Modeling Longitudinal Dynamics of Comorbidities

Basil Maag
ETH Zurich
Zurich, Switzerland
basilmaag@gmail.com

Stefan Feuerriegel
ETH Zurich
Zurich, Switzerland
sfeuerriegel@ethz.ch

Mathias Kraus
ETH Zurich
Zurich, Switzerland
mathiaskraus@ethz.ch

Maytal Saar-Tsechansky
University of Texas at Austin
Austin, Texas, USA
maytal.saar-
tsechansky@mcombs.utexas.edu

Thomas Züger
Inselspital, Bern, University Hospital,
University of Bern, Bern, Switzerland
ETH Zurich, Zurich, Switzerland
thomas.zueger@insel.ch

ABSTRACT

In medicine, comorbidities refer to the presence of multiple, co-occurring diseases. Due to their co-occurring nature, the course of one comorbidity is often highly dependent on the course of the other disease and, hence, treatments can have significant spill-over effects. Despite the prevalence of comorbidities among patients, a comprehensive statistical framework for modeling the longitudinal dynamics of comorbidities is missing. In this paper, we propose a probabilistic model for analyzing comorbidity dynamics over time in patients. Specifically, we develop a coupled hidden Markov model with a personalized, non-homogeneous transition mechanism, named Comorbidity-HMM. The specification of our Comorbidity-HMM is informed by clinical research: (1) It accounts for different disease states (i. e., acute, stable) in the disease progression by introducing latent states that are of clinical meaning. (2) It models a coupling among the trajectories from comorbidities to capture co-evolution dynamics. (3) It considers between-patient heterogeneity (e. g., risk factors, treatments) in the transition mechanism. Based on our model, we define a spill-over effect that measures the indirect effect of treatments on patient trajectories through coupling (i. e., through comorbidity co-evolution). We evaluated our proposed Comorbidity-HMM based on 675 health trajectories where we investigate the joint progression of diabetes mellitus and chronic liver disease. Compared to alternative models without coupling, we find that our Comorbidity-HMM achieves a superior fit. Further, we quantify the spill-over effect, that is, to what extent diabetes treatments are associated with a change in the chronic liver disease from an acute to a stable disease state. To this end, our model is of direct relevance for both treatment planning and clinical research in the context of comorbidities.

CCS CONCEPTS

• **Applied computing** → **Health informatics**; • **Mathematics of computing** → *Bayesian computation*; *Markov processes*.

KEYWORDS

Bayesian analysis, Hidden Markov model, Longitudinal data analysis, Disease dynamics, Comorbidity

ACM Reference Format:

Basil Maag, Stefan Feuerriegel, Mathias Kraus, Maytal Saar-Tsechansky, and Thomas Züger. 2021. Modeling Longitudinal Dynamics of Comorbidities. In *ACM Conference on Health, Inference, and Learning (ACM CHIL '21)*, April 8–10, 2021, Virtual Event, USA. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3450439.3451871>

1 INTRODUCTION

In medicine, comorbidities refer to the simultaneous occurrence of several diseases affecting patient health [70]. Recent estimates suggest that 50 % of hospitalizations in the U.S. involve patients where eleven and more different disease codes are present [71]. Similarly, multiple diseases and, especially, chronic ones simultaneously affect the health of a large proportion of elderly patients [5, 78]. This calls for a better understanding regarding the interaction of diseases, that is, the dynamics behind comorbidities.

Comorbidities are characterized by a concomitant or concurrent evolution, whereby the course of one disease is highly dependent on the course of the other disease [70]. An example are chronic pain and depression. Both often co-occur and their evolution is further highly dependent. Chronic pain often causes a depression and, vice versa, a depression increases typically the perceived level of pain [4, 27]. Another common example of comorbidities are diabetes and chronic liver diseases (e. g., non-alcoholic fatty liver disease) [34, 37, 67]. The evolution of both diseases is known to be subject to profound interactions and metabolic alterations. This is also reflected in a recent proposal to change the nomenclature from non-alcoholic fatty liver disease to metabolic associated fatty liver disease (MAFLD) [21]. To this end, a better understanding of the co-evolution could inform treatment planning and is thus of direct clinical relevance. However, this requires models where the co-evolution is considered, that is, where disease interactions over time are subject to statistical modeling.

Statistical models that are tailored for comorbidities are scarce. One stream of literature has studied which diseases are often found

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ACM CHIL '21, April 8–10, 2021, Virtual Event

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8359-2/21/04...\$15.00

<https://doi.org/10.1145/3450439.3451871>

together [32, 40, 73]. However, this question relates to unsupervised learning (e. g., clustering) and does not analyze the longitudinal interactions of the diseases. Another stream of literature is concerned with the impact of comorbidities on patient health [e. g., 11, 27, 51, 59, 66] or readmission risk [71]. However, in these works, comorbidities are merely used as independent variables but not as dependent variables, that is, outcomes of interest. Because of this, such models cannot describe the co-evolution among comorbidities. This motivates our research: *how does treating disease A affect the course (disease state) of disease B?* To answer this question, statistical inferences are needed that allow for **interpretability**.

In medicine, the course of diseases is commonly classified by different trajectory phases called “disease states”, whereby diseases undergo different phases of *acute* and *stable* states. These disease states underlie decision-making in a clinical practice as treatment plans and regimens differ depending on whether a disease is in an *acute* or in a *stable* state (cf. the Corbin-Strauss trajectory framework [63]). However, the disease states cannot be directly observed [14]. Specifically, lab measurements provide only noisy realizations of the disease state and, therefore, both measurements (e. g., glucose levels or even other forms of symptoms) and disease states are related only stochastically. To reflect this, a common approach is to model health trajectories in a way that disease states are formalized through the use of latent states and thereby represent hidden Markov models (HMMs) [e. g., 6, 50, 54, 60, 62, 65]. In this work, we adhere to the so-called trajectory framework [14, 63] and describe the evolution of disease based on disease states, yet we extend the HMM-based framework from a single-disease to a multi-disease setting.

Model: In this work, we develop a probabilistic model for analyzing the longitudinal progression of comorbidities in patients. Specifically, we propose a tailored coupled hidden Markov model (CHMM) with a personalized, non-homogeneous transition mechanism, named Comorbidity-HMM. Informed by clinical practice, we specify the Comorbidity-HMM as follows: (1) We account for different states (i. e., *acute*, *stable*) in the disease progression by introducing latent states that are of clinical meaning. (2) We model a coupling among the trajectories from comorbidities to capture co-evolution dynamics. (3) We consider between-patient heterogeneity (e. g., risk factors, treatments) in the transition mechanism. We further provide a fully Bayesian estimation framework through Markov chain Monte Carlo (MCMC). Based on our model, we formalize and estimate a spill-over effect which quantifies the indirect effect of treatments through the co-evolution of the diseases. This allows for statistical inferences that ensure interpretability.

