu_{\phi^{(A)}(s_t)}^{(A)}, \sigma_{(A)}^2) \times \mathcal{N}(y_{it}^{(B)}; \mu_{\phi^{(B)}(s_t)}^{(B)}, \sigma_{(B)}^2) \right], \end{aligned} \quad (7)$$

where $\phi^{(A)}(s_t)$ and $\phi^{(B)}(s_t)$ map/represent the global state and corresponding individual states of each disease (see Table 1). In order to develop an efficient computation scheme, we

used the forward algorithm [84, pp. 36–39]. Furthermore, to avoid label switching and improve identifiability, the state-dependent means were arranged, i.e., $\mu_1^{(A)} < \mu_2^{(A)}, \dots, \mu_1^{(B)} < \mu_2^{(B)}$. On normal hardware, the runtime was roughly 12 hours.

Priors: For the emission component and the initial state probabilities, implying the necessity of accommodating different disease we employed weakly informative priors:

$$\mu_s^{(A)}, \mu_s^{(B)} \sim \mathcal{N}(0, 10^2) \quad (8)$$

$$\sigma_{(A)}, \sigma_{(B)} \sim \mathcal{N}(0, 1) \quad (9)$$

$$\pi_s \sim \text{Dirichlet}(1, 1, 1, 1) \quad (10)$$

The intercepts α_{jk} were chosen with slightly narrow priors in order to stabilise the model fit. Priors were set on the transformed parameters $\tilde{\beta}_{jk}$, which emerged from a thinned QR decomposition of the centred covariates x_{it} , rather than on the co-efficients β_{jk} .

$$\alpha_{jk} \sim \mathcal{N}(0, 2.5^2), \quad (11)$$

$$\alpha_{jk} \sim \mathcal{N}(0, 1). \quad (12)$$

Diagnostics: We examine the convergence of the Markov chains using typical Bayesian modelling guidelines [28]. We also run posterior predictive checks to see if we can simulate plausible new observations using the model parameters we've estimated [25]. Both are reported in Sec. B as part of the model diagnostics. We also tested if our Coupled-HMM could recover the needed values from completely simulated data. We later present the posterior mean, as well as the 50% and 90% credible intervals, unless otherwise noted (CrI).

4 EMPIRICAL SETUP

4.1 Data

We use two co-occurring diseases to assess our HMM approach: (A) diabetes mellitus type 2 and (B) chronic liver disease. Both are frequently seen as comorbidities [34, 37, 67]. Diabetes is one of the most common chronic diseases, affecting millions of individuals throughout the world. Diabetes also costs a lot of money when it comes to providing care [79]. Chronic liver disease, on the other hand, is a medical disorder that frequently occurs in conjunction with diabetes [34, 37, 67]. We used blood glucose levels as lab measurements for diabetes and the level of alanine aminotransferase, an enzyme found primarily in the liver [24, 74], as lab measurements for chronic liver disease in our study.

Following the treatment of elevated blood glucose levels with Metformin, we compute spill-over effects. Metformin is a first-line therapy for type 2 diabetes that decreases blood glucose levels overall. Metformin is also thought to be useful for patients with liver disease in the literature [3, 10, 34]. As a result, the question of how much the spill-over impact affects the disease state of chronic liver disease arises.

We used a longitudinal dataset of patients with prediabetes in our study. The dataset is made up of annual electronic health records [1, 82] that contain lab measurements, therapies, and risk factors, among other things (see next section). There are 675 patients in total in the dataset. Each patient has a time series of recordings spanning four to ten years. There are 3253 observations in total.

4.2 Model Variables

The coupled HMM is estimated using the following variables (Table 2).

The observations $Y_{it}^{(A)}$ and $Y_{it}^{(B)}$ represent the glucose level and the alanine aminotransferase level, respectively, as measurements of diabetes and chronic liver disease. Due to their right-skewed distributions, we log-transform both variables.

Between-patient heterogeneity is defined by the covariates x_{it} . These are risk factors and treatment information. The body mass index (BMI), sex, age (at time step $t=1$), and time since prediabetes are all used in this study. We explicitly split temporal data into age and time-since-prediabetes in order to isolate age as a risk factor and possible trends. Metformin medication provides information on treatments. This is a binary variable that indicates whether Metformin (Glucomin or Glucophage) was prescribed.

Table 2: Overview of model variables

Model component	Variables
Observations	
Diabetes $Y_{it}^{(A)}$	Blood glucose level (log)
Chronic liver disease $Y_{it}^{(B)}$	Alanine aminotransferase level (log)
Transition mechanism	
Covariates x_{it}	BMI, sex, age, time-since-prediabetes, Metformin treatment (centered), Metformin treatment (lag=1; centered)

To determine Metformin's so-called *within* effect, the following approach was used. Metformin is generally expected to reduce or stabilise blood glucose levels in patients [3]. Patients who use Metformin, on the other hand, may have greater glucose levels than those who do not. As a result, the so-called within and between effects may be at odds. As a result, in order to infer about the inside impact, the Metformin variable is centred with the relevant patient specific mean. This strategy is in line with earlier research [7, 17, 20, 33, 38, 76]. A lagged version of the Metformin variable was included in addition to the centred version.

4.3 Summary Statistics

Table 7 contains summary statistics for the model variables. Blood glucose levels are higher than recommended by the World Health Organization in this population of people with (pre-) diabetes, as is to be expected. Similarly, population of individuals had alanine aminotransferase readings as high as 143.40 u/l, indicating the existence of liver disease [e.g., 34]. Overall, patients have a relatively high BMI, indicating a significant prevalence of obesity. This is to be expected, as a high BMI is a known diabetes risk factor. Metformin was only given to a small percentage of the patients.

Both measurements, notably blood glucose and alanine aminotransferase levels, show a substantial association. The Pearson's correlation coefficient is 0.11 ($p < 0.001$). As a result, a higher blood glucose level coincides to a larger amount of alanine aminotransferase. This suggests that the two diseases may interact, but it does not allow for statistical inferences about the underlying longitudinal dynamics.

