Artificial Intelligence In Medicine 122 (2021) 102211

Contents lists available at ScienceDirect

Artificial Intelligence In Medicine

journal homepage: www.elsevier.com/locate/artmed

Interpreting clinical latent representations using autoencoders and probabilistic models☆

David Chushig-Muzo [a], Cristina Soguero-Ruiz [a], Pablo de Miguel-Bohoyo [b], Inmaculada Mora-Jiménez [a,*]

[a] Department of Signal Theory and Communications, Telematics and Computing Systems, Rey Juan Carlos University, Fuenlabrada 28943, Spain
[b] University Hospital of Fuenlabrada, Fuenlabrada 28943, Spain

ARTICLE INFO

Keywords:
Autoencoder
Learning latent representations
Gaussian mixture model
Clustering
Chronic diseases
Electronic health records

ABSTRACT

Electronic health records (EHRs) are a valuable data source that, in conjunction with deep learning (DL) methods, have provided important outcomes in different domains, contributing to supporting decision-making. Owing to the remarkable advancements achieved by DL-based models, autoencoders (AE) are becoming extensively used in health care. Nevertheless, AE-based models are based on nonlinear transformations, resulting in *black-box models* leading to a lack of interpretability, which is vital in the clinical setting. To obtain insights from AE latent representations, we propose a methodology by combining probabilistic models based on Gaussian mixture models and hierarchical clustering supported by Kullback-Leibler divergence. To validate the methodology from a clinical viewpoint, we used real-world data extracted from EHRs of the University Hospital of Fuenlabrada (Spain). Records were associated with healthy and chronic hypertensive and diabetic patients. Experimental outcomes showed that our approach can find groups of patients with similar health conditions by identifying patterns associated with diagnosis and drug codes. This work opens up promising opportunities for interpreting representations obtained by the AE-based model, bringing some light to the decision-making process made by clinical experts in daily practice.

1. Introduction

Currently, due to advances in medicine and technology, people live longer, with higher rates of individuals over 65 years. This aging phenomenon is particularly evident in developed countries, with an increasing number of elderly individuals more prone to develop chronic conditions such as diabetes and hypertension [1]. Individuals suffering from these diseases substantially increase the economic burden and health service demand due to hospital readmissions, adverse events and constant visits to primary and specialized care [2]. Beyond costs, the quality of life for patients with chronic conditions usually worsens, which can cause high rates of disability and mortality for acute cases [3]. Both health stakeholders and policy-makers are intensifying efforts to foster prevention and awareness, and searching for advanced approaches based on machine learning (ML) to provide an efficient method to identify chronic patients early and improve their care and treatment [4].

Among the state-of-the-art ML approaches, deep learning (DL) methods have achieved excellent outcomes across different domains in the last decade [5]. Inspired by the significant strides shown in other disciplines, several DL-based approaches have been conducted in clinical settings [6]. The rapid adoption of electronic health records (EHRs) has generated vast registers of data that offer great opportunities for clinical research [7]. In conjunction with EHRs, DL-based models have

Abbreviations: EHR, Electronic Health Record; KLD, Kullback Leibler divergence; AE, Autoencoder; DAE, Denoising Autoencoder; GMM, Gaussian Mixture Model; AHC, Agglomerative Hierarchical Clustering; CRG, Clinical Risk Group.

☆ This work has been partly supported by the Spanish National Research projects PID2019-106623RB-C41 (BigTheory), TEC2016–75361-R (Klinilycs), PID2019-107768RA-I00 (AAVis-BMR), DTS17/00158, by the Project ReCovid (Ref M2408) funded by Rey Juan Carlos University-Banco Santander, and by the Erasmus Mundus Program (Action 2 Partnerships, AESOP+ by Agreement No. 2014–0878/001–001). Found action by the Community of Madrid in the framework of the Multiannual Agreement with Rey Juan Carlos University in line of action 1, "Encouragement of Young Phd students investigation" Project Mapping-UCI (Ref F661).

* Corresponding author.
E-mail addresses: cd.chushig@alumnos.urjc.es (D. Chushig-Muzo), cristina.soguero@urjc.es (C. Soguero-Ruiz), pablo.miguel@salud.madrig.org (P. de Miguel-Bohoyo), inmaculada.mora@urjc.es (I. Mora-Jiménez).

https://doi.org/10.1016/j.artmed.2021.102211
Received 5 April 2021; Received in revised form 29 October 2021; Accepted 1 November 2021
Available online 9 November 2021

been extensively used for extracting knowledge and discovering patterns in a myriad of works, supporting the clinical decision-making [8,9]. However, EHR data can be heterogeneous, noisy, sparse and characterized by a high number of features (high-dimensional data), making it more difficult for knowledge discovery and the application of conventional data-driven methods on the original data space.

To circumvent these challenges, many dimensionality reduction techniques have been proposed [10]. These encompass different linear and nonlinear approaches transforming the original high-dimensional features into a new set of features, which capture relevant information from data [10]. Due to the success achieved by DL-based models, the use of the autoencoder (AE) [11] has been extended to the clinical setting [6]. AE relies on a nonlinear transformation of the original data space into a lower-dimensional space, known as latent representation (LR) space. Several studies have demonstrated that LRs can capture relevant clinical insights from raw data [12,13], with the denoising autoencoder (DAE) being effective in creating a robust a lower-dimensional space [14].

Despite the undeniable success of DL approaches, nonlinear transformations create models with a lack of interpretability (referred to as *black-box models*) that cause important shortcomings in extracting knowledge, particularly in the clinical domain [15]. Interpretation is defined as the process of generating human-understandable explanations on outcomes provided by computational models [16]. There is a growing concern, especially in the DL research community, about the interpretation of models beyond results, focusing efforts on the development of *interpretable models* [15]. Consequently, the number of papers related to interpretability has rapidly increased in recent years [17,18,19,20]. Model interpretability methods are classified according to several criteria, being one of the most accepted divisions into two categories: intrinsic and *post hoc* methods. In intrinsic methods, interpretability is present in the inner structure of the models, namely, self-explanatory models. Some examples are tree-based models and generalized linear models [21]. In the *post hoc* methods, the strategies of interpretation are the next step to training, extracting knowledge from trained models through an auxiliary method [22]. Although several works for model interpretability have been proposed in the literature, the majority of them aim to understand how models make predictions (supervised learning methods) [23,24], and few studies have interpreted unsupervised learning methods such as AEs.

To interpret unsupervised approaches, four research lines can be distinguished [25,26,27]: *(i)* to select explanatory or saliency features, *(ii)* to build disentangled representations, *(iii)* to visualize data in a lower-dimensional space by finding similar samples by proximity and *(iv)* to find corresponding relationships between raw and transformed features. For explanatory features, methods in the literature aim to find specific features that capture relevant properties for a specific problem [25]. Recent advances have sought to improve features generated during the unsupervised learning process, building the called disentangled representations that gather an intrinsic meaning [26] (*e.g.,* with face pictures, extracting hair color, nose size or others). In clinical settings, several works have looked for interpretation with the projection of data into a lower-dimensional space (commonly two or three dimensions), aiming to find patterns by the proximity of projected patients [12,13]. Remark that in health care, model interpretability acquires great importance, since a lack of interpretability may hamper decision-making [6]. It is paramount to develop a reasonable comprehension of the specific clinical problem and to gain valuable insights from analyzed data.

This paper presents a methodology based on probabilistic methods for interpreting AE-based models, with application in clinical settings. This approach aims to find hidden patterns to distinguish groups of patients with common clinical conditions. Our work lies within the context of *post hoc* interpretability methods by applying probabilistic approaches, in particular Gaussian mixture models (GMM), to gain insights into the characterization of the patient's health status.

The rest of the paper is structured as follows. A data description and the proposed methodology for interpreting AE-based models are presented in Section 2. Experimental results are detailed in Section 3. Finally, the discussion and conclusions are presented in Sections 4 and 5, respectively.

2. Materials and methods

2.1. Data collection and preprocessing

Data for this study were extracted from EHRs collected by the University Hospital of Fuenlabrada (UHF). UHF is a public hospital located in Madrid (Spain) that provides medical services to six health centers and encompasses approximately 225,000 inhabitants. We focus on chronic patients due to their importance in developed societies with aging populations. Diabetes and hypertension are among the most prevalent chronic conditions, especially among elderly people. It is estimated that 75% of people will suffer from chronic diseases worldwide [1].