Findings: Our Comorbidity-HMM is evaluated based on a large longitudinal dataset. Specifically, we leverage electronic health records from 675 patients over the course of 10 years. The electronic health records document the joint progression of (A) diabetes and (B) chronic liver disease. Patients at risk of developing diabetes mellitus type 2 were chosen as this disease is among the most common diseases worldwide and its presence is often associated with a chronic liver disease as a concomitant comorbidity [34, 37, 67]. We obtain three main findings. First, modeling disease interactions through means coupling results in a superior model fit. Second, an *acute* state in one of the two diseases is associated with an increased probability of transitioning to an *acute* state in the other disease.

Third, a treatment targeting diabetes has considerable spill-over effects: it is not only associated with a decrease of the risk of an *acute* disease state in diabetes but also in chronic liver disease. Altogether, the results demonstrate empirically the importance of modeling the co-evolution among comorbidities.

Contributions:¹ Our work advances existing research on modeling longitudinal disease dynamics in the following ways.

- (1) To the best of our knowledge, our Comorbidity-HMM is the first statistical model that is specifically tailored to capture longitudinal dynamics among comorbidities.
- (2) For this purpose, we develop a coupled hidden Markov model with a non-homogeneous transition mechanism for longitudinal data and, by deriving the likelihood, provide a Bayesian estimation procedure.
- (3) We formalize and estimate a spill-over effect that measures the indirect effect of treatments through disease interactions. This allows for direct interpretability as it quantifies empirically how treating disease A affects the course (disease state) of disease B. Thereby, we generate new knowledge for clinical practice (e. g., treatment planning) and research.

2 RELATED WORK

2.1 Comorbidities

In medicine, the term “comorbidity” refers to co-occurring diseases (see [70]). In the literature, comorbidities are sometimes classified into a primary disease and secondary diseases, yet such a classification is not universally applicable [23, 69]. Hence, we later make no such distinction and simply refer to disease A and disease B. Both diseases may simply co-occur by chance, yet they might also be subject to a causal relationship or due to correlated or shared risk factors [69]. This means that interactions among diseases occur. Specifically, by treating one disease, one oftentimes can also expect a better course of co-occurring diseases.

Comorbidities are widespread in clinical practice. In particular, patients who are hospitalized often suffer from multiple comorbidities (e. g., every second hospitalization lists eleven or more different disease codes [71]). Due to their prevalence, comorbidities are also of direct relevance for clinical decision-making. In particular, the presence of comorbidities might worsen health outcomes and, hence, clinical practice emphasizes that their treatment should be given special consideration [30].

In clinical practice, there are numerous examples of comorbidities. For instance, depression and chronic pain occur often jointly, especially their severity is often dependent [4, 27]. Similarly, up to 40 % of cancer patients also suffer from psychological distress such as anxiety or depression [81]. Further, patients with diabetes are also frequently diagnosed with additional comorbidities [41], such as hypertension, cardiovascular, or, as studied in this paper, chronic liver diseases (e. g., non-alcoholic fatty liver disease) [34, 37, 67].

2.2 Statistical Inferences using Comorbidity Data

Despite the importance of comorbidities in medicine, there are only few studies that have incorporated their structure into the model.

¹Our code is available via <https://github.com/mb2019/Comorbidity-HMM>

We provide a summary of key references below. While they also use data on comorbidities, their objective is different from our work.

One stream in the literature has used statistical models to detect co-occurrence patterns among diseases [e. g., 32, 40, 73]. This allows to make associations of which diseases co-occur and should thus be regarded as comorbidities. However, these studies address the question of which comorbidities appear frequently together but they do not model their course. For this different models can be used. For instance, comorbidities can be represented as networks in which nodes correspond to symptoms and edges to potentially causal relationships [16]. Based on the network, shared or overlapping symptoms of distinct disorders then indicate the onset of comorbidities. This hypothesis has also been tested using dynamic structural equation models [9, 31] or deep diffusion processes [56]. This answers the question of *which* diseases co-occur but not *how*, that is, without providing insights into longitudinal disease interactions. Similarly, Bayesian networks have been used to investigate comorbidities and their temporal dependencies [22, 46, 47]. However, these works offer only little insights into the actual disease progression such as underlying disease states.

Other works use data on comorbidities as input when making inferences regarding patient outcomes [e. g., 11, 27, 59, 66, 83] including readmission risk [71]. This allows to control for the fact that the co-occurrence of diseases often introduces associations with patient outcomes that are of nonlinear form. For instance, patient survival might be lowered when there are many co-occurring diseases. In clinical practice, this is reflected in different measures such as, e. g., the Charlson comorbidity index [13]. However, the aforementioned works consider comorbidities as independent variables and not as dependent variables.

2.3 Hidden Markov Models

Hidden Markov models (HMMs) are a flexible statistical technique for modeling time series. HMMs model a sequence of latent states that follow a Markov chain, based on which, observations are emitted [29]. HMMs have found widespread application in modeling human behavior. Examples include clickstream data [35], urban activities [80], and hiring decisions [45].

HMMs have previously been applied to data on patient health [e. g., 2, 42, 43, 54, 60, 65]. In these works, HMMs are further adapted in order to reflect patient trajectories through following specification. First, the latent states commonly correspond to different diseases states (i. e., phases) in the trajectory [see 14]. Second, covariates as well as random effects are incorporated in the transition mechanism in order to account for patient heterogeneity and other sources of unobserved heterogeneity [2]. However, the aforementioned HMMs describe the evolution of single-disease outcomes but not the evolution in multi-disease settings, that is, comorbidities.

There exist several variants of HMMs [e. g., 64, pp. 76–77]. For instance, coupled HMMs are an extension of HMMs that could potentially model the evolution of multiple interacting processes. These CHMMs allow the latent states of distinct but possibly correlated processes to be conditional dependent [8]. Here the latent states of different sequences can be linked using a Cartesian product, which gives the basis of a so-called Cartesian product CHMM. In a simulation study, such a Cartesian product CHMM was shown to

outperform ordinary and multivariate HMMs when the underlying processes are dependent [55].