5 NUMERICAL RESULTS

5.1 Model Performance

Table 3 contains the findings of the performance comparison. On the basis of the expected pointwise predictive density, the models are evaluated (elpd). In a Bayesian context, the elpd evaluates predictive accuracy and, as a result, the overall fit of a model [28, 72]. A higher elpd suggests that the fit is better. The proposed coupled HMM for comorbidities has the greatest elpd (2,148.78) out of all the models that converged. As a result, it obtains the best overall model fit.

We employ baseline models that can be used to model numerous time series at the same time. We use the baselines from [55] in particular. Alternative models are usually developed for single time series or don't have a method to account for illness interactions (see Sec. 2). As a result, their selection for the purpose of our paper is ruled out.

What should the number of latent states be? We compare our HMM to baselines with a simplified latent state space, that is, a coupled-HMM with disease **A** or **B** latent states are removed. Here, simplified A refers to a diabetic variant with a single latent state, while simplified B refers to a chronic liver disease variant with a single latent state. As a result, the diseases are believed to have only one state (despite the fact that this contradicts the Corbin-Straus trajectory paradigm [14, 63]). The suggested coupled-HMM outperforms both the simplified A (elpd: 1,620.96) and simple B (elpd: 1,450.83) by a significant margin.

Taken together, this emphasizes the significance of accounting for different trajectory phases by integrating latent states (i. e., disease states). As a result, we give estimation results from the suggested coupled HMM with 22 states in the following sections, as this model is favored in the model comparison.

How should patient heterogeneity be represented? We run more comparisons to find an answer. To begin, we use a naive linked HMM [8]. For modelling between-patient heterogeneity, this model does not employ any covariates (i.e., neither in the emission component nor in the transition mechanism). Clearly, the model is inadequate. Second, we employ Pohle et al emission's coupled HMM [55]. In the emission, this model incorporates variables (whereas our model includes covariates in the transition). However, during MCMC sampling, the chains did not converge, preventing interpretability of the model coefficients. Moreover, the estimate of the elpd may be unreliable. Third, as a remedy, we tweaked Pohle et al.'s emission coupled HMM [55]. In order to avoid label-switching across latent states, we introduce an ordering of the estimated means in addition to ordered intercepts, set initial values for intercepts based on a k-means procedure [18], and follow common practice in Bayesian modelling by subjecting the covariates to a QR decomposition. The latter should remove posterior correlations among covariates, allowing the MCMC chains to converge more quickly [28]. To validate the accuracy of our implementation, we tested simulated data through the improved emission coupled HMM, where MCMC sampling converged and the model coefficients were correctly recovered. For 12 states, the model converged effectively (but not for the other state combinations). In comparison to our proposed model, the model exhibits a lower match overall. Overall, this suggests that variables should be considered in the transition mechanism (rather than the emission component).

What is the value of disease interaction modelling? Coupled HMMs with only one latent state cannot account for disease interactions by definition. As a result, we must now evaluate the associated benefit. We find that the model with 2 X 2 states works best for the naive coupled HMM, implying that capturing disease interactions is helpful. (Due to chain

divergence, we are unable to remark on the emission coupled HMM.) We also discover that the model with 2 X 2 states is the best for the suggested coupled- HMM. As a result, both diabetes and chronic liver disease must be anticipated to have longitudinal illness interactions.

5.2 Estimation Results

Emission component: Table 4 reports the estimated parameters from the emission component. The stable condition of diabetes is defined by an average blood glucose level of only $\exp(4.55) = 94.63$ mg/dl. The **acute** condition, as expected, has a larger value, with an average of $\exp(4.70) = 109.95$ mg/dl. The credible intervals for both do not overlap, indicating that they are statistically significant.

The mean of the stable state in chronic liver disease is $\exp(2.86) = 17.46$ u/l. This is lower than the acute state's average of 30.88 u/l. Because the credible intervals do not overlap, the posterior means of the stable and acute states are statistically significant level.

Transition mechanism: The following are parameter estimates for state transition probability. We refrain from presenting all of the coefficients due to high number, instead we focus on how Metformin treatment affects the transition likelihood. As a result, the direct treatment effect of Metformin on the disease state can be quantified. We anticipate the following, based on clinical research. (1) Metformin is a type of diabetes medication. As a result, it should solely affect diabetes and not chronic liver disease. This is confirmed (no zero in any of the related CrI). (2) Because metformin is only effective for a brief period of time, we anticipate to detect treatment effects in the non-lagged variable. However, again the approach yields positive evidence. (3) We anticipate a positive treatment impact in the transition (**acute**, **acute**) phase (**stable**, **acute**). A considerable fraction of the probability mass is above zero in this case, indicating that Metformin is having a beneficial treatment impact. In this case, the posterior mean equates to a 3.29 log-odds increase, resulting in an $\exp(3.29) = 26.84$ odds ratio for a one unit rise in the centred Metformin variable, making the aforementioned transition more likely than remaining in the state (**acute**, **acute**).

Coupling: We now look into the relationship between the two comorbidities. For this, we present the transition mechanism's baseline probability (i.e., the intercept α_{jk}). It calculates the transformation based on both illness **A** and **B**'s latent stages. All covariates were set to their respective sample mean to estimate this. In Eq. 3, this equates to $x_{it} = \bar{x}_{it}$.

The diabetic coupling is shown in Figure 2. The plot depicts the likelihood of disease A (diabetes) transitioning from a stable to an acute latent state, based on the latent states of both diseases. If both diseases are unrelated, the credible intervals for the chance of a conditional transition should encompass zero. Stronger connection is indicated by higher conditional transition probability values. We can see from Fig. 2 that an acute condition for chronic liver disease is linked to a higher risk of switching from stable to acute diabetes; see $(1, 2) \rightarrow (2, 2)$. Simply put, if chronic liver disease is acute, diabetes is likely to go from stable to acute as well. Overall, there is significant evidence that diabetes and chronic liver disease are linked.

When exchanging **A** and **B**, the understanding is analogous. Similar effects may be seen in Fig. 3 as well. Simply put, if diabetes is already **acute**, chronic liver disease is likely to become **acute** as well. However, the effect is less apparent, and the credible intervals of 50% and 90% overlap to a large extent. Nonetheless, these findings show that the two comorbidities do actually interact.