To determine the patient health status, we considered a population classification system named clinical risk groups (CRGs), version 1.8 [28]. CRGs are oriented toward chronic conditions and have been internationally validated by the health care community [29,30]. CRGs consider patient encounters with the health system at a limited time period (usually one year) and assign patients to one group according to their health status. CRGs consider both demographic data (age and sex) and clinical codes (with the corresponding timestamp), including diagnoses, procedures and medical drugs. There are nine core health status groups, which are hierarchically organized from 1 (healthy patients) to 9 (the worst health status). Intermediate values correspond to acute diseases, a single chronic condition, or several chronic conditions, providing more than one thousand groups or health statuses.

Health statuses are identified by five numbers: the first indicates the core health group; the next three numbers in the middle identify the chronic condition number, and the four previous digits refer to the CRGs; the last digit refers to the severity level. In this work, we only considered the first four digits due to the limitation in the number of patients in some CRGs. We focused on three CRGs, namely, CRG-1000 (healthy), CRG-5192 (hypertensive) and CRG-5424 (diabetic), with 46,835, 12,447, and 2166 patients for one natural year, respectively. Since the number of patients associated with each CRG is highly imbalanced and ML algorithms may be affected by the class imbalance learning problem [31], we followed a random sampling strategy as in [32,33]. Thus, the number of patients considered per CRG is limited by CRG-5424 (2166 patients).

In this work, EHRs contain diagnoses from primary and specialized care, as well as pharmaceutical drug dispensation. Diagnoses were coded according to the International Classification of Diseases-9th Revision Clinical Modification (ICD9-CM) [34]. Drugs were coded according to the Anatomical Therapeutic Chemical (ATC) Classification System [35]. Both ICD9-CM and ATC codes consist of a determined number of alpha-numeric characters organized in a hierarchical manner. Some ICD9-CM codes directly related to the chronic conditions analyzed in this paper are codes 250.01 (diabetes mellitus without mention of complications, type I) and 401.1 (benign essential hypertension). Some ATC codes included in drug therapy for chronic conditions are A10BA02 (metformin) and C09AA02 (enalapril). Since each code can be considered a potential feature, the dimensionality of the feature space can notably increase, reaching 13,000 potential codes for diagnoses and 3430 for drugs.

To circumvent the curse of the dimensionality problem [36], we proposed in previous works [32,33] to limit the number of features by reducing the detail of the aforementioned codes. For ICD9-CM codes, we only worked with the category code (three alpha-numeric elements), while for ATC codes, we retained up to the chemical level (five alpha-numeric elements). This approach resulted in a total of 2263 features,

with 1517 and 746 features for diagnoses and drugs, respectively. Therefore, we propose to represent the encounters of a patient by a binary vector of 2263 features (see Panel (a) in Fig. 2). A value of '1' in a feature means that the corresponding code is present in at least one patient EHR, while a value of '0' means that the code was not registered.

To provide a graphical characterization of a specific group of patients, we compute the profiles associated with diagnosis/drug codes. In previous works [32,37], we defined the profile as a representation where the x-axis shows ICD9-CM and ATC codes and the y-axis represents the presence rate of each code when considering patients associated with a group. For each group of patients (assume the group is composed of N patients), the group profile is obtained from the concatenation of D elements (codes of diagnoses and drugs). For the d-th element in the profile (*e.g.,* assume it is associated with the hypertension diagnosis), its value is computed as the proportion of N patients presenting the particular diagnosis code of hypertension (ICD9-CM '401'). Thus, profiles with elements close to 1 indicate that almost all patients in that group have the same code (either diagnosis or drug). This profile provides an overall vision of more representative diagnoses and drugs linked to one group of patients. Taking CRG-5424 as an example, its diagnosis profile (see green line in Fig. 1 (a)) shows the prevalence of ICD9-CM '250' (diabetes mellitus) with a value over 0.8, indicating that approximately 80% of patients in CRG-5424 present this diagnosis code in the EHR. Concerning the drug profile, Fig. 1 (b) shows the drugs most dispensed, corresponding to ATC codes: 'A10BA' (biguanides) and 'C10AA' (HMG CoA reductase inhibitors). The first is frequently prescribed for diabetes treatment, and the second one is utilized for lowering cholesterol. Regarding CRG-5192, ICD9-CM '401' (essential hypertension) stands out among other codes, being the most important diagnostic code used (with a presence rate of 0.8) for hypertension (see orange line in Fig. 1). In terms of drugs, the most prevalent drugs for CRG-5192 were ATC 'C09AA' (ACE inhibitors) and 'C10AA' (HMG CoA reductase inhibitors), drugs commonly used in hypertension treatment.

2.2. Proposed methodology for interpreting AE-based models

In this section, we introduce the proposed methodology for interpreting AE-based models. From a dataset $\mathscr{D} = \{\mathbf{x}_i\}_{i=1}^{n}$ consisting of n patients, the i–th patient is represented by a vector of D random variables $\mathbf{x} = [x_1, \ldots, x_D] \in \mathbb{R}^D$. By using DAE, we transform an input vector $\mathbf{x}$ into a new vector $\mathbf{h} = [h_1, \ldots, h_d] \in \mathbb{R}^d$ composed of d latent dimensions (LDs), often called latent representation (LR). Note that the dimensionality of the LR is lower than that of the original dataset ($D > d$).

Our work was carried out by following a methodology consisting of two stages (a schematic is given in Fig. 2). In the first stage, we estimated the probability density function (PDF) for each LD by applying both kernel density estimation (KDE) [38] and GMM [36]. Through the use of the Gaussian components of the GMM on each LD, we identified what we called regions of interest (ROIs). In the second stage, clustering of similar LDs was conducted by using Kullback-Leibler divergence (KLD) as a dissimilarity measure. We used the ROIs of each LD belonging to the same cluster to identify hidden patterns and characterize complex clinical conditions.

2.3. Autoencoders and latent dimensions

The simplest AE is a three-layer (input, intermediate, output) network with an encoder, latent space and a decoder (see Panel (b) in Fig. 2). The encoder transforms an input vector $\mathbf{x}$ into a dimensionality-reduced representation $\mathbf{h} \in \mathbb{R}^d$ guided by the mapping $\mathbf{h} = f(\mathbf{x}) = f(\mathbf{W}\mathbf{x}^T + \mathbf{b})$, with dimensionality $d < D$. The parameter $\mathbf{W}$ represents a weight matrix $d \times D$, and $\mathbf{b}$ corresponds to a bias vector. The decoder reconstructs an approximated representation of $\mathbf{x}$ guided by $\hat{\mathbf{x}} = g(\mathbf{h}) = g(\mathbf{W}'\mathbf{h}^T + \mathbf{b}')$. $\mathbf{W}'$ is a weight matrix $D \times d$ and $\mathbf{b}'$ is a bias vector associated with the decoder. Functions $f(\cdot)$ and $g(\cdot)$ are nonlinear activation functions (commonly a sigmoid or hyperbolic function). The AE is trained to find optimal parameters $\theta = \{\mathbf{W}, \mathbf{b}, \mathbf{W}', \mathbf{b}'\}$ by minimizing a loss function.

Many AE variants have been proposed [14,39]. For instance, DAE was proposed to generate more robust LRs [14]. The central idea behind the DAE is to corrupt a fraction (FN) of input data by considering stochastic noise, with zero-masking noise (ZMN) and salt-and-pepper noise (SPN) being the most frequently used. Both ZMN and SPN randomly change a fraction of input values to their minimum value in ZMN and the maximum/minimum value (according to a fair coin flip) in SPN. Note that for binary data, minimum and maximum values correspond to 0 and 1, respectively.

2.4. Density estimation techniques for finding regions of interest

The PDF associated with each LD is estimated by considering probabilistic density-based models, specifically KDE and GMM. KDE, also known as the Parzen-Rosenblatt window method [38], is a nonparametric technique to estimate the PDF of a continuous random variable l as:

$$p\left(l\right) = \frac{1}{nh}\sum_{i=1}^{n} K_h(l, l_i) \tag{1}$$

where $K_h(\cdot)$ is a local kernel function and h is the corresponding width.