CHMMs have been applied in various areas, including speech recognition [53], epileptic seizure detection [15], and sleep staging [58]. Furthermore, they have been used to model the spreading of infectious diseases [19, 68]. In [55], a Cartesian product CHMM is developed for modeling the joint progression of lab measurements in a single-disease setting from intensive care units. In ecology, a CHMM was utilized to investigate the interdependence among voles from different forest regions [61]. More specifically, it was examined if the occurrence of one disease makes the contagion of another one more likely. Further, a multi-chain Markov switching model akin to a CHMM was employed to check if there is volatility spillover between geographical financial markets [26].

Research gap: The importance of comorbidities for clinical practice has been widely recognized, and yet statistical frameworks for modeling their co-evolution over time have received little attention. However, understanding the underlying disease dynamics could be of significant value for treatment planning. To this end, we develop a Comorbidity-HMM that is tailored to capture comorbidity dynamics. Specifically, we model the disease interactions among comorbidities in patients based on longitudinal data.

3 PROPOSED MODEL: Comorbidity-HMM

3.1 Problem Statement

We aim to model the co-evolution of comorbidities. For this, we make use of a longitudinal dataset with health information from patients. We further consider two diseases as given by disease **A** and disease **B**. Here we explicitly assume that the evolution of both is characterized by disease interactions. This motivates a better understanding of *how* the course of both diseases interact. Specifically, we seek to answer the question whether there are potential spill-over effects: *how does treating disease A affect the state of disease B?* For instance, to what extent does Metformin as diabetes treatment influence the risk of an *acute* state of a co-occurring liver disease? Formally, we want to measure the influence of one disease on the state of the other. We thus develop a model that ensures interpretability in order to allow for such statistical inferences.

To be of clinical relevance, our model specification accommodates further properties, namely (1) disease states, (2) coupling, and (3) between-patient heterogeneity, as described in the following.

(1) Disease states. The course of many diseases (especially chronic or otherwise long-lasting diseases) undergoes different phases over time, where the disease is in an “*acute*” or “*stable*” phase [14]. In medicine, the phases are typically termed “disease states”. These states themselves cannot be directly identified and, instead, can only be recovered from measurements [14]. To accommodate this, we follow prior literature [e. g., 6, 50, 54, 60, 62, 65] and model disease states through latent states. We later name the latent states “*acute*” and “*stable*” to reflect their clinical meaning. Consistent with prior literature [e. g., 6, 50, 54, 60, 62, 65], we thus build upon the HMM-based framework [57] for modeling the course of diseases **A** and **B**. Intuitively, we represent both through two separate HMMs.

(2) Coupling. We expect disease interactions among comorbidities and, therefore, explicitly assume the states of disease **A** and **B**

to be dependent. As an example, when a disease **A** transitions from a **stable** to an **acute** state, it should also increase the probability of disease **B** to move from a **stable** to an **acute** state. For this reason, we later introduce a coupling between the HMMs from disease **A** and **B**.

(3) **Between-patient heterogeneity.** Disease dynamics are known to vary extensively across patients [e. g., 52, 75]. In medicine, this is described via risk factors (e. g., age, gender) and the presence of treatments [e. g., 36]. Therefore, our model accommodates between-patient heterogeneity as follows. On the one hand, we incorporate patient-level covariates that denote risk factors and treatments. These are entered in the transition mechanism of the HMMs, so that risk factors (as well as treatments) affect the disease dynamics: this allows risk factors (as well as treatments) to increase the propensity of **acute** vs. **stable** disease states.² On the other hand, we account for unobserved heterogeneity. Here we follow prior literature [7, 17, 20, 33, 38, 76] and control for within effects in the treatment variable (see Sec. 4.2 for details).

In the following, we develop a model named Comorbidity-HMM. Our Comorbidity-HMM accommodates the above properties (1)–(3). Based on the Comorbidity-HMM, we define a spill-over effect which estimates the influence of a treatment from disease **A** on the state from disease **B**. Here we follow clinical guidelines that base treatment planning primarily on disease states instead of measurements or symptoms [14, 48, 63]. For that reason, we define the spill-over effect with regard to disease states as opposed to measurements or symptoms. This allows us to make statistical inferences of how a treatment targeting disease **A** render the risk of disease **B** to become **acute** vs. **stable**.

3.2 Model Specification

Our Comorbidity-HMM consists of the following five components: (i) latent states representing the disease states; (ii) observations representing the measurements or symptoms; (iii) an emission component linking states and observations; (iv) a transition mechanism; and (v) a coupling. Based on it, we model the evolution for each of two diseases **A** and **B** through a HMM-based framework that is additionally subject to coupling. The components (i)–(v) are specified in the following:

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

(ii) Latent states: The latent states represent the different disease states (i. e., **acute** vs. **stable**). Here we introduce latent states $S_{it}^{(A)} \in \mathcal{S}^{(A)}$ and $S_{it}^{(B)} \in \mathcal{S}^{(B)}$ for patient i at time step t , where the latent states belong to disease **A** and **B**, respectively. As such, each disease

has a separate latent state and can thus attain a different state. For each disease, we assume two different states “1” and “2”, i. e., $\mathcal{S}^{(A)} = \mathcal{S}^{(B)} = \{1, 2\}$. For ease of notation, we refer to state 1 as “**stable**” and state 2 as “**acute**”. In total, this allows for $|\mathcal{S}^{(A)}| \times |\mathcal{S}^{(B)}|$ different combinations.

Table 1: Mapping of latent states.

Disease states		Latent states	Global state
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

We now define a global latent state based on the Cartesian product of the two separate latent states. Formally, we define $S_{it}^{(G)} = (S_{it}^{(A)}, S_{it}^{(B)})$. This is later necessary when we detail the coupling as part of component (v). The underlying mapping is listed in Tbl. 1. For instance, a global state $S_{it}^{(G)} = (2, 2)$ means that both 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 emission component defines the probability of an observation conditional on the latent state. In our case, it defines the probability of observing a measurement or symptom conditional on disease state. For this, both observations $Y_{it}^{(A)}$ and $Y_{it}^{(B)}$ of the two diseases are assumed to depend on the latent states $S_{it}^{(A)}$ and $S_{it}^{(B)}$, respectively. Given a latent state s , we model them to be conditionally normally distributed via

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

with mean $\mu_s^{(A)}, \mu_s^{(B)}$ and variance $\sigma_{(A)}^2, \sigma_{(B)}^2$. Here the observations are modeled with an individual mean per state s and disease. Hence, an **acute** state from disease **A** might have, on average, larger values as observed measurements or symptoms. The variance is disease-specific but independent of the state. In the above specification, we used a normal distribution to model the observations as, in our experiments, the measurements (e. g., blood glucose levels) follow a normal distribution.