5.3 Estimated Spill-Over Effect

Lastly spill-over consequences are discussed. Through the co-evolution of the two comorbidities, we specifically analyze the influence of Metformin on liver disease. The state transitions $(2, 2) \rightarrow (1, 2) \rightarrow (1, 1)$. It reflects the progression of diabetes from **acute** to stable, followed by the progression of liver disease from **acute** to **stable**. We specifically mention the following: (1) Metformin, i.e. it (z), has been prescribed. (2) The centered variable is set to 0.5 in this case, and the remaining variables are considered to be equal to the sample mean as well. (3) The spillover impact, which is equivalent to the absolute difference between the two. (4) There is a quotient. When comparing Metformin vs. no Metformin, it indicates how many times more frequent a spill-effect is.

Note: The probability of diabetes moving from a stable to an acute latent state are shown here, i.e. $(1, \square)$ and $(2, \square)$. As a result, the transition probability between stable chronic liver disease (top) and acute chronic liver disease (bottom) is compared (bottom). The posterior means (thick line) and the 50% credible interval (shaded region) are shown in the estimates. During estimation, all covariates are assigned to the sample mean.

Figure 2: Estimation results for coupling conditional on diabetes.

Table 3: Model performance comparison

	Disease interactions	#Latent states	elpd	SE(elpd)	WAIC
Naïve coupled HMM [8]	×	(1×2)	1,622.54	150.1	-3,245.17
Naïve coupled HMM [8]	×	(2×1)	1,430.12	140.83	-2,860.39

Naïve coupled HMM [8]	✓	(2×2)	2,133.57	148.64	-4,267.34
Emission coupled HMM (unmodified) [55]	✗	(1×2)	(1775.55)†	(152.19)†	(-3491.69)†
Emission coupled HMM (unmodified) [55]	✗	(2×1)	(1462.99)†	(148.61)†	(-2483.17)†
Emission coupled HMM (unmodified) [55]	✓	(2×2)	(2280.02)†	(153.16)†	(-4509.16)†
Emission coupled HMM (modified) [55]	✗	(1×2)	1,932.06	148.19	-3,864.83
Emission coupled HMM (modified) [55]	✗	(2×1)	(1383.15)†	(142.6)†	(-2528.34)†
Emission coupled HMM (modified) [55]	✓	(2×2)	(1424.63)†	(150.14)†	(-2563.72)†
Coupled-HMM (simplified A)	✗	(1×2)	1,620.96	149.61	-3,241.99
Coupled -HMM (simplified B)	✗	(2×1)	1,450.83	139.70	-2,902.97
Coupled-HMM (proposed)	✓	(2×2)	2,148.78	147.29	-4,298.73

† MCMC sampling did not converge

elpd: expected log pointwise predictive density (higher values are better); SE: standard error

WAIC: widely-acceptable information criterion (lower values are better)

Table 4: Posterior estimates of emission component

	Latent state	Mean	Credible Interval	
			5 %	95 %
Diabetes (blood glucose level; in log)				
$\mu_1^{(A)}$	stable	4.55	4.54	4.55
$\mu_2^{(A)}$	acute	4.70	4.69	4.70
$\sigma_{(A)}$	----	0.09	0.09	0.09
Chronic liver disease (alanine aminotransferase level; in log)				
$\mu_1^{(A)}$	stable	2.86	2.85	2.88
$\mu_2^{(A)}$	acute	3.43	3.41	3.45
$\sigma_{(A)}$	-----	0.30	0.29	0.31

Table 5 reports the corresponding quantiles of the transition probabilities of these two scenarios, as well as their difference.

In general, we discover the following. If no Metformin is provided, the situation of liver illness through the underlying coupling seldom changes. All estimated quantiles are close to 0 in this case. When Metformin is administered, there is a non-zero chance that the state of

chronic liver disease will improve as well. We can see that the distribution for it (z), i.e., Metformin prescribed, has a very large tail by comparing the different quantiles. The spill-over impact is minor or absent for many patients on the left-tail (e.g., due to non-adherence or simply because the treatment is ineffective). Based on our findings, it cannot be ruled out that Metformin treatment has unfavorable consequences for a co-occurring liver condition in the 5% quantile.

However, there is a considerable estimated spill-over effect for patients on the right-tail. For the 50 percent quantile, prescribing Metformin (as opposed to not prescribing) increases the likelihood of a better liver state by 2.898 times. The relative quotient even rises to 13.695 at the 75th percentile. Our methodology, as illustrated below, allows us to estimate the impact of diabetes on liver damage caused by Metformin use. In conclusion, we find evidence that Metformin stabilizes not only diabetes but, through the coupling, also a co-occurring chronic liver disease.

Note: Here shown are the transitions probabilities that chronic liver disease moves from a stable to an acute latent state, i. e., $(\square, 1) \rightarrow (\square, 2)$. Hence, this compares the transition probability when exiting a stable diabetes (top) and an acute diabetes (bottom). Estimations show the posterior means (thick line), as well as the 50 % (shaded area) credible interval. All covariates are set to the respective sample mean during estimation.

Figure 3: Estimation results for coupling conditional on chronic liver disease

6 DISCUSSION

Summary of findings: (1) According to our findings, our Coupled-HMM beats other model specifications. It is especially superior to variations in which no illness interaction is modelled. We get significant gains in model performance as assessed by Bayesian modelling predictive accuracy (i.e., the anticipated pointwise predictive density, elpd). (2) We present evidence that both diabetes and chronic liver disease are interconnected. This is revealed by the fact that transitions from one disease to the next are statistically significant dependent on the state of the other disease. Thereby, we confirm the importance of coupling in our model formation. (3) Our model returns empirical evidence characterizing the spill-over effect of a diabetes treatment on the state from chronic liver disease.

Table 5: Estimated spill-over effects

	5%	25%	50%	75%	95%
Metformin prescribed ($\xi it(z)$)	0.000 %	0.002 %	0.017 %	0.089 %	0.627 %
No Metformin prescribed ($\xi it(z')$)	0.000 %	0.002 %	0.006 %	0.016 %	0.044 %
Difference	-0.019 %	-0.002 %	0.008%	0.077%	0.619%
Quotient	0.022	0.463	2.898	13.695	91.276

Note: The quantiles of the probabilities of the transition $(2,2) \rightarrow (1,2) \rightarrow (1,1)$ are reported. Metformin prescribed assumes that the centered Metformin variables is equal to 0.5 and the remaining covariates are equal to the sample mean. In contrast, no Metformin prescribed corresponds to a centred Metformin variable equal to 0 and the remaining covariates also set to the sample mean. Additionally, the quantiles of the difference of these transition probabilities are given.