A Gaussian kernel function defined as $K_h\left(u, v\right) = e^{-\frac{1}{h^2}\left\|u-v\right\|^2}$ is

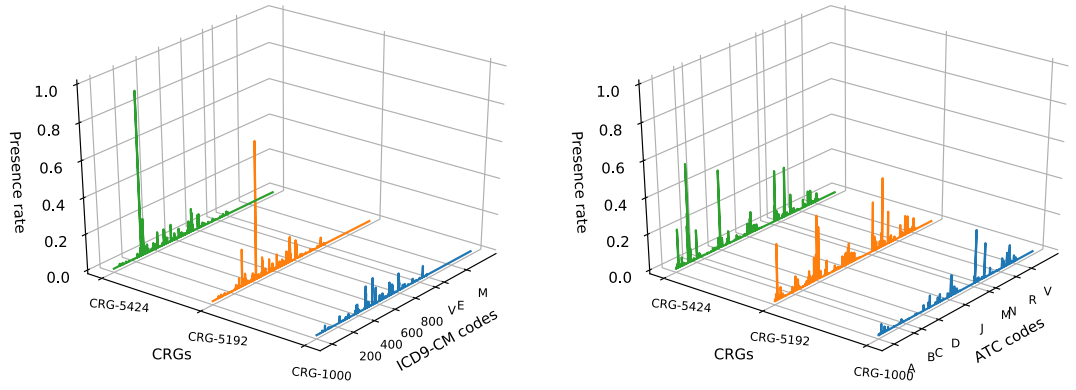

(a) (b)

Fig. 1. Profiles associated with CRG-1000, CRG-5192, and CRG-5424: (a) diagnosis profiles; (b) drug profiles.

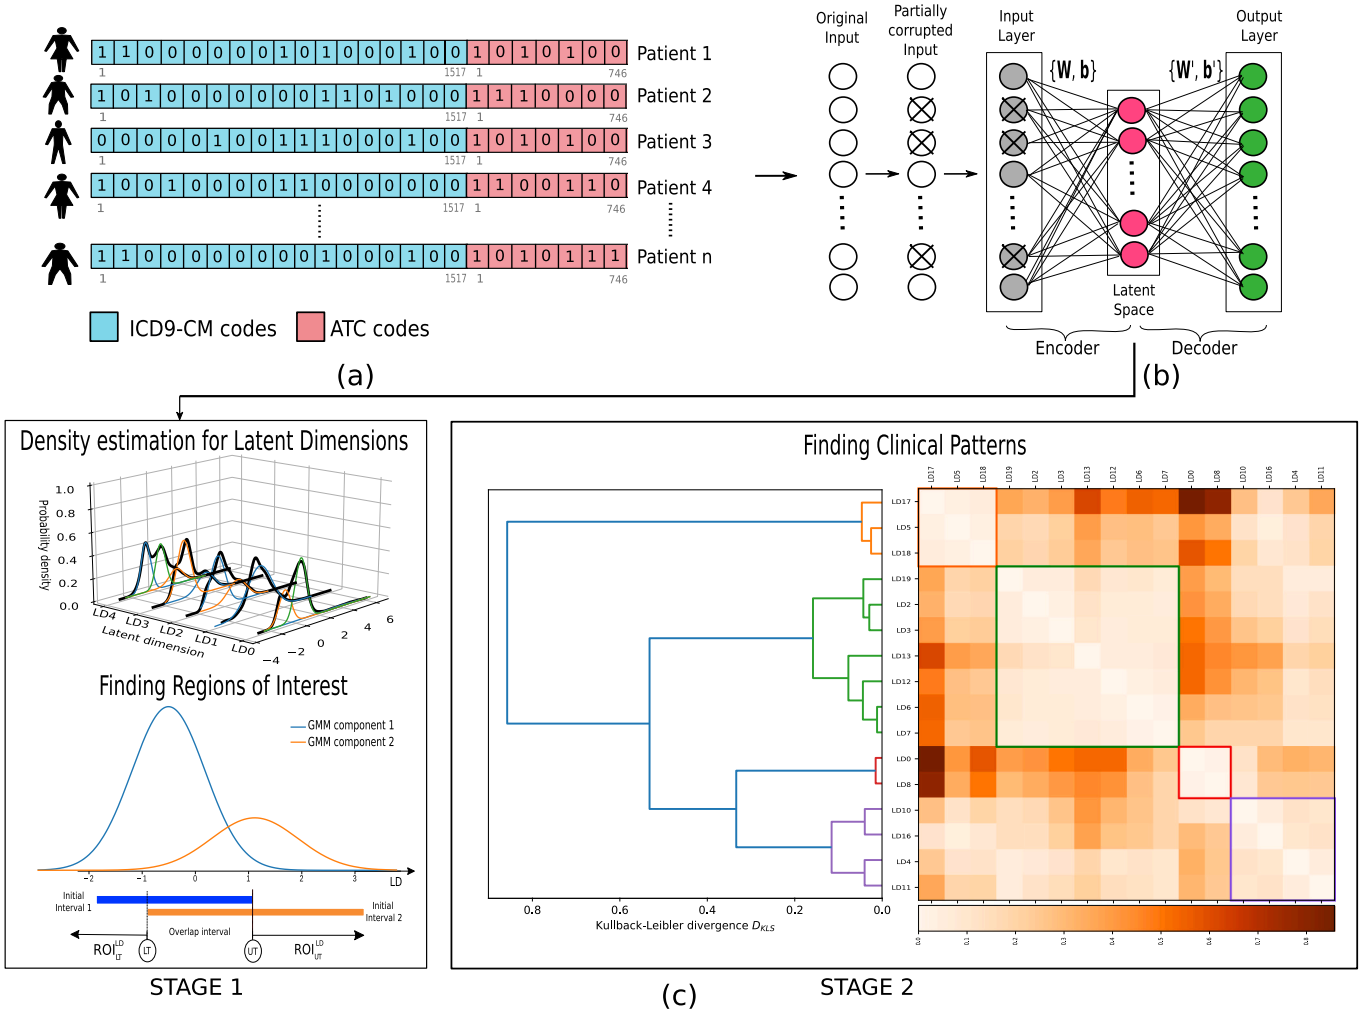

Fig. 2. Workflow of the proposed approach: (a) representation of n patient vectors; (b) schematic of DAE; (c) two-stage methodology for interpreting clinical latent representations.

considered. According to [40], the choice of h is significant to control the smoothness of the PDF estimation. Large values of h provide over-smoothed PDF estimations, while small values can generate functions very adapted to the particular samples used in PDF estimation. To select the optimal width, several methods have been proposed in the literature, considering Silverman's method herein [41].

A GMM is a parametric probabilistic method for estimating the PDF, defined as a weighted sum of M Gaussian distributions, called GMM components [36]:

$$p\left(x\right) = \sum_{m=1}^{M} \omega_i \phi_i\left(x, \mu_i, \sigma_i^2\right), \quad \text{with} \quad \sum_{i=1}^{M} \omega_i = 1 \tag{2}$$

where $\omega_i \in [0,1]$ is the weight (known as prior probability) of the i–th Gaussian distribution, indicating the relative contribution of each component to the mixture model. $\phi(\cdot)$ is a Gaussian function with parameters mean μ_i and variance σ_i^2. These parameters are estimated by optimizing the likelihood function in Eq. (2) through the iterative expectation-maximization algorithm [42]. To determine the appropriate number of M Gaussian components, we propose the use of both the Bayesian inference criterion (BIC) [43] and the silhouette coefficient index (SCI) [44]. The most appropriate value of M is selected according to the minimum value of the BIC and the maximum value of the SCI.