(iv) Transition mechanism: The transition mechanism specifies the probability of moving from one latent disease state to another latent disease state. As typical for HMMs, we assume the latent states to follow a Markov process. However, we accommodate additional covariates (i. e., risk factors, treatments) whereby we yield a non-homogeneous transition mechanism that is personalized to patient profiles.

Formally, we model the probabilities of the global states via a multinomial logit link function [2, 49]. Hence, the transition probability $\gamma_{it}^{(j \rightarrow k)}$ denotes the probability of patient i transitioning from

²The choice that risk factors (and treatments) should enter not the emission but the transition mechanism is due to clinical research. The reasoning is the following: variables in the emission only affect observed measurements (e. g., they affect pain resistance) as they moderate how **acute** vs. **stable** disease states link to measurements; yet they cannot affect the course of the disease progression. This choice is later validated in our numerical experiments.

global state j at time t to global state k at time $t + 1$. We specify the transition probability via

$$\begin{aligned} \gamma_{it}^{(j \rightarrow k)} &= P\left(S_{i,t+1}^{(G)} = k \mid S_{it}^{(G)} = j\right) \\ &= \frac{\exp(\eta_{jk})}{\sum_{l=1}^{|S^{(G)}|} \exp(\eta_{jl})}, \end{aligned} \quad (2)$$

where

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

The above logit makes use of further variables: 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, treatments). The coefficients of the covariates x_{it} are given by β_{jk} . These capture an elevated probability of **acute** states (over **stable** states) due to risk factors. Covariates from the treatment are further processed to control for within effects in the treatment variable (see Sec. 4.2). Coefficients for $j = k$ are set to zero in order to ensure identifiability. Hence, transitions where the same latent disease state is maintained (i. e., recurrent transitions) serve as the reference category; see Tbl. 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 model the interaction of the two diseases by introducing a coupling among their latent states. Specifically, we build upon the concept of a Cartesian product coupled HMM [55]. For this, we draw upon the Cartesian product of the underlying latent states from above. 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 Comorbidity-HMM.

Mathematically, by introducing the global states, the coupling is captured as part of the transition mechanism. The reason is that the transition mechanism is now based on the global state, so that the probability of $S_{it}^{(A)}$ does not only depend on the previous latent

state $S_{i,t-1}^{(A)}$ but also on that from the other disease $S_{i,t-1}^{(B)}$. The same holds true analogously when exchanging disease **A** and **B**. Hence, the transition mechanism accounts for potential co-movements from the states of both diseases.

The coupling among the latent states is best seen in an example. Recall that the transition mechanism depends on an intercept but it is also influenced by further covariates (i. e., risk factors). We now discuss how the transitions of one disease differ depending on the other diseases. For instance, it might be the case that the probability of a transition from a **stable** to an **acute** state for disease **B** is more likely given that disease **A** is in an **acute** state. Such a case could, for instance, arise if the corresponding intercept α_{34} [for transition $(2=\text{acute}, 1=\text{stable}) \rightarrow (2=\text{acute}, 2=\text{acute})$] is substantially larger than the intercept α_{12} [for transition $(1=\text{stable}, 1=\text{stable}) \rightarrow (1=\text{stable}, 2=\text{acute})$]. This means that an **acute** state for disease **A** is associated with an increased propensity of disease **B** also moving to an **acute** state (rather than a **stable** state). Thus, comparing the estimated intercepts of such transitions can offer insights about the underlying disease interactions. Hence, we later report our estimations results for α_{jk} as this indicates the strength of the underlying coupling, i. e., the underlying disease interaction.

3.3 Spill-Over Effect

Treatments such as medications can also be included in the covariates. If such a treatment is specifically targeted at one comorbidity, it will have a direct effect on the state transitions of that disease. For instance, a drug designed for disease **A** may prevent the transition from state $(1, 1)$ to $(2, 1)$ (i. e., prevent **A** from becoming **acute**). However, the same drug might also implicitly effect the state transitions of disease **B** through their coupling. We will therefore refer to the former as direct treatment effect and to the latter as spill-over effects.

An example of the a spill-over effect is the following. Let us consider the state transitions $(2, 2) \rightarrow (1, 2) \rightarrow (1, 1)$. It represents the scenario in which a treatment makes disease **A** transition from an **acute** state to a **stable** one (i. e., direct effect). Subsequently, disease **B** also transitions to a **stable** state in the following time period (i. e., as an additional indirect effect).

We formalize the above spill-over effect as follows. Let z denote a treatment which is targeted at disease **A** and affects the transition probabilities via the covariates x_{it} as denoted in Equation 3. Then, the probability of the aforementioned state transitions of patient i under treatment z is given by

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

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

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

Thus, if the treatment z increases the probability $\xi_{it}(z)$ but is specifically designed for disease **A**, it also affects disease **B** positively through their coupling. 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: We evaluate the fit of our Comorbidity-HMM using standard performance metrics for Bayesian modeling [28], that is, the expected log pointwise predictive density (elpd). The elpd can be efficiently computed using the help of Bayesian leave-one-out cross-validation and Pareto smoothed importance sampling [72]. In our case, we measure model performance for the hold-out sample at patient level (rather than at observation level). For this, we used the loo package as part the statistical software R. Furthermore, we report the widely-applicable information criterion (WAIC).

Model variants: We compare our proposed Comorbidity-HMM against alternative model variants. The model variants: (1) A simplified variant where disease **A** has only a single latent state, i. e., $|\mathcal{S}^{(A)}| = 1$. (2) A simplified variant where disease **B** has only a single latent state, i. e., $|\mathcal{S}^{(B)}| = 1$. These represent simplified disease dynamics with only reduced set of states. Note that both model variants have access to identical data as our Comorbidity-HMM. As we shall see later, the proposed Comorbidity-HMM with $|\mathcal{S}^{(A)}| = |\mathcal{S}^{(B)}| = 2$ is superior, thus implying the necessity of accommodating different disease states.