The role of Metformin in patients with chronic liver disease has been controversial in the past [3, 10, 34], especially because a direct effect is not always supported by the molecular mechanism of action. We present empirical evidence that Metformin has a direct influence on the condition of diabetes, but no evidence that it has a direct effect on the state of chronic liver disease. Instead, we find an indirect effect through the coupling.

Interpretability: In order to warrant interpretability, our model is based on a parsimonious specification. We draw on a central concept in medicine: diseases form a so-called "trajectory" that goes through several phases, including **acute** and **stable** disease states [14, 63]. These disease states are crucial in clinical decision-making because different treatment approaches should be used based on whether a patient is in an **acute** or **stable** state. It suggests, in particular, that clinical decision-making should adjust therapy to the underlying illness state (rather than measurements or symptoms, which are just noisy observations of the state) [48]. As a result, previous research has used latent states to simulate health trajectories [e.g., 6, 50, 54, 60, 62, 65]. In keeping with this, we create latent illness states and utilize them to link the states of various diseases, as well as characterize our spill-over effect in terms of latent states (rather than measures or symptoms).

Limitations: Our research, like any medical research, has limitations, and we propose recommendations for future research. To begin, we have a one-of-a-kind, large-scale collection of electronic health records from a Western health provider. Our approach could be replicated in future studies with different cohorts. Second, we used a large number of risk factors based on earlier research. Other variables, on the other hand, could reflect risk factors, and future study could examine how the model's performance changes when they are included. Third, our research is centered on a specific scenario, namely the interaction between diabetes and chronic liver disease while taking Metformin. Future study could confirm our hypothesis by looking at other comorbidities. By increasing the underlying latent state space, the model can be expanded to comorbidities involving three (or more) diseases if desired.

Implications: Our model offers insights that can be used in clinical treatment and research. Spill-over effects can be used to assess to what extent treatments for disease **A** also affect/change the course (disease state) of disease **B**. As a result, when attempting to stabilize the course of disease **B**, clinical practitioners can choose a treatment that has a direct effect on **B** or, alternatively, a treatment plan for **A** that has an indirect influence on **B**, resulting in disease **B** being stabilized. Our approach provides statistical insights in this area, allowing for

evidence-based treatment planning decisions.

Diabetes and chronic liver disease were used to test our approach. However, in order to ensure broad applicability, our model is constructed in a general manner. Our clinical co-author envisions a variety of applications: Our approach could be used to investigate comorbidities that are poorly understood or have contradictory findings in the clinical literature. Given the vast number of comorbidities, this gives rich opportunities for follow-up research.

7 CONCLUSION

Patients with comorbidities, which are defined as the presence of numerous co-occurring diseases, are common, but full statistical frameworks for modelling the longitudinal dynamics of comorbidities are uncommon. Such models would provide new insights into how comorbidities co-evolve throughout time. This, in turn, would enable us to answer the question: *how does treating disease A affect the status of disease B?*

We developed a probabilistic model for longitudinal data analysis to solve this: a linked hidden Markov model with individualized, non-homogeneous transitions (coupled-HMM).

Our coupled-HMM is the first statistical model that is especially optimized to capture longitudinal dynamics among comorbidities, to the best of our knowledge. Our methodology enhances interpretability even further by providing clinically relevant insights. We presented a spill-over effect, which evaluates a treatment's indirect effect on disease states via the co-development of comorbidities. A longitudinal dataset of 675 patients was used to test our model.

8 DECLARATIONS

Ethics approval and consent to participate Not applicable, as the data is openly accessible in the form of MIMIC-III.

Consent for publication Not applicable

Availability of data and material All data generated or analysed during this study are included in this manuscript. The data consists of electronic health records from Medical Information Mart for Intensive Care (MIMIC III) database which is openly accessible. [84]

Competing interests The authors declare that they have no competing interests.

Authors' contributions Conceptualization, Data curation, Investigation, Methodology, Writing—draft, review and editing : M. BADKI

REFERENCES

- [1] Ardavan Afshar, Ioakeim Perros, Haesun Park, Christopher deFilippi, Xiaowei Yan, Walter Stewart, Joyce Ho, and Jimeng Sun. 2020. TASTE: Temporal and Static Tensor Factorization for Phenotyping Electronic Health Records. In Proceedings of the ACM Conference on Health, Inference, and Learning (Toronto, Ontario, Canada) (CHIL '20). Association for Computing Machinery, New York, NY, USA, 193–203.
- [4] Matthew J Bair, Rebecca L Robinson, Wayne Katon, and Kurt Kroenke. 2003. Depression and Pain Comorbidity: A Literature Review. *Archives of Internal Medicine* 163, 20 (2003), 2433–2445.
- [2] Rachel MacKay Altman. 2007. Mixed Hidden Markov Models. *J. Amer. Statist. Assoc.* 102, 477 (2007), 201–210.