In this paper, GMM components are used to obtain the ROIs associated with the considered LDs. In the simplest scenario, an ROI is an interval in the LD domain that potentially allows us to identify groups of patients with a specific health condition. The procedure to compute the ROI was inspired by [45] and follows the following steps: *(i)* setting the initial interval for each GMM component as the range of values encompassing the majority of the area, defined as $[\mu_i - 3\sigma_i, \mu_i + 3\sigma_i]$; *(ii)* establishing the overlap interval of the initial intervals; *(iii)* setting the lower end of the overlap interval as the lower threshold (LT) and the upper end as the upper threshold (UT); and *(iv)* establishing $\text{ROI}_{\text{LT}}^{\text{LD}}$ for

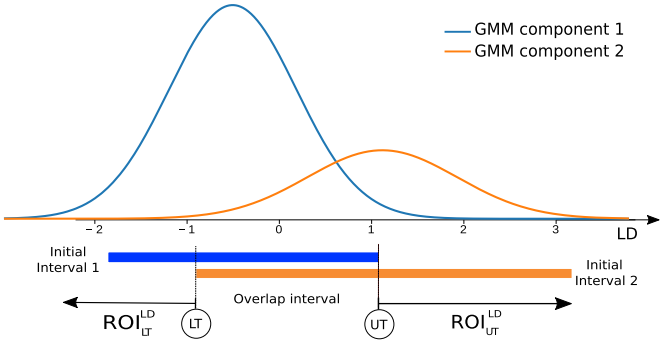

Fig. 3. Components of the GMM for a specific LD. By discarding the overlap interval of both components, two ROIs are defined from thresholds LT and UT: $\text{ROI}_{\text{LT}}^{\text{LD}}$ for LD values lower than LT and $\text{ROI}_{\text{UT}}^{\text{LD}}$ for LD values greater than UT.

LD values lower than LT and setting ROI_{UT}^{LD} for LD values greater than UT. An illustrative example is shown in Fig. 3 for an LD with two GMM components: 1 (blue) and 2 (orange), representing their initial intervals in the same color. The ROIs are identified as the part of the initial intervals disjoint with the overlap interval, and thresholds (LT, UT) are depicted in a circle.

2.5. Identifying similarity between latent dimensions

Methods to group samples based on similarity criteria have been successfully used in different scientific disciplines, including health care [46]. Among these methods, agglomerative hierarchical clustering (AHC) [47] is notably considered because *(i)* no prior knowledge of cluster centers is required and *(ii)* the clustering result is presented as a tree-like structure (dendrogram), hence contributing to the interpretation of the underlying structure in the data. Generally, the AHC algorithm is based on a similarity/dissimilarity measure. In this paper, we consider the KLD to quantify the dissimilarity D_{KL} between the PDF p and q associated with two LDs [48]. Note that $D_{KL} \approx 0$ when both distributions are similar. To perform clustering techniques adequately, symmetric measures are required. Therefore, we compute the symmetric KLD defined as $D_{KLS} = D_{KL}(p \parallel q) + D_{KL}(q \parallel p)$. We consider it in the AHC by using several linkage criteria (*single, complete and average*). The AHC method was applied to group similar LDs into clusters on the assumption that these clusters can be clinically relevant.

We proceeded next to use the resulting clusters of LDs and corresponding ROIs to identify patients with similar clinical characteristics. The ROIs may serve as a base for identifying patients and characterizing populations with specific clinical conditions. To achieve this, we considered the parallel coordinates plot (PCP) [49], a tool that allows us to visualize multidimensional variables and to find trends and patterns. With the ROIs established, groups of patients were identified by using PCP. The clinical patterns associated with each group were obtained by computing the presence rate of every ICD9-CM/ATC code, identifying the most representative ones.

It may be noted that each LD has two associated ROIs when $M>1$ (ROI_{LT}^{LD}, ROI_{UT}^{LD}), and each cluster is usually composed of a set of LDs. When the clinical complexity increases, a single LD might not be sufficient to identify patients with similar clinical characteristics. Then, to identify patients with a specific health condition, we combined ROIs of the LDs belonging to the same cluster. For selecting the ROIs, we consider all combinations of ROI_{LT}^{LD}, ROI_{UT}^{LD} linked to LDs and identify such a combination that allows us to gather a representative number of patients for each group.

3. Experimental results

This section presents the experimental setup, the results linked to every stage of the methodology sketched in Fig. 2, and the characterization of the groups of patients found.

3.1. Experimental setup

This subsection describes the data used for training and validation and the hyperparameter setup considered for training DAEs. In this work, AE-based models were implemented with Keras and Theano (as the backend), and both Python libraries were extensively used by the ML community. The binary cross-entropy was chosen as the cost function, which was optimized by applying a minibatch gradient descent procedure with early stopping and adaptive learning rates [50]. As activation functions, we considered a rectified linear activation function named LeakyReLU [51] for the neurons in the intermediate layer and a sigmoid function for the output layer. Several neurons {5,10,20,30} (*i.e.,* the number of LDs) for the intermediate layer were explored in DAEs, as well as several types of noise (ZMN and SPN), considering a different corruption level {10,50,100}. Empirical analysis showed that SPN with a corruption level of 10 and 20 latent dimensions was the most appropriate network configuration to extract patterns of chronic diseases and to establish groups with specific clinical characteristics.

As described in subsection 2.1, CRGs present remarkable differences in the number of patients, emerging the class imbalance problem. To address this issue, we considered an undersampling approach, which modifies the class distribution of majority classes, taking as reference the number of patients in the minority class [31]. In our case, CRG-5424 corresponded to the minority class (2166 patients). Considering CRG-1000, CRG-5192 and CRG-5424, a dataset with 6498 patients was obtained. This resulting dataset was split into a training subset (4873 patients) and an evaluation subset (1625 patients). The training subset was used to design the DAEs as proposed in Section 2.3, whereas the evaluation subset was considered for determining the consistency of results provided by applying our methodology.

3.2. Density estimation and clustering of latent dimensions

In this subsection, we estimate the PDF and find the ROIs associated with each LD. Then, we perform the clustering of LDs using the AHC method with the KLD as the similarity measure aiming to group similar LDs. To identify the LDs, we named them using identifiers from 0 to 19 in the first column of Table 1. For the PDF estimation, we used the KDE and GMM techniques and compared both with the D_{KLS}, showing that the D_{KLS} values are close to zero (see the second column of Table 1). Since both methods provide similar density estimations, we henceforth used the GMM. Fig. 4 shows, for each LD, both the PDF estimation using GMM (in black) and the corresponding components (in different colors). To select the number M of GMM components (see the third column in Table 1), we used the minimum BIC and the maximum SCI. Our hypothesis is that LDs with $M>1$ may be adequate for characterizing groups of patients with a specific clinical condition. Thus, LDs that did not fulfill this condition (LD1, LD9, LD14, LD15) were not considered in the next stages. With the GMM distributions associated with each LD, we obtained the ROIs according to the method detailed in Section 2.2. These thresholds are shown in the fourth column of Table 1.

As stated in subsection 2.5, with a single LD, it might not be straightforward to identify patients with similar clinical characteristics. In these cases, it might be adequate to group LDs to obtain an appropriate characterization of different kinds of clinical conditions. For this purpose, similar LDs are grouped into clusters through the AHC method

Table 1
Details of PDF estimation and thresholds for each LD using GMM: (First column) LD identifier; (second column) KLD symmetric D_{KLS}; (third column) number of GMM components; and (fourth column) thresholds {LT, UT}, where NA stands for "not applicable".

LD identifier	D_{KLS}	# GMM components	{LT; UT}
0	0.0235	3	{−1.05; 1.78}
1	0.0136	1	NA
2	0.0203	2	{−1.20; 1.26}
3	0.0390	2	{−0.82; 1.12}
4	0.0290	3	{−0.92; 1.43}
5	0.0286	2	{−1.16; 0.52}
6	0.0271	2	{−1.59; 1.16}
7	0.0356	2	{−1.99; 1.26}
8	0.0343	4	{−1.39; 1.08}
9	0.0413	1	NA
10	0.0186	3	{−1.10; 1.62}
11	0.0098	3	{−0.99; 1.48}
12	0.0158	2	{−1.16; 0.92}
13	0.0194	2	{−1.21; 1.58}
14	0.0214	1	NA
15	0.0068	1	NA
16	0.0245	2	{−1.89; −0.12}
17	0.0272	2	{−0.89; 0.24}
18	0.0184	2	{−0.90; 0.62}
19	0.0181	2	{−1.19; 1.40}