3.5 Estimation Procedure

Sampling: In the Comorbidity-HMM, all model parameters are obtained via a fully Bayesian estimation [28], specifically through the use of Markov chain Monte Carlo (MCMC) sampling. For this, we used the probabilistic language Stan [12]. Based on it, we obtained posterior estimates from the Hamiltonian Markov chain algorithm together with the No-U-Turn sampler [39]. We ran four chains with 1500 iterations each (1500 additional iterations were discarded as part of a warm-up) to obtain posterior estimates. For this, we first derived the likelihood $\mathcal{L}$ of the Comorbidity-HMM via

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

where $\phi^{(A)}(s_t)$ and $\phi^{(B)}(s_t)$ map the global state to the corresponding individual states of each disease (see Tbl. 1). We further applied

the forward algorithm [84, pp. 36–39] in order to obtain an efficient calculation scheme. Moreover, the state-dependent means were ordered, i. e., $\mu_1^{(A)} < \mu_2^{(A)}, \dots, \mu_1^{(B)} < \mu_2^{(B)}$, to avoid label-switching and to promote identifiability [see 44]. The runtime on typical hardware amounted to around 12 hours.

Priors: We used weakly informative priors for the emission component and the initial state probabilities:

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

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

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

Slightly narrow priors were chosen for the intercepts α_{jk} with the aim to stabilize the model fit. Instead of specifying priors for the coefficients β_{jk} , they were placed on the transformed parameters $\tilde{\beta}_{jk}$, which resulted from a thinned QR decomposition of the centered covariates x_{it} . This corresponds to

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

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

Diagnostics: We follow common guidelines in Bayesian modeling [28] and inspect the convergence of the Markov chains. We further conduct posterior predictive checks of whether we can simulate reasonable new observations based on the estimated model parameters [25]. Both are reported as part of the model diagnostics in Sec. B. In addition to that, we checked if our Comorbidity-HMM was able to recover the specified parameters from entirely simulated data. If not stated otherwise, we later report the posterior mean, as well as the 50 % and 90 % credible intervals (CrI).

4 EMPIRICAL SETUP

4.1 Data

We evaluate our Comorbidity-HMM based on two co-occurring diseases: (**A**) diabetes mellitus type 2 and (**B**) chronic liver disease. Both often co-occur as comorbidities [34, 37, 67]. Diabetes is one of the most common chronic disease, impairing the lives of millions of people worldwide. Diabetes further causes substantial cost during the provision of care [79]. On the other hand, chronic liver disease is a medical condition often co-occurring with diabetes [34, 37, 67]. In our analysis, we use blood glucose level as lab measurements for diabetes and the level of alanine aminotransferase, an enzyme predominantly found in the liver [24, 74], as a lab measurement for chronic liver disease.

We later compute spill-over effects following the treatment of increased blood glucose level with Metformin. Metformin is the first line treatment for diabetes type 2 and lowers overall blood glucose levels. In the literature, Metformin is also considered to have a beneficial impact for patients with liver disease [3, 10, 34]. Hence, this raises the question to what extent spill-over effect influence the disease state for chronic liver disease.

In our analysis, we utilize a longitudinal dataset of patients with prediabetes. The dataset consists of annual electronic health records [1, 82] that also include recordings on lab measurements, treatments, and risk factors (see next section). Overall, the dataset includes 675 patients. Each patient has a time series with recordings between 4 and 10 consecutive years. This amounts to 3253 observations.

4.2 Model Variables

Our Comorbidity-HMM is estimated using the following variables (Tbl. 2).

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

The covariates x_{it} characterize between-patient heterogeneity. These represent risk factors and data on treatments. Here we use the body mass index (BMI), sex, age (at time step $t = 1$), and time-since-prediabetes. We explicitly decompose temporal information into both age and time-since-prediabetes as this allows us to isolate age as a risk factor and potential trends. Data on treatments is given by medication with Metformin. This is a binary variable of whether a Metformin treatment (i. e., 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)

The following procedure was used to infer the so-called *within* effect of Metformin. Metformin is generally expected to reduce or stabilize the glucose levels within patients [3]. However, patients who take Metformin might have higher glucose levels compared to other ones. Thus, the so-called *within* and *between* effects might be conflicting. For that reason, the Metformin variable is centered with the corresponding patient specific mean in order to infer about the *within* effect. This approach is consistent with previous literature [7, 17, 20, 33, 38, 76]. Besides the centered Metformin variable, a lagged version of it was also included. These two Metformin variables will be used to estimate potential treatment and spill-over effects over time.

4.3 Summary Statistics

Summary statistics of the model variables are given in Tbl. 7. Blood glucose levels tend to be elevated as compared to recommendations from the World Health Organization as expected in this population of individuals with (pre-)diabetes. Similarly, alanine aminotransferase values up to 143.40 u/l, suggesting the presence of liver disease in some of the patients [e.g., 34]. Overall, patients are characterized by rather high BMI implying a frequent presence of obesity. This is expected as a high BMI is a known risk factor for diabetes. Only a subset of patients received a Metformin treatment.

We observe a significant correlation between both measurements, namely blood glucose and alanine aminotransferase levels. The Pearson's correlation coefficient amounts to 0.11 ($p < 0.001$). Hence, a higher blood glucose level coincides with a higher alanine

aminotransferase level. This points towards potential interactions among both diseases, yet it does not allow for statistical inferences regarding the underlying longitudinal dynamics.

5 NUMERICAL RESULTS

5.1 Model Performance

The results of the performance comparison are reported in Tbl. 3. The models are evaluated based on the expected pointwise predictive density (elpd). The elpd assesses predictive accuracy in a Bayesian framework and, therefore, the overall fit of a model [28, 72]. A higher elpd indicates a better fit. Among the models which converged, the highest elpd (2,148.78) is registered for the proposed Comorbidity-HMM. It thus achieves the best overall model fit.

We use baseline models that are suitable for jointly modeling multiple time series. Specifically, we draw upon the baselines from [55]. Alternative models are typically designed for single time series or lack a mechanism for disease interactions (see Sec. 2). Hence, their choice for the purpose of our paper is precluded.

Which number of latent states should be selected? We compare our Comorbidity-HMM against baselines with a simplified latent state space, that is, a Comorbidity-HMM where the latent states for either disease **A** or **B** are removed. Here simplified **A** refers to variant with a single latent state for diabetes and, analogously, simplified **B** to a variant with a single latent state for chronic liver disease. As a result, the diseases are assumed to have only a single state (yet which would oppose the Corbin-Straus trajectory framework [14, 63]). Both the simplified **A** (elpd: 1,620.96) and a simplified **B** (elpd: 1,450.83) are outperformed by the proposed Comorbidity-HMM by a considerable margin. Taken together, this confirms the importance of including latent states and thus accounting for different trajectory phases (i. e., disease states). In subsequent sections, we thus report estimation results from the proposed Comorbidity-HMM with 2×2 states as this model is favored in the model comparison.