- [5] Karen Barnett, Stewart W Mercer, Michael Norbury, Graham Watt, Sally Wyke, and Bruce Guthrie. 2012. Epidemiology of multi morbidity and implications for health care, research, and medical education: a cross-sectional study. *The Lancet* 380, 9836 (2012), 37–43.
- [3] Clifford J. Bailey and Robert C. Turner. 1996. Metformin. *New England Journal of Medicine* 334, 9 (1996), 574–579. PMID: 8569826.
- [7] Andrew Bell and Kelvyn Jones. 2015. Explaining Fixed Effects: Random Effects Modeling of Time-Series Cross-Sectional and Panel Data. *Political Science Research and Methods* 3, 1 (2015), 133–153.
- [6] Nicola Bartolomeo, Paolo Trerotoli, and Gabriella Serio. 2011. Progression of liver cirrhosis to HCC: an application of hidden Markov model. *BMC medical research methodology* 11, 1 (2011), 38–38. Modeling Longitudinal Dynamics of Comorbidities
- [9] Laura F. Bringmann, Nathalie Vissers, Marieke Wichers, Nicole Geschwind, Peter Kuppens, Frenk Peeters, Denny Borsboom, and Francis Tuerlinckx. 2013. A Network Approach to Psychopathology: New Insights into Clinical Longitudinal Data. *PLOS ONE* 8, 4 (04 2013), 1–13.
- [8] M. Brand, N. Oliver, and A. Pentland. 1997. Coupled hidden Markov models for complex action recognition. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*. 994–999.
- [10] Elisabetta Bugianesi, Elena Gentilcore, Rita Manini, Stefania Natale, Ester Vanni, Nicola Villanova, Ezio David, Mario Rizzetto, and Giulio Marchesini. 2005. A Randomized Controlled Trial of Metformin versus Vitamin E or Prescriptive Diet in Nonalcoholic Fatty Liver Disease. *The American Journal of Gastroenterology* 100, 5 (2005), 1082.
- [12] Bob Carpenter, Andrew Gelman, Matthew Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell. 2017. Stan: A Probabilistic Programming Language. *Journal of Statistical Software, Articles* 76, 1 (2017), 1–32.
- [11] Mercedes R. Carnethon, Mary L. Biggs, Joshua I. Barzilay, Nicholas L. Smith, Viola Vaccarino, Alain G. Bertoni, Alice Arnold, and David Siscovick. 2007. Longitudinal Association Between Depressive Symptoms and Incident Type 2 Diabetes Mellitus in Older Adults: The Cardiovascular Health Study. *Archives of Internal Medicine* 167, 8 (04 2007), 802–807.
- [15] Jeff Craley, Emily Johnson, and Archana Venkataraman. 2018. A Novel Method for Epileptic Seizure Detection Using Coupled Hidden Markov Models. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2018*, Alejandro F. Frangi, Julia A. Schnabel, Christos Davatzikos, Carlos Alberola-López, and Gabor Fichtinger (Eds.). Springer International Publishing, Cham, 482–489.
- [13] Mary E. Charlson, Peter Pompei, Kathy L. Ales, and C. Ronald MacKenzie. 1987. A new method of classifying prognostic comorbidity in longitudinal studies: Development and validation. *Journal of Chronic Diseases* 40, 5 (1987), 373–383.
- [16] Angélique O. J. Cramer, Lourens J. Waldorp, Han L. J. van der Maas, and Denny Borsboom. 2010. Comorbidity: A network perspective. *Behavioral and Brain Sciences* 33, 2-3 (2010), 137–150.
- [14] J M Corbin and A Strauss. 1991. A nursing model for chronic illness management based upon the Trajectory Framework. *Scholarly inquiry for nursing practice* 5, 3 (1991), 155–174.
- [17] Patrick J. Curran and Daniel J. Bauer. 2011. The Disaggregation of Within-Person and Between-Person Effects in Longitudinal Models of Change. *Annual Review of Psychology* 62, 1 (2011), 583–619. PMID: 19575624.
- [19] Wen Dong, Alex Sandy Pentland, and Katherine A. Heller. 2012. Graph-Coupled HMMs for Modeling the Spread of Infection. In *Proceedings of the Twenty-Eighth*

Conference on Uncertainty in Artificial Intelligence (Catalina Island, CA) (UAI'12). AUAI Press, Arlington, Virginia, USA, 227–236.

[18] Luis Damiano, Brian Peterson, and Michael Weylandt. 2018. A Tutorial On Hidden Markov Models Using Stan.

[20] Craig K. Enders and Davood Tofighi. 2007. Centering Predictor Variables in Cross Sectional Multilevel Models: A New Look at an Old Issue. *Psychological Methods* 12, 2 (2007), 121–138.

[22] Syed Hasib Akhter Faruqui, Adel Alaeddini, Carlos A. Jaramillo, Jennifer S. Potter, and Mary Jo Pugh. 2018. Mining patterns of comorbidity evolution in patients with multiple chronic conditions using unsupervised multi-level temporal Bayesian network. *PLOS ONE* 13, 7 (07 2018), 1–22.

[21] Mohammed Eslam, Arun J. Sanyal, Jacob George, Arun Sanyal, Brent Neuschwander-Tetri, Claudio Tiribelli, David E. Kleiner, Elizabeth Brunt, Elisa-betta Bugianesi, Hannele Yki-Järvinen, Henning Grønbaek, Helena Cortez-Pinto, Jacob George, Jiangao Fan, Luca Valenti, Manal Abdelmalek, Manuel Romero- Gomez, Mary Rinella, Marco Arrese, Mohammed Eslam, Pierre Bedossa, Philip N. Newsome, Quentin M. Anstee, Rajiv Jalan, Ramon Bataller, Rohit Loomba, Silvia Sookoian, Shiv K. Sarin, Stephen Harrison, Takumi Kawaguchi, Vincent Wai- Sun Wong, Vlad Ratziu, Yusuf Yilmaz, and Zobair Younossi. 2020. MAFLD: A Consensus-Driven Proposed Nomenclature for Metabolic Associated Fatty Liver Disease. *Gastroenterology* 158, 7 (2020), 1999 – 2014.e1. Nonalcoholic Fatty Liver Disease in 2020.

[23] Alvan R. Feinstein. 1970. The pre-therapeutic classification of co-morbidity in chronic disease. *Journal of Chronic Diseases* 23, 7 (1970), 455–468.

[25] Jonah Gabry, Daniel Simpson, Aki Vehtari, Michael Betancourt, and Andrew Gelman. 2019. Visualization in Bayesian workflow. *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 182, 2 (2019), 389–402.

[24] Abigail Fraser, Ross Harris, Naveed Sattar, Shah Ebrahim, George Davey Smith, and Debbie A. Lawlor. 2009. Alanine Aminotransferase, γ -Glutamyltransferase, and Incident Diabetes. *Diabetes Care* 32, 4 (2009), 741–750.

[26] Giampiero M. Gallo and Edoardo Otranto. 2008. Volatility spillovers, interdependence and comovements: A Markov Switching approach. *Computational Statistics & Data Analysis* 52, 6 (2008), 3011–3026.