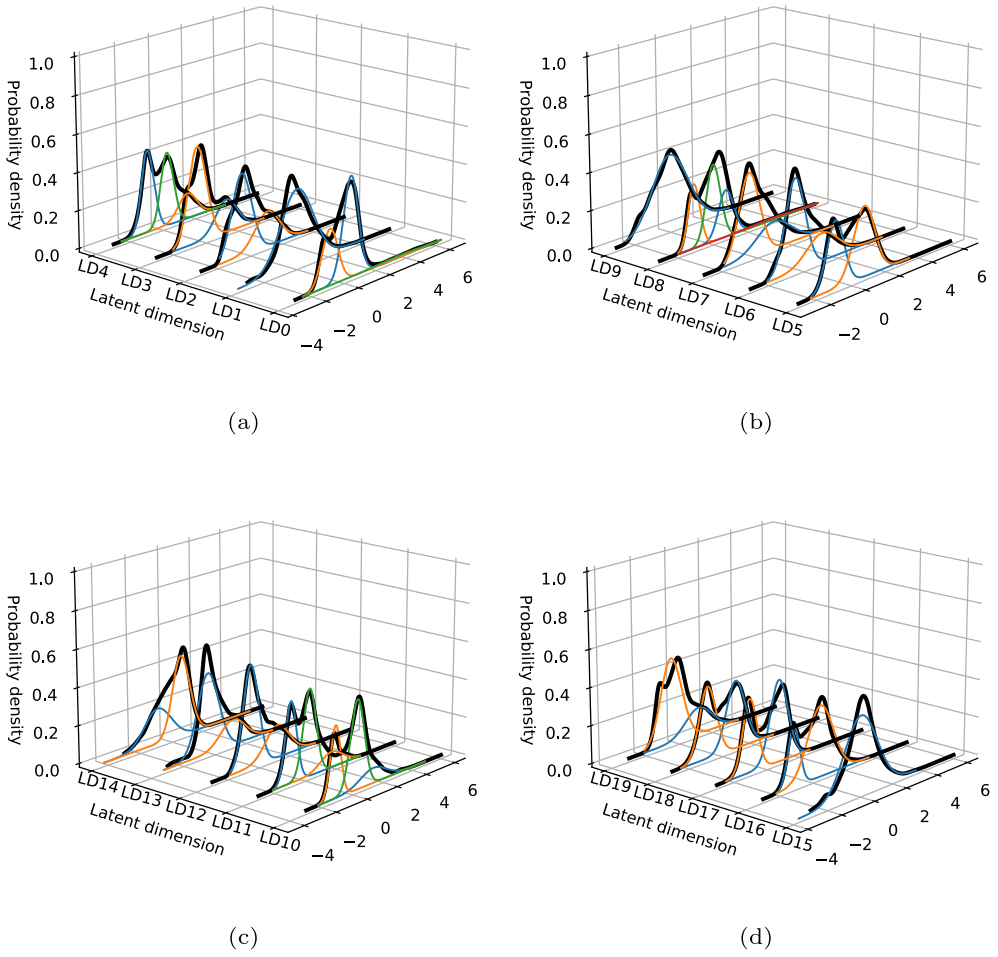

Fig. 4. The PDF and Gaussian distributions of GMM for each LD. For a better display of LDs, we split them into four plots where each plot contains five LDs identified as (a) 0–4; (b) 5–9; (c) 10–14; and (d) 15–19.

by using D_{KLS} as the similarity measurement. Three linkage criteria were tested (simple, complete, average), with *complete* linkage providing reasonable results. The number of clusters was selected by using the SCI [52], which provided K=4 as the number of clusters. The clustering configuration result using the AHC with complete linkage is shown in Fig. 5 (a), $C_{D1} = \{LD5, LD17, LD18\}$, $C_{D2} = \{LD2, LD3, LD6, LD7, LD12, LD13, LD19\}$, $C_{D3} = \{LD0, LD8\}$, $C_{D4} = \{LD4, LD10, LD11, LD16\}$. In addition, we show in Fig. 5 (b) a heatmap representing the D_{KLS} values between every pair of LDs. Lighter colors are associated with more similar LDs.

Once clusters of LDs were found, we used the PCP to visually identify groups of patients by establishing the ROIs for each LD belonging to a cluster. In PCP, each variable was represented by a vertical axis, and values for a particular sample (patient) were linked by a line connecting the specific value of each variable. An example is provided in Fig. 6, where Panel (a) shows all patients belonging to CRG-1000 (blue), CRG-5192 (green), and CRG-5424 (red), while panels (b), (c) and (d) show the result of progressively selecting patients by considering just one ROI per LD in the cluster C_{D1}. In particular, Fig. 6 (b) shows the selected patients when using one specific ROI in LD5 (named $\text{ROI}_{\text{UT}}^{\text{LD5}}$). Note that it is not feasible to distinguish patients with a specific clinical condition (mixture of patients from CRG-1000, CRG-5192, CRG-5424). By using one of the ROIs associated with LD17 (see Fig. 6 (c), $\text{ROI}_{\text{UT}}^{\text{LD17}}$), we identified patients belonging to CRG-1000 and some of CRG-5192. Finally, by considering one of the ROIs of LD18 (named $\text{ROI}_{\text{UT}}^{\text{LD18}}$ in Fig. 6 (d)), the identification of patients belonging to CRG-1000 was better achieved. Note that for each cluster of LDs, ROI identification led to the characterization of patients with similar clinical conditions.

3.3. Analysis and characterization of groups of patients

In this subsection, we characterize the groups of patients found (using clinical data) by considering the presence rate of the diagnosis and drug profiles. As argued, each LD has associated two ROIs (identified by $\text{ROI}_{\text{LT}}^{\text{LD}}$, $\text{ROI}_{\text{UT}}^{\text{LD}}$) when M>1, and each cluster is composed of a set of LDs. Thus, since we identified four clusters of LDs in subsection 3.2, the number of groups of patients who could appear, in our case, was eight. We checked that the relationship in terms of the diagnosis and drug profiles was high for some pairs of groups. Therefore, to simplify the analysis, we focus here on those groups with the most significant clinical findings. The profiles of these groups (henceforth identified by $\{G_i\}_{i=0}^{4}$) are presented in Fig. 7. Note that one profile for diagnoses (ICD9-CM codes, Fig. 7 (a)) and another profile for drugs (ATC codes, Fig. 7 (b)) for each group of patients are depicted. The five highest values in the profile of diagnoses and drugs per group when considering the training subset are presented in Table 2, whereas those corresponding to the evaluation subset are shown in Table 3.

By inspecting the corresponding profiles of Groups G0 and G1, the clinical characteristics shown as prevalent diagnosis code the ICD9-CM '250' (see two first columns in Table 2) and significant differences in the drug consumption. Patients in Group G0 presented a high consumption of different types of insulin, especially 'A10AB' (fast-acting action) and 'A10AE' (long-acting action). These drugs are commonly recommended