How should between-patient heterogeneity be modeled? To answer this, we run additional comparisons. First, we draw upon a naïve coupled HMM [8]. This model does not use any covariates for modeling between-patient heterogeneity (i. e., neither in the emission component nor in the transition mechanism). Evidently, the model is inferior. Second, we use the emission coupled HMM from Pohle et al. [55]. This model includes covariates in the emission (whereas our model includes covariates in the transition). However, the chains during MCMC sampling did not converge and, because of this, interpretability for model coefficients is precluded. Moreover, the estimate of the elpd may be unreliable. Third, as a remedy, we modified the emission coupled HMM from Pohle et al. [55]. Here we additionally introduce an ordering of the estimated means besides ordered intercepts in order to prevent label-switching across latent states, set initial values for intercepts based on a k -means procedure [18], and follow common practice in Bayesian modeling whereby the covariates are subject to a QR decomposition. The latter should remove posterior correlations among covariates and thus facilitate convergence of the MCMC chains [28]. To confirm the accuracy of our implementation, we tested the modified emission coupled HMM based on simulated data, where MCMC sampling converged and where the model coefficients were recovered successfully. The

model converged successfully for 1×2 states (but not for the other state combinations). Overall, the model has a poorer fit compared to our proposed model. Altogether, this implies that that covariates should be included in the transition mechanism (and not in the emission component).

What is the utility of modeling disease interactions? By definition, coupled HMMs with only a single latent state cannot account for disease interactions. Hence, we now assess the corresponding benefit. For the naïve coupled HMM, we find that model with 2×2 states performs best, implying that the capturing disease interactions is beneficial. (We cannot comment on the emission coupled HMM due to divergence in the chains.) For the proposed Comorbidity-HMM, we also find that the model with 2×2 states is best. Hence, it must be assumed that both diabetes and chronic liver disease are subject to longitudinal disease interactions. This is later also tested statistically by interpreting the transition probability.

5.2 Estimation Results

Emission component: The estimated parameters from the emission component are reported in Tbl. 4. For diabetes, the *stable* state is associated with an average blood glucose level of only $\exp(4.55) = 94.63$ mg/dl. As expected, the *acute* state has a higher value, with a mean of $\exp(4.70) = 109.95$ mg/dl. For both, the credible intervals do not overlap and are thus different at a statistically significant level.

For chronic liver disease, the mean of the *stable* state amounts to $\exp(2.86) = 17.46$ u/l. This is lower than the mean of the *acute* state (30.88 u/l). The credible intervals do not overlap and, hence, the posterior means of the *stable* and *acute* states are different at the statistically significant level.

Transition mechanism: Estimates of the parameters regarding the state transition probabilities as follows. Due to high number of coefficients, we refrain from presenting them all and rather focus on how Metformin treatment influences the transition probability. This thus quantifies the direct treatment effect of Metformin on the disease state. Informed by clinical research, we expect the following. (1) Recall that Metformin is a diabetes treatment. Hence, it should only have an effect on diabetes and not on chronic liver disease. This is confirmed (all corresponding CrI do not include zero). (2) Metformin has a short-term effectiveness and, hence, we expect to see treatment effects in the non-lagged variable (rather than in the lagged). Again, we yield positive evidence. (3) We expect a positive treatment effect which would be seen in the transition (*acute, acute*) $\rightarrow$ (*stable, acute*). Here a large portion of the probability mass is above zero, indicating a positive treatment effect of Metformin. Here the posterior mean corresponds to an increase of the log-odds of 3.29, resulting in a $\exp(3.29) = 26.84$ odds ratio for a one unit increase in the centered Metformin variable and, hence, making the aforementioned transition more likely than staying in the state (*acute, acute*).

Coupling: We now investigate the coupling among the two comorbidities. For that, we report the baseline probability in the transition mechanism (i. e., the intercept α_{jk}). It quantifies the transition conditional on the latent states of both disease **A** and **B**. To estimate this, all covariates were set to respective sample mean. This corresponds to $x_{it} = \bar{x}_{it}$ in Eq. 3.

Fig. 2 shows the coupling for diabetes. The plot gives the probability of disease **A** (diabetes) to transition from a *stable* to an *acute* latent state *conditional* on the latent state of both diseases. If both disease are independent, then the credible intervals for the conditional transition probability should include zero. Larger conditional transition probability values indicate a stronger coupling. From Fig. 2, we find that an *acute* state for chronic liver disease is associated with an increased probability of moving from *stable* to *acute* in diabetes; see $(1, 2) \rightarrow (2, 2)$. Put simply, if chronic liver disease is *acute*, then it is likely that diabetes also transitions from *stable* to *acute*. Overall, we find strong evidence that diabetes is coupled with chronic liver disease.

The interpretation holds analogously when exchanging **A** and **B**. Here similar effects can also be detected in Fig. 3. Put simply, if diabetes is already *acute*, then it is somewhat likely that chronic liver disease also becomes *acute*. However, the effect is less pronounced and the 50% and 90% credible interval overlap to a great extent. Nevertheless, these findings suggest that the two comorbidities indeed interact with each other.

Note: Here shown are the transitions probabilities that diabetes moves from a *stable* to an *acute* latent state, i. e., $(1, \square) \rightarrow (2, \square)$. Hence, this compares the transition probability when exiting a *stable* chronic liver disease (top) and an *acute* chronic liver disease (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 2: Estimation results for coupling conditional on diabetes.