[28] Andrew Gelman, John B Carlin, Hal S Stern, David B Dunson, Aki Vehtari, and Donald B Rubin. 2013. *Bayesian data analysis*. CRC press.

[27] Sandra W. Geerlings, Jos W. R. Twisk, Aartjan T. F. Beekman, Dorly J. H. Deeg, and Willem van Tilburg. 2002. Longitudinal relationship between pain and depression in older adults: sex, age and physical disability. *Social Psychiatry and Psychiatric Epidemiology* 37, 1 (2002), 23–30.

[31] Robin N. Groen, Oisín Ryan, Johanna T. W. Wigman, Harriëtte Riese, Brenda W. J. H. Penninx, Erik J. Giltay, Marieke Wichers, and Catharina A. Hartman. 2020. Comorbidity between depression and anxiety: assessing the role of bridge mental states in dynamic psychological networks. *BMC Medicine* 18, 1 (2020), 308.

[29] Zoubin Ghahramani. 2001. An Introduction to Hidden Markov Models and Bayesian Networks. *International Journal of Pattern Recognition and Artificial Intelligence* 15, 1 (2001), 9–42.

[32] Marina Guisado-Clavero, Albert Roso-Llorach, Tomàs López-Jimenez, Mariona Pons-Vigués, Quintí Foguet-Boreu, Miguel Angel Muñoz, and Concepción Violán. 2018. Multimorbidity patterns in the elderly: a prospective cohort study with cluster analysis. *BMC Geriatrics* 18, 1 (2018), 16.

[30] Ronald Gijsen, Nancy Hoeymans, François G. Schellevis, Dirk Ruwaard, William A.

- Satariano, and Geertrudis A.M. van den Bos. 2001. Causes and consequences of comorbidity: A review. *Journal of Clinical Epidemiology* 54, 7 (2001), 661–674.
- [33] Ellen L. Hamaker and Raoul P. P. Grasman. 2015. To center or not to center? Investigating inertia with a multilevel autoregressive model. *Frontiers in Psychology* 5 (2015), 1492.
- [34] Elizabeth H. Harris. 2005. Elevated Liver Function Tests in Type 2 Diabetes. *Clinical Diabetes* 23, 3 (2005), 115–119.
- [38] Lesa Hoffman and Robert S. Stawski. 2009. Persons as Contexts: Evaluating Between-Person and Within-Person Effects in Longitudinal Analysis. *Research in Human Development* 6, 2-3 (2009), 97–120.
- [36] Tobias Hatt and Stefan Feuerriegel. 2021. Estimating Average Treatment Effects via Orthogonal Regularization. *arXiv preprint arXiv:2101.08490* (2021).
- [44] A. Jasra, C. C. Holmes, and D. A. Stephens. 2005. Markov Chain Monte Carlo Methods and the Label Switching Problem in Bayesian Mixture Modelling. *Statist. Sci.* 20, 1 (2005), 50–67.
- [39] Matthew D. Homan and Andrew Gelman. 2014. The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research (JMLR)* 15, 1 (Jan. 2014), 1593–1623.
- [37] Ingrid J. Hickman and Graeme A. Macdonald. 2007. Impact of Diabetes on the Severity of Liver Disease. *The American Journal of Medicine* 120, 10 (2007), 829–834.
- [42] Christopher H. Jackson and Linda D. Sharples. 2002. Hidden Markov models for the onset and progression of bronchiolitis obliterans syndrome in lung transplant recipients. *Statistics in Medicine* 21, 1 (2002), 113–128.
- [40] Chien-Fang Huang, Ju-Chi Liu, Hui-Chuan Huang, Shao-Yuan Chuang, Chang-I Chen, and Kuan-Chia Lin. 2017. Longitudinal transition trajectory of gouty arthritis and its comorbidities: a population-based study. *Rheumatology International* 37, 2 (2017), 313–322.
- [43] Christopher H. Jackson, Linda D. Sharples, Simon G. Thompson, Stephen W. Duffy, and Elisabeth Couto. 2003. Multistate Markov models for disease progression with classification error. *Journal of the Royal Statistical Society: Series D (The Statistician)* 52, 2 (2003), 193–209.
- [48] Pamala D Larsen. 2017. *Chronic Illness: Impact and Intervention* (10 ed.). Burlington: Jones & Bartlett Learning.
- [45] Marios Kokkodis. 2018. Dynamic Recommendations for Sequential Hiring Decisions in Online Labor Markets. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (London, United Kingdom) (KDD '18)*. ACM SIGKDD Conference on Knowledge Discovery & Data Mining (KDD), New York, NY, USA, 453–461.
- [41] Kristy Iglay, Hakima Hannachi, Patrick Joseph Howie, Jinfei Xu, Xueying Li, Samuel S. Engel, Lori M. Moore, and Swapnil Rajpathak. 2016. Prevalence and co-prevalence of comorbidities among patients with type 2 diabetes mellitus. *Current Medical Research and Opinion* 32, 7 (2016), 1243–1252. PMID: 26986190.
- [46] Martijn Lappenschaar, Arjen Hommersom, Peter J.F. Lucas, Joep Lagro, and Stefan Visscher. 2013. Multilevel Bayesian networks for the analysis of hierarchical health care data. *Artificial Intelligence in Medicine* 57, 3 (2013), 171–183.
- [49] Vianey Leos-Barajas and Théo Michelot. 2018. An Introduction to Animal Movement Modeling with Hidden Markov Models using Stan for Bayesian Inference. *arXiv.org* (2018).
- [50] Andrea Martino, Giuseppina Guatteri, and Anna Maria Paganoni. 2020. Multi-variate Hidden Markov Models for disease progression. *Statistical Analysis and Data Mining: The ASA Data Science Journal* 13, 5 (2020), 499–507.
- [52] Katherine L. Musliner, Trine Munk-Olsen, Thomas M. Laursen, William W. Eaton,