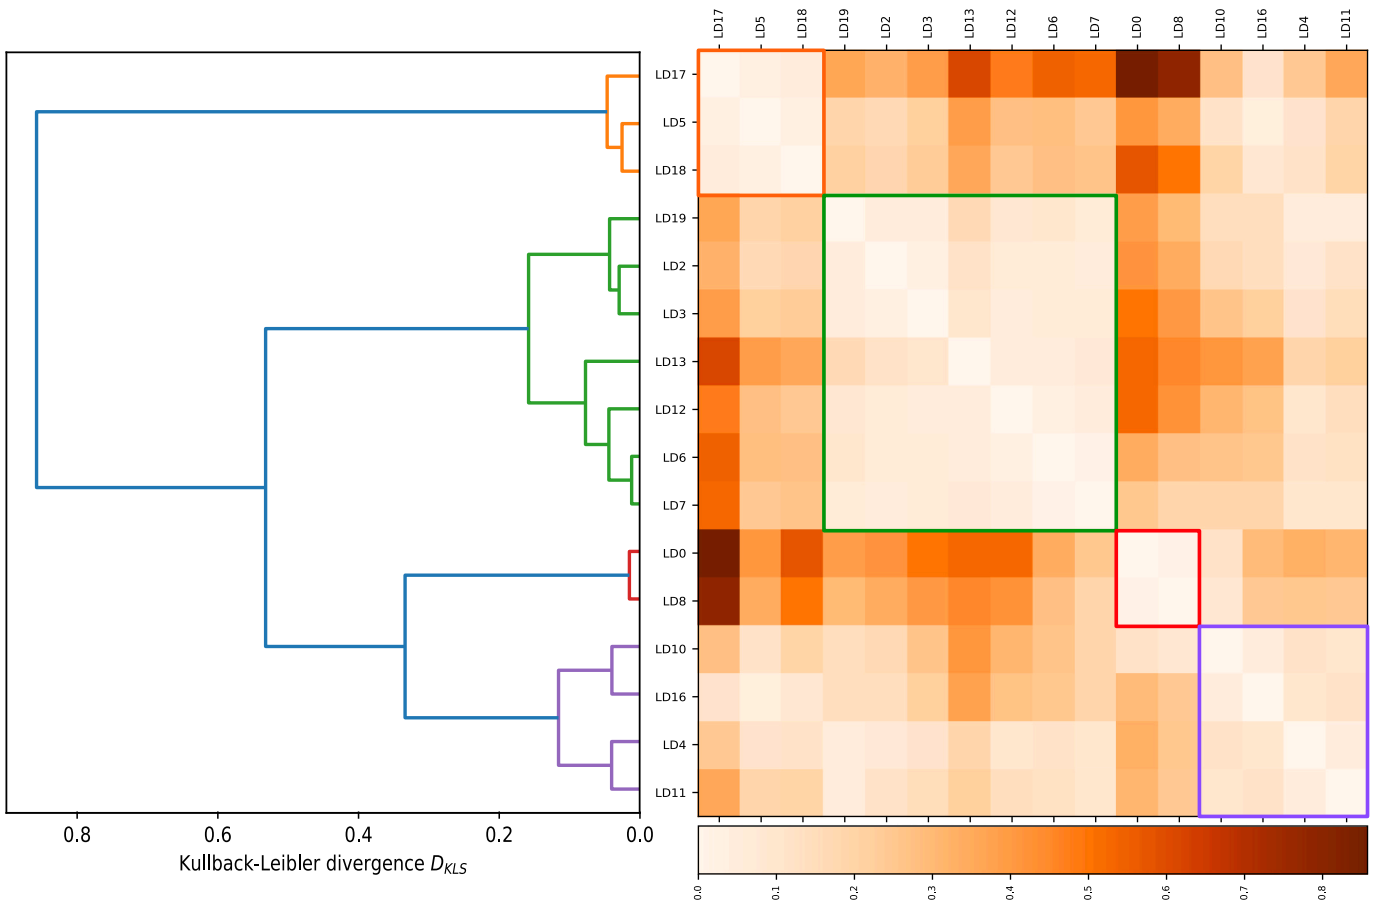

Fig. 5. (a) AHC dendrogram using complete linkage on LDs. (b) Heatmap of D_{KLS} between every pair of LDs.

for treating type 1 diabetes [53]. In addition, in its diagnosis profile, other diagnosis codes with values under 0.1 appeared (ICD9-CM '465', '526', '692'), neither of which is directly linked to diabetes. Patients associated with G1 presented a prevalent consumption of ATC 'A10BA' (biguanides) and 'C10AA' (statins). Biguanides are oral anti-hyperglycemic drugs that aim to improve glucose tolerance and are the first line of pharmaceutical treatment used for type 2 diabetes. ATC 'C10AA' is used to lower serum cholesterol, and it is widely used to prevent cardiovascular disease [54]. In addition, patients of G1 also consumed ATC 'B01AC' (platelet aggregation inhibitors), which is frequently prescribed for people over 65, and used for preventing cardiovascular complications [55]. Notably, Group G1 was the most similar in terms of the diagnosis/drug profile to CRG-5424 (see profiles in Fig. 1). We may characterize these two groups based on their most predominant diagnosis and drug codes as follows: G0 as insulin-dependent diabetic patients and G1 as mostly noninsulin-dependent diabetic patients.

By continuing the analysis, the diagnosis profile of Group G2 showed that the ICD9-CM code with the highest presence rate was '401' (hypertension). Other frequent codes were '272', '300', '526' and '719'. Regarding the drug profile, patients belonging to Group G2 were characterized by consumption of the following ATC codes: 'C10AA' (0.65) and 'C09AA' (0.63). It is interesting to remark that the ATC code 'C09AA' (angiotensin-converting enzyme (ACE) inhibitors) is one of the most commonly used treatments for hypertension and is also prescribed for various cardiovascular and renal diseases [56]. Additionally, in the drug profile, ATC 'N05BA' appeared to be one the most frequent. Note that anxiolytics are used to treat sleep disorders and anxiety disorders [57], showing a link with the ICD9-CM '300' (anxiety disorders).

Group G3 includes pregnant women with ICD9-CM '401' and 'V27' (outcome delivery) as the most representative diagnosis codes. Moreover, the diagnosis profile of G3 shows other codes related to pregnancy (see the fourth column in Table 2). In particular, the ICD9-CM code '648' was quite frequent in this group, which indicates that pregnant women suffer from complications during pregnancy such as gestational diabetes. The drug profile of G3 showed the consumption of drugs related to pregnancy 'B03AA' (iron bivalent), 'H03CA' (iodine therapy) and analgesics 'N02BB' and 'N02BE'. This insight is in accordance with the clinical literature [58]. Unlike the drug profile associated with the hypertensive population, ATC 'C09AA' or other drugs used for hypertension treatment are not frequent in G3 since they are not recommended in pregnancy. For instance, ACE inhibitors and angiotensin receptor blockers are contraindicated because of the adverse fetal effects reported [59].

Finally, the most common codes associated with G4 are shown in the last column of Table 2, with very low rates in comparison with previous groups. For drugs, we observed that the most common ATC codes found were: 'N02BE' (anilides) and 'M01AE' (propionic acid derivatives), which are analgesics. The exploration of profiles did not show great differences with clinical codes of CRG-1000. Thus, Group G4 could be characterizing a healthy population.

We proceeded next to validate the generalization capability of interpreting the clinical latent representations by analyzing the profiles when considering test data. The profiles of the groups of test patients are shown in Fig. 8, with the five most frequent codes presented in Table 3. As it can be visually checked, training and test profiles are quite similar in terms of diagnosis and drugs. A quantitative measure of the similarity between pairs of training/test profiles in the same group reveals Pearson coefficient values greater than 0.9 for the four groups, supporting visual intuition.

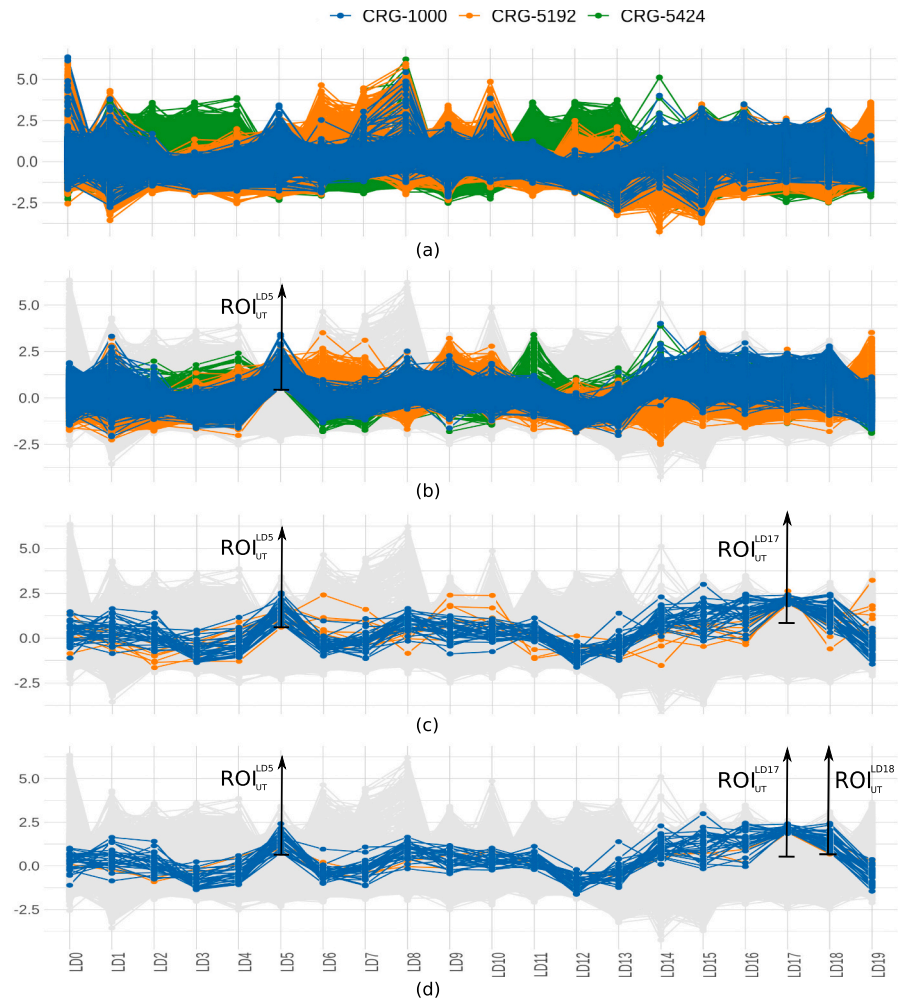

Fig. 6. Process of considering ROIs (black arrows) in several LDs and resulting PCP: (a) patients belonging to CRG-1000, CRG-5192, CRG5424; (b) patients selected using ROI_{UT}^{LD5}; (c) patients selected with the previous ROI and ROI_{UT}^{LD17}; (d) patients selected using the two previous ROIs and ROI_{UT}^{LD18}.