5.3 Estimated Spill-Over Effect

Lastly spill-over effects are presented. Specifically, we assess the effect of Metformin on liver disease through the co-evolution of the two comorbidities. For this, we consider the state transitions $(2, 2) \rightarrow (1, 2) \rightarrow (1, 1)$. It reflects the transition from *acute* to *stable* in diabetes followed by the transition from *acute* to *stable* in liver disease. In particular, we report the following: (1) Metformin has been prescribed, i. e., $\xi_{it}(z)$. Here the centered variable is set to 0.5, whereby the remaining covariates are also assumed to be equal to the sample mean. (2) No Metformin has been prescribed, i. e., $\xi_{it}(z')$. This corresponds to a centered Metformin variable equal to zero and all remaining covariates equal to the respective sample

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) [†]
Comorbidity-HMM (simplified A)	✗	(1 × 2)	1,620.96	149.61	−3,241.99
Comorbidity-HMM (simplified B)	✗	(2 × 1)	1,450.83	139.70	−2,902.97
Comorbidity-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-applicable 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)				
$\mu_1^{(B)}$	stable	2.86	2.85	2.88
$\mu_2^{(B)}$	acute	3.43	3.41	3.45
$\sigma_{(B)}$	—	0.30	0.29	0.31

mean. (3) The spill-over effect, which amounts to the absolute difference between both. (4) A quotient. It denotes how many times more likely a spill-effect is when comparing Metformin vs. no Metformin. The corresponding quantiles of the transition probabilities of these two scenarios as well as their difference are reported in Tbl. 5.

Overall, we find the following. If no Metformin is prescribed, there is hardly any change in the state for liver disease through the underlying coupling. Here all estimated quantiles are close to zero. If Metformin is prescribed, we observe a non-zero probability that the state from chronic liver disease also improves. By comparing the different quantiles, we observe that the distribution for $\xi_{it}(z)$, i. e., Metformin prescribed, has a very long tail. For many patients on the left-tail, the spill-over effect is small or absent (e. g., due to non-adherence or simply because the treatment is not effective). Based on our results, it cannot be ruled out that, for the 5 % quantile, the Metformin treatment has also negative implications for a co-occurring liver disease.

However, for the patients on the right-tail, there is a large estimated spill-over effect. For instance, for the 50 % quantile, prescribing Metformin (as opposed to non-prescribing) makes a improvement in the liver state 2.898 times more likely. In the 75 % quantile, the relative quotient increases even to 13.695. As shown here, our

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.

model allows to quantify the spill-over effect of diabetes on liver disease due to the use of Metformin. In sum, we find evidence that Metformin stabilizes not only diabetes but, through the coupling, also a co-occurring chronic liver disease.

6 DISCUSSION

Summary of findings: (1) Our results find that our Comorbidity-HMM outperforms alternative model specifications. In particular, it is superior over variants where no disease interaction is modeled. Here we obtain substantial improvements in the model performance as measured by predictive accuracy for Bayesian modeling (i. e., the expected pointwise predictive density, elpd). (2) We provide empirical evidence that both diabetes and chronic liver disease are dependent. This is revealed by the fact that transitions of one

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 centered 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.

disease depend on the state of the other disease at a statistically significant level. 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.

Prior literature suggested mixed findings on the role of Metformin for patients with chronic liver disease [3, 10, 34], especially as a direct effect would not necessarily be supported by the molecular mechanism of action. Here we contribute empirical evidence whereby Metformin has a direct effect on the state for diabetes, while we find no support for a direct effect on the state for chronic liver disease. Instead, we find an indirect effect through the coupling.

Interpretability: Our model is based on a parsimonious specification in order to warrant for interpretability. Here we build upon a central concept from medicine: diseases form a so-called trajectory which undergoes different phases with **acute** and **stable** disease states [14, 63]. These disease states are directly relevant for clinical decision-making as different treatment plans should be applied depending on whether a patient is in an **acute** vs. a **stable** state. Specifically it suggests that clinical decision-making should tailor care to the underlying disease state (and not measurements or symptoms as these are only noisy observations of the state) [48]. Owing to this, prior research has modeled health trajectories through the use of latent states [e. g., 6, 50, 54, 60, 62, 65]. In line with this, we also introduce latent disease states and use them for coupling the states of different diseases, and, on top of that, we define our spill-over effect with regard to latent states (as opposed to measurements or symptoms).

Limitations: As with other research in medicine, ours is not free of limitations and we thus make suggestions for future research. First, our data comprises a unique, large-scale collection of electronic health records from a Western health provider. Future research could replicate our approach with other cohorts. Second, we used an extensive set of risk factors informed by prior literature. Nevertheless, other variables might also represent risk factors and, hence, future research could test how the model performance changes by including them. Third, our analysis is based on specific setting, namely the interaction among diabetes and chronic liver disease under Metformin treatment. Here future research could validate our model based on other comorbidities. If desired, the model can also be extended to comorbidities comprising of three (or more) diseases by expanding the underlying latent state space.

Implications: Our model offers insights for clinical practice and research. By inferring spill-over effects, one can assess to what extent treatments for disease **A** also change the course (disease state) of disease **B**. Hence, when seeking to stabilize the course of disease **B**, clinical practitioners can thus decide upon a treatment with a direct effect on **B** or, by contrast, choose a treatment plan for **A** that also has an indirect effect such that disease **B** is stabilized. Here our model offers statistical insights, allowing for evidence-based decisions regarding treatment planning.

Our model was tested based on diabetes and chronic liver disease. However, our model is designed in a generic manner in order to ensure widespread applicability. The clinical co-author from our work foresees several use cases: Here our model could analyze comorbidities that are poorly understood or where there are mixed findings in the clinical literature. Given the vast number of comorbidities, this gives rich opportunities for follow-up research.

7 CONCLUSION

Comorbidities, defined as the presence of multiple co-occurring diseases, are widespread among patients, and yet comprehensive statistical frameworks for modeling the longitudinal dynamics of comorbidities are rare. Such models would allow for novel insights into the co-evolution of comorbidities over time. Specifically, this would allow to answer the question: *how does treating disease A change the state of disease B?* To address this, we developed a probabilistic model for longitudinal data analysis: a coupled hidden Markov model with personalized, non-homogeneous transitions (Comorbidity-HMM). To the best of our knowledge, our Comorbidity-HMM is the first statistical model that is specifically tailored to capture longitudinal dynamics among comorbidities. Our model further ensures interpretability by offering clinically-relevant inferences. In particular, we presented a spill-over effect that measures the indirect effect of a treatment on disease states through the co-development of comorbidities. We evaluated our model based on a longitudinal dataset from 675 patients. We found that Comorbidity-HMM outperforms alternative model specifications in terms of model fit.

ACKNOWLEDGMENTS

Financial supported from the Swiss National Science Foundation (SNSF) through an Eccellenza grant (186932) is acknowledged.