- Peter P. Zandi, and Preben B. Mortensen. 2016. Heterogeneity in 10-Year Course Trajectories of Moderate to Severe Major Depressive Disorder: A Danish National Register-Based Study. *JAMA Psychiatry* 73, 4 (04 2016), 346–353.
- [51] Michael Mueller-Peltzer, Stefan Feuerriegel, Anne Molgaard Nielsen, Alice Kongsted, Werner Vach, and Dirk Neumann. 2020. Longitudinal healthcare analytics for disease management: Empirical demonstration for low back pain. *Decision Support Systems* (2020), Article 113271.
- [47] Martijn Lappenschaar, Arjen Hommersom, Peter J.F. Lucas, Joep Lagro, Stefan Visscher, Joke C. Korevaar, and François G. Schellevis. 2013. Multilevel temporal Bayesian networks can model longitudinal change in multimorbidity. *Journal of Clinical Epidemiology* 66, 12 (2013), 1405–1416.
- [53] A. V. Nefian, L. Liang, X. Pi, L. Xiaoxiang, C. Mao, and K. Murphy. 2002. A coupled HMM for audio-visual speech recognition. In *2002 IEEE International Conference on Acoustics, Speech, and Signal Processing*, Vol. 2. II–2013–II–2016.
- [56] Zhaozhi Qian, Ahmed M. Alaa, Alexis Bellot, Jem Rashbass, and Mihaela van der Schaar. 2020. Learning Dynamic and Personalized Comorbidity Networks from Event Data using Deep Diffusion Processes. *arXiv:2001.02585 [cs.LG]*
- [61] Chris Sherlock, Tatiana Xifara, Sandra Telfer, and Mike Begon. 2013. A coupled hidden Markov model for disease interactions. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 62, 4 (2013), 609–627.
- [55] Jennifer Pohle, Roland Langrock, Mihaela van der Schaar, Ruth King, and Frants Havmand Jensen. 2020. A primer on coupled state-switching models for multiple interacting time series. *Statistical Modelling* (2020).
- [54] Yilmazcan Oez yurt, Mathias Kraus, Tobias Hatt, and Stefan Feuerriegel. 2021. AttDMM: An Attentive Deep Markov Model for Risk Scoring in Intensive Care Units. *arXiv preprint arXiv:2102.04702* (2021).
- [57] L. R. Rabiner. 1989. A tutorial on hidden Markov models and selected applications in speech recognition. *Proc. IEEE* 77, 2 (Feb 1989), 257–286.
- [59] Lisa K. Richardson, Leonard E. Egede, Martina Mueller, Carrae L. Echols, and Mulugeta Gebregziabher. 2008. Longitudinal effects of depression on glycemic control in veterans with Type 2 diabetes. *General Hospital Psychiatry* 30, 6 (2008), 509–514.
- [60] Steven L Scott, Gareth M James, and Catherine A Sugar. 2005. Hidden Markov Models for Longitudinal Comparisons. *J. Amer. Statist. Assoc.* 100, 470 (2005), 359–369.
- [62] Kenneth E. Shirley, Dylan S. Small, Kevin G. Lynch, Stephen A. Maisto, and David W. Oslin. 2010. Hidden Markov models for alcoholism treatment trial data. *The Annals of Applied Statistics* 4, 1 (2010), 366–395.
- [58] I. Rezek and S. J. Roberts. 2000. Estimation of coupled hidden Markov models with application to biosignal interaction modelling. In *Neural Networks for Signal Processing X. Proceedings of the 2000 IEEE Signal Processing Society Workshop (Cat.No.00TH8501)*, Vol. 2. 804–813 vol.2.
- [63] Anselm L Strauss and Juliet M Corbin. 1988. *Unending work and care: managing chronic illness at home*. San Francisco, Calif. : Jossey-Bass.
- [66] Claudio Terzano, Vittoria Conti, Fabio Di Stefano, Angelo Petroianni, Daniela Ceccarelli, Elda Graziani, Salvatore Mariotta, Alberto Ricci, Antonio Vitarelli, Giovanni Puglisi, Corrado De Vito, Paolo Villari, and Luigi Allegra. 2010. Comorbidity, Hospitalization, and Mortality in COPD: Results from a Longitudinal Study. *Lung* 188, 4 (2010), 321–329.
- [68] Panayiota Touloupou, Bärbel Finkenstädt, and Simon E. F. Spencer. 2020. Scalable Bayesian Inference for Coupled Hidden Markov and Semi-Markov Models. *Journal of Computational and Graphical Statistics* 29, 2 (2020), 238–249.