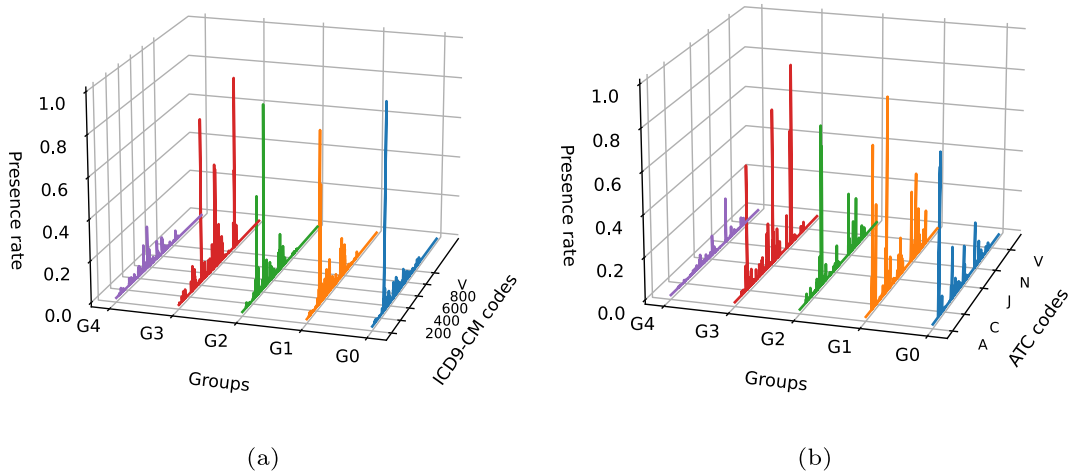

Fig. 7. Group profiles obtained when considering training patients: (a) diagnosis profiles, and (b) drug profiles.

4. Discussion

In this paper, the two-stage methodology proposed for interpreting clinical LRs obtained by AE-based models showed great potential to discover hidden patterns and groups of patients in an unsupervised way. In the first stage, the PDF estimation of each LD using GMM played a key role in characterizing patients with specific clinical characteristics. This PDF estimation allowed us to visualize how the values of an LD are distributed, leading to the identification of ROIs, which are intervals within the domain of an LD where patterns of a specific condition are

Table 2
Description of the ICD9-CM and ATC codes with the five highest values in the profile of diagnoses and drugs per group (training subset).

Code	Description	G0	G1	G2	G3	G4
'250'	Diabetes mellitus	0.98	0.83			
'272'	Disorders of lipoid metabolism	0.10	0.57	0.48		
'300'	Anxiety, dissociative disorders			0.12		
'401'	Essential hypertension			0.91	0.79	
'460'	Acute nasopharyngitis		0.15			
'462'	Acute pharyngitis					0.12
'465'	Acute upper respiratory infections	0.07	0.15			
'526'	Diseases of the jaws	0.08		0.17		0.20
'558'	Unspecified gastroenteritis and colitis					0.08
'648'	Other current complication in pregnancy				0.44	
'650'	Normal delivery				0.50	
'664'	Trauma to vulva				0.41	
'692'	Contact dermatitis	0.07				
'719'	Other and unspecified disorders of joint		0.18	0.09		0.05
'780'	General symptoms					0.08
'V27'	Outcome of delivery				0.84	
'A10AB'	Insulins for injection fast-acting	0.75				
'A10AD'	Insulins for injection intermediate-acting	0.19				
'A10AE'	Insulins for injection long-acting	0.70				
'A10BA'	Biguanides		0.75			
'B01AC'	Platelet aggregation inhibitors		0.46			
'B03AA'	Iron bivalent, oral preparations				0.58	
'C07AB'	Beta blocking agents, selective			0.14		
'C09AA'	ACE inhibitors plain			0.75		
'C10AA'	HMG CoA reductase inhibitors		0.91	0.65		
'J01CA'	Penicillins with extended spectrum					0.09
'J01CR'	Combinations of penicillins					0.06
'H03CA'	Iodine therapy				0.73	
'M01AE'	Propionic acid derivatives	0.24	0.33	0.29		0.21
'N01BB'	Amides				0.55	
'N02BB'	Pyrazolones				0.80	
'N02BE'	Anilides	0.11	0.40		0.87	0.24
'N05BA'	Benzodiazepine derivatives			0.23		
'R06AX'	Other antihistamines					0.09

Table 3
Description of the ICD9-CM and ATC codes with the five highest values in the profile of diagnoses and drugs per group (evaluation subset).

Code	Description	G0	G1	G2	G3	G4
'250'	Diabetes mellitus	0.94	0.95			
'272'	Disorders of lipoid metabolism		0.48	0.38		
'401'	Essential hypertension			0.92	0.77	
'460'	Acute nasopharyngitis	0.10	0.11			0.09
'463'	Acute tonsilitis					0.10
'526'	Diseases of the jaws	0.09				0.21
'558'	Unspecified gastroenteritis and colitis	0.07				0.08
'648'	Other current complication in pregnancy				0.32	
'650'	Normal delivery				0.48	
'664'	Trauma to vulva				0.48	
'692'	Contact dermatits and other eczema					0.11
'719'	Other and unspecified disorders of joint		0.31	0.15		
'724'	Other and unspecified disorders of back			0.08		
'780'	General symptoms	0.06	0.20	0.08		
'V27'	Outcome of delivery				0.88	
'A10AB'	Insulins for injection fast-acting	0.61				
'A10AD'	Insulins for injection intermediate-acting	0.17				
'A10AE'	Insulins for injection long-acting	0.70				
'A10BA'	Biguanides		0.77			
'B01AC'	Platelet aggregation inhibitors		0.51			
'B03AA'	Iron bivalent, oral preparations				0.55	
'C08CA'	Dihydropyridine derivatives			0.25		
'C09AA'	ACE inhibitors plain			0.53		
'C10AA'	HMG CoA reductase inhibitors		0.85	0.54		
'J01CA'	Penicillins with extended spectrum					0.07
'J01CR'	Combinations of penicillins					0.04
'H03CA'	Iodine therapy				0.62	
'M01AE'	Propionic acid derivatives	0.28	0.39			0.12
'N01BB'	Amides				0.70	
'N02BB'	Pyrazolones				0.88	
'N02BE'	Anilides	0.18	0.39	0.25	0.96	
'N05BA'	Benzodiazepine derivatives			0.23		
'R01AD'	Corticosteroids					0.05
'R06AX'	Other antihistamines					0.12

found. In the second stage, the discovery of clinical patterns was improved by clustering similar LDs with the AHC based on the symmetric KLD between every pair of PDFs linked to LDs. Our approach was validated on groups of healthy and chronic patients, characterized both with diagnoses and drugs. The proposed methodology allowed us to reasonably identify in an unsupervised way groups of patients with specific clinical conditions. For instance, we identified diabetic patients who were mainly distinguished by the pattern of drug consumption: insulin-dependent and noninsulin-dependent patients. Another group comprised mainly women with a complication during pregnancy, especially those affected by hypertension and diabetes. The literature review noted that one of the major pregnancy complications is the high probability of the onset of hypertension and diabetes [60]. The clinical patterns and the groups of patients discussed in this work are based on the results obtained in [32]. To the best of our knowledge, the novelty of the proposed methodology is an unsupervised approach for identifying patients with specific clinical conditions in a latent space. In addition, as mentioned previously, in the health care domain, interpretability is vital for the application of data-driven models in real scenarios. In this line, our approach could support the application of AE-based models in health care, improving their interpretation.