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.
- [2] Rachel MacKay Altman. 2007. Mixed Hidden Markov Models. *J. Amer. Statist. Assoc.* 102, 477 (2007), 201–210.
- [3] Clifford J. Bailey and Robert C. Turner. 1996. Metformin. *New England Journal of Medicine* 334, 9 (1996), 574–579. PMID: 8569826.
- [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.
- [5] Karen Barnett, Stewart W Mercer, Michael Norbury, Graham Watt, Sally Wyke, and Bruce Guthrie. 2012. Epidemiology of multimorbidity and implications for health care, research, and medical education: a cross-sectional study. *The Lancet* 380, 9836 (2012), 37–43.

- [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.
- [7] Andrew Bell and Kelvin 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.
- [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.
- [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.
- [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.
- [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.
- [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.
- [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.
- [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.
- [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.
- [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.
- [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.
- [18] Luis Damiano, Brian Peterson, and Michael Weylandt. 2018. A Tutorial On Hidden Markov Models Using Stan.
- [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.
- [20] Craig K. Enders and Davood Tofghi. 2007. Centering Predictor Variables in Cross-Sectional Multilevel Models: A New Look at an Old Issue. *Psychological Methods* 12, 2 (2007), 121–138.
- [21] Mohammed Eslam, Arun J. Sanyal, Jacob George, Arun Sanyal, Brent Neuschwander-Tetri, Claudio Tiribelli, David E. Kleiner, Elizabeth Brunt, Elisabetta 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.
- [22] Syed Hasib Akhter Faruqi, 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.
- [23] Alvan R. Feinstein. 1970. The pre-therapeutic classification of co-morbidity in chronic disease. *Journal of Chronic Diseases* 23, 7 (1970), 455–468.
- [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.
- [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.
- [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.
- [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.
- [28] Andrew Gelman, John B. Carlin, Hal S. Stern, David B. Dunson, Aki Vehtari, and Donald B. Rubin. 2013. *Bayesian data analysis*. CRC press.
- [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.
- [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.
- [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.
- [32] Marina Guisado-Clavero, Albert Roso-Llorach, Tomàs López-Jimenez, Mariona Pons-Vigués, Quinti 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.
- [33] Ellen L. Hamaker and Raoul P. 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.
- [35] Tobias Hatt and Stefan Feuerriegel. 2020. Early Detection of User Exits from Clickstream Data: A Markov Modulated Marked Point Process Model. In *Proceedings of The Web Conference 2020 (Taipei, Taiwan) (WWW '20)*. International Conference of the World Wide Web (WWW), New York, NY, USA, 1671–1681.
- [36] Tobias Hatt and Stefan Feuerriegel. 2021. Estimating Average Treatment Effects via Orthogonal Regularization. *arXiv preprint arXiv:2101.08490* (2021).
- [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.
- [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.
- [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.
- [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.
- [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.
- [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.
- [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.
- [44] A. Jasra, C. C. Holmes, and D. A. Stephens. 2005. Markov Chain Monte Carlo Methods and the Label Switching Problem in Bayesian Mixture Modeling. *Statist. Sci.* 20, 1 (2005), 50–67.
- [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.
- [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.
- [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.
- [48] Pamela D Larsen. 2017. *Chronic Illness: Impact and Intervention* (10 ed.). Burlington: Jones & Bartlett Learning.
- [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 Guatterri, and Anna Maria Paganoni. 2020. Multivariate Hidden Markov Models for disease progression. *Statistical Analysis and Data Mining: The ASA Data Science Journal* 13, 5 (2020), 499–507.
- [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.
- [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.
- [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.
- [54] Yilmazcan Oezuyurt, 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).
- [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).
- [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]*
- [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.
- [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.
- [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.
- [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.
- [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.
- [63] Anselm L. Strauss and Juliet M. Corbin. 1988. *Unending work and care : managing chronic illness at home*. San Francisco, Calif. : Jossey-Bass.
- [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.
- [66] Claudio Terzano, Vittoria Conti, Fabio Di Stefano, Angelo Petrianni, 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.
- [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.
- [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.
- [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.
- [71] Eva van Weenen and Stefan Feuerriegel. 2020. Estimating Risk-Adjusted Hospital Performance. *IEEE Big Data* (2020).
- [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.
- [73] Concepción Violán, Sergio Fernández-Bertolín, Marina Guisado-Clavero, Quinti 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.
- [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.
- [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.
- [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.
- [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.
- [79] World Health Organization et al. 2016. Global report on diabetes. (2016).
- [80] Tong Xia, Yue Yu, Fengli Xu, Funing Sun, Diansheng Guo, Depeng Jin, and Yong Li. 2019. Understanding Urban Dynamics via State-Sharing Hidden Markov Models. In *The World Wide Web Conference* (San Francisco, CA, USA) (WWW '19). Association for Computing Machinery, New York, NY, USA, 3363–3369.
- [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.
- [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.
- [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.
- [84] Walter Zucchini, Iain I. MacDonald, and Roland Langrock. 2017. *Hidden Markov models for time series: an introduction using R*. CRC press.

A ADDITIONAL TABLES

Table 6: Transition matrix (with corresponding α).

Latent states <i>from</i>	<i>to</i>	1	2	3	4
		(1,1)	(1,2)	(2,1)	(2,2)
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

We follow common practice in Bayesian modeling [28] by performing the following model diagnostics. This is to ensure convergence of the MCMC algorithm and thus precise estimates. First, we inspected the effective sample size n_{eff} , indicating that the number of MCMC samples is sufficient. Second, we calculated the Gelman-Rubin convergence diagnostic $\hat{R}$ of all model parameters. The $\hat{R}$ is below the critical threshold of 1.1, suggesting convergence of the MCMC chain. Third, we retrieved traceplots (Fig. 4). The traceplots show the MCMC draws of the emission mean across the latent states and difference diseases (i. e., $\mu_s^{(A)}$ and $\mu_s^{(B)}$, $s \in \{1, 2\}$). The traceplots suggest that the chains have mixed well. Altogether, the previous diagnostics provide ample evidence that the MCMC algorithm has converged.

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

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

We further assessed our model by conducting posterior predictive checks following recommendations on Bayesian modeling [25, 28]. For this, new observations were simulated based on the posterior draws of the model parameters as well as the most probable latent states which were estimated using the Viterbi algorithm. The results are in Fig. 5. Based on it, we compare the credible intervals of the generated observations against the actual observations. The results suggest that our model fits the data well. We further checked whether Comorbidity-HMM is able to recover the specified parameters from simulated data. Here we obtain confirmatory results, suggesting that it can successfully capture the underlying dynamics.

(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).