- [64] Sucar. 2015. Probabilistic Graphical Models. Springer London. [65] R. Sukkar, E. Katz, Y. Zhang, D. Raunig, and B. T. Wyman. 2012. Disease progression modeling using Hidden Markov Models. In 2012 Annual International Conference of the IEEE Engineering in Medicine and Biology Society. 2845–2848.
- [67] Keith G. Tolman, Vivian Fonseca, Anthony Dalpiaz, and Meng H. Tan. 2007. Spectrum of Liver Disease in Type 2 Diabetes and Management of Patients With Diabetes and Liver Disease. *Diabetes Care* 30, 3 (2007), 734–743.
- [69] Jose M. Valderas, Barbara Starfield, Bonnie Sibbald, Chris Salisbury, and Martin Roland. 2009. Defining Comorbidity: Implications for Understanding Health and Health Services. *The Annals of Family Medicine* 7, 4 (2009), 357–363.
- [70] Marjan van den Akker, Frank Buntinx, and J André Knottnerus. 1996. Comorbidity or multimorbidity. *European Journal of General Practice* 2, 2 (1996), 65–70.
- [72] Aki Vehtari, Andrew Gelman, and Jonah Gabry. 2017. Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Statistics and Computing* 27, 5 (2017), 1413–1432.
- [71] Eva van Weenen and Stefan Feuerriegel. 2020. Estimating Risk-Adjusted Hospital Performance. *IEEE Big Data* (2020).
- [73] Concepción Violán, Sergio Fernández-Bertolín, Marina Guisado-Clavero, Quintí Foguet-Boreu, Jose M. Valderas, Josep Vidal Manzano, Albert Roso-Llorach, and Margarita Cabrera-Bean. 2020. Five-year trajectories of multimorbidity patterns in an elderly Mediterranean population using Hidden Markov Models. *Scientific Reports* 10, 1 (2020), 16879.
- [75] Iris Walraven, M. Ruth Mast, Trynke Hoekstra, A. P. Danielle Jansen, Amber A. W. A. van der Heijden, Simone P. Rauh, Femke Rutters, Esther van 't Riet, Petra J. M. Elders, Annette C. Moll, Bettine C. P. Polak, Jacqueline M. Dekker, and Giel Nijpels. 2015. Distinct HbA1c trajectories in a type 2 diabetes cohort. *Acta Diabetologica* 52, 2 (2015), 267–275.
- [76] Lijuan (Peggy) Wang and Scott E. Maxwell. 2015. On Disaggregating Between-Person and Within-Person Effects With Longitudinal Data Using Multilevel Models. *Psychological Methods* 20, 1 (2015), 63–83.
- [78] Jennifer L. Wolff, Barbara Starfield, and Gerard Anderson. 2002. Prevalence, Expenditures, and Complications of Multiple Chronic Conditions in the Elderly. *Archives of Internal Medicine* 162, 20 (11 2002), 2269–2276.
- [74] Barbora Vozarova, Norbert Stefan, Robert S. Lindsay, Aramesh Saremi, Richard E. Pratley, Clifton Bogardus, and P. Antonio Tataranni. 2002. High Alanine Amino-transferase Is Associated With Decreased Hepatic Insulin Sensitivity and Predicts the Development of Type 2 Diabetes. *Diabetes* 51, 6 (2002), 1889–1895.
- [79] World Health Organization et al. 2016. Global report on diabetes. (2016).
- [77] Peiming Wang and Martin L. Puterman. 2001. Analysis of Longitudinal Data of Epileptic Seizure Counts – A Two-State Hidden Markov Regression Approach. *Biometrical Journal* 43, 8 (2001), 941–962.
- [81] James Zabora, Karlynn BrintzenhofeSzoc, Barbara Curbow, Craig Hooker, and Steven Piantadosi. 2001. The prevalence of psychological distress by cancer site. *Psycho-Oncology* 10, 1 (2001), 19–28.
- [83] Hamed Majidi Zolbanin, Dursun Delen, and Amir Hassan Zadeh. 2015. Predicting overall survivability in comorbidity of cancers: A data mining approach. *Decision Support Systems* 74 (2015), 150 – 161.
- [82] Wei Zhang, Zhaobin Kuang, Peggy Peissig, and David Page. 2020. Adverse Drug Reaction Discovery from Electronic Health Records with Deep Neural Networks. In *Proceedings of the ACM Conference on Health, Inference, and Learning* (Toronto, Ontario, Canada) (CHIL '20). Association for Computing Machinery, New York, NY, USA, 30–39.

requires [84] Walter Zucchini, Iain L MacDonald, and Roland Langrock. 2017. Hidden Markov models for time series: an introduction using R. CRC press.
 [84] Johnson, A., Pollard, T., Shen, L. *et al.* MIMIC-III, a freely accessible critical care database. *Sci Data* **3**, 160035 (2016). <https://doi.org/10.1038/sdata.2016.35>

A ADDITIONAL TABLES

Table 6: Transition matrix (with corresponding α).

Latent states from	to 1	2	3	4	
1	(1,1)	—	α_{12}	α_{13}	α_{14}
2	(1,2)	α_{21}	—	α_{23}	α_{24}
3	(2,1)	α_{31}	α_{32}	—	α_{34}
4	(2,2)	α_{41}	α_{42}	α_{43}	—

Table 7: Summary statistics

Model variable	Overall (N=3253)
Blood glucose level [in mg/dl]	
Mean (SD)	99.59 (11.68)
Range	62.00 – 211.00
Alanine aminotransferase level [in u/l]	
Mean (SD)	24.23 (11.43)
Range	4.05 – 143.40
BMI	
Mean (SD)	30.55 (5.86)
Range	15.48 – 66.67
Sex	
Male	1194 (36.7%)
Female	2059 (63.3%)
Age [in years]	
Mean (SD)	55.34 (8.27)
Range	24.00 – 70.00
Metformin treatment	
Not prescribed (= 0)	3080 (94.7%)
Prescribed (= 1)	173 (5.3%)

Omitted (for brevity): time-since-prediabetes, Metformin (lag=1)

SD: standard deviation

B MODEL DIAGNOSTICS

The following model diagnostics are performed in accordance with standard Bayesian modelling technique [28]. This is done to assure the MCMC algorithm's convergence and, as a result, precise estimates. First, we looked at the effective sample size n_{eff} , which showed that the number of MCMC samples was adequate. Second, for each model parameter, we calculated the Gelman- Rubin convergence diagnostic $\hat{R}$. The $\hat{R}$ is below the critical threshold of 1.1, indicating that the MCMC chain is convergent. Finally, we obtained trace plots (Fig. 4). The trace plots illustrate the MCMC draws of the emission mean for different latent states

and illnesses (i.e., $\mu_s^{(A)}$ and $\mu_s^{(B)}$, $s \in \{1, 2\}$). The trace plots indicate that the chains are nicely intermingled. The previous diagnostics, taken together, show that the MCMC algorithm has converged.

Note: Trace plot for 1500 iterations of each chain (warmup of MCMC algorithm omitted).

Figure 4: Trace plots of emission means $\mu_s^{(A)}$ and $\mu_s^{(B)}$

We further assessed our model by performing posterior predictive checks in accordance with Bayesian modelling recommendations [25, 28]. New observations were simulated using the Viterbi algorithm, which was based on the posterior draws of the model parameters as well as the most likely latent states. Figure 5 shows the results. We use it to compare the generated observations' credible intervals to the real observations. Our model appears to suit the data well, based on the results. We also tested if Coupled-HMM could recover specific parameters from simulated data. The underlying dynamics being successfully captured are represented with the below results that are confirmatory.

(a) Observation $Y_{it}^{(A)}$: Blood glucose level (in log)

(b) Observation $Y_{it}^{(B)}$: Alanine aminotransferase level (in log)

Note: The plots compare the 50 % (light color) and 90 % (dark color) credible intervals of simulated observations against the actual observations. Reported are blood glucose levels (top) and alanine aminotransferase levels (bottom).

Figure 5: Posterior predictive checks (indicating the accuracy of the model on simulated data).