This methodology has also several limitations. Our approach showed a promising capacity for obtaining knowledge from LRs associated with clinical data, but its application is directly related to the number of LDs in the AE-based models. Further research in this line may explore new approaches that disentangle the potential relationship between LDs. Furthermore, the proposed methodology is established to work with continuous variables (either raw features or transformed features). An extension of this methodology for dealing with heterogeneous data will be considered in future work. It is also worth noting that we used EHR data during a natural year. The inclusion of new patient information during a long period of time would allow creating a tool for identifying health trajectories and monitoring the patient's health status evolution, potentially characterizing the dynamic course of a disease.

5. Conclusions

In this paper, we propose a novel unsupervised methodology that combined probabilistic methods for interpreting AE-based models with application in the clinical setting. This work can support clinical experts in daily practice, bringing some light to the decision-making process

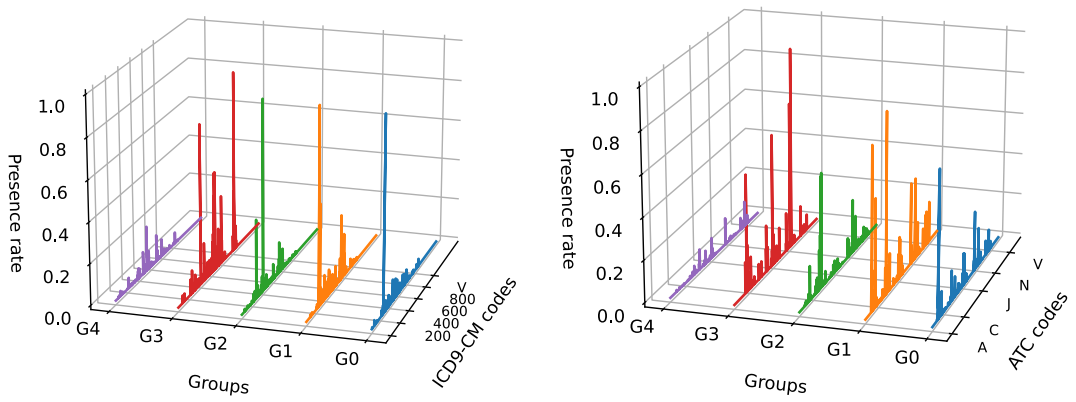

Fig. 8. Group profiles obtained when considering evaluation patients: (a) diagnosis profiles and (b) drug profiles.

when using nonlinear methods. Furthermore, our proposal has proven to be useful for visualizing patterns and extracting the most significant clinical features associated with specific clinical conditions. Our work opens up promising opportunities in the health care domain since the methodology can also be extrapolated to more complex and sophisticated scenarios or applications, allowing us to find groups of patients with complex health conditions, for instance, patients suffering from multimorbidity (two or more chronic conditions at the same time).

Availability of data and materials

The data used for this research comprise confidential patient health information, which is protected and may not be released unless approved by the Committee of Ethics of the UHF.

Declaration of competing interest

The authors declare no conflict of interest.

Acknowledgment

This paper resulted from the efforts made by many people of the clinical area, including Angel Gómez, Isabel Caballero, Javier Rivas and José Luis Rodríguez from University Hospital of Sureste and University Hospital of Fuenlabrada. David Chushig-Muzo wishes to thank Wouter Labuschagne, Marissa Labuschagne, Nuno da Costa and Frederick Atiah for their support during his research stay in Pretoria.

References

[1] Hall M, Dondo TB, Yan AT, Mamas MA, Timmis AD, Deanfield JE, et al. Multimorbidity and survival for patients with acute myocardial infarction in England and wales: latent class analysis of a nationwide population-based cohort. PLoS Med 2018;15(3).
[2] Glynn LG, Valderas JM, Healy P, Burke E, Newell J, Gillespie P, et al. The prevalence of multimorbidity in primary care and its effect on health care utilization and cost. Fam Pract 2011;28(5):516–23.
[3] Pefoyo AJK, Bronskill SE, Gruneir A, Calzavara A, Thavorn K, Petrosyan Y, et al. The increasing burden and complexity of multimorbidity. BMC Public Health 2015;15(1):415.
[4] Epping-Jordan J, Pruitt S, Bengoa R, Wagner EH. Improving the quality of health care for chronic conditions. BMJ Quality & Safety 2004;13(4):299–305.
[5] Guo Y, Liu Y, Oerlemans A, Lao S, Wu S, Lew MS. Deep learning for visual understanding: a review. Neurocomputing 2016;187:27–48.
[6] Shickel B, Tighe PJ, Bihorac A, Rashidi P. Deep ehr: a survey of recent advances in deep learning techniques for electronic health record (ehr) analysis. IEEE J Biomed Health Inform 2017;22(5):1589–604.
[7] Brelsford KM, Spratt SE, Beskow LM. Research use of electronic health records: patients' perspectives on contact by researchers. J Am Med Inform Assoc 2018;25(9):1122–9.
[8] Zhang Q, Hann E, Werys K, Wu C, Popescu I, Lukaschuk E, et al. Deep learning with attention supervision for automated motion artefact detection in quality control of cardiac t1-mapping. Artif Intell Med 2020;110:101955.
[9] Bernal J, Kushibar K, Asfaw DS, Valverde S, Oliver A, Mart R, et al. Deep convolutional neural networks for brain image analysis on magnetic resonance imaging: a review. Artif Intell Med 2019;95:64–81.
[10] Van Der Maaten L, Postma E, Van den Herik J. Dimensionality reduction: a comparative. J Mach Learn Res 2009;10:66–71.
[11] Hinton GE, Salakhutdinov RR. Reducing the dimensionality of data with neural networks. Science 2006;313(5786):504–7.
[12] Miotto R, Li L, Kidd BA, Dudley JT. Deep patient: an unsupervised representation to predict the future of patients from the electronic health records. Sci Rep 2016;6:26094.
[13] Beaulieu-Jones BK, Greene CS, et al. Semi-supervised learning of the electronic health record for phenotype stratification. J Biomed Inform 2016;64:168–78.
[14] Vincent P, Larochelle H, Bengio Y, Manzagol P-A. Extracting and composing robust features with denoising autoencoders. In: Proceedings of the 25th International Conference on Machine Learning, ACM; 2008. p. 1096–103.
[15] Samek W, Wiegand T, Müller K-R. Explainable artificial intelligence: understanding, visualizing and interpreting deep learning models. In: ITU Journal: ICT Discoveries - Special Issue 1 - The Impact of Artificial Intelligence (AI) on Communication Networks and Services 1; 2017. p. 1–10.
[16] Montavon G, Samek W, Müller K-R. Methods for interpreting and understanding deep neural networks. In: Digital Signal Processing. 73; 2018. p. 1–15.
[17] Guidotti R, Monreale A, Ruggieri S, Turini F, Giannotti F, Pedreschi D. A survey of methods for explaining black box models. In: ACM computing Surveys (CSUR). 51 (5); 2019. p. 93.
[18] da Cruz HF, Pfahringer B, Martensen T, Schneider F, Meyer A, Böttinger E, et al. Using interpretability approaches to update "black-box". In: Clinical prediction models: An external validation study in nephrology, artificial intelligence in medicine; 2020. p. 101982.
[19] Rebane J, Samsten I, Papapetrou P. Exploiting complex medical data with interpretable deep learning for adverse drug event prediction. Artif Intell Med 2020;109:101942.
[20] S. N. Payrovnaziri, Z. Chen, P. Rengifo-Moreno, T. Miller, J. Bian, J. Chen, X. Liu, Z. He, Explainable artificial intelligence models using real-world electronic health record data: a systematic scoping review, J Am Med Inform Assoc 27.
[21] Du M, Liu N, Hu X. Techniques for interpretable machine learning. Commun ACM 2019;63(1):68–77.
[22] Z. Lipton, The mythos of model interpretability, Communications of the ACM 61.
[23] Ribeiro MT, Singh S, Guestrin C. "Why should i trust you?" explaining the predictions of any classifier. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining; 2016. p. 1135–44.
[24] Lundberg SM, Lee S-I. A unified approach to interpreting model predictions. In: Proceedings of the 31st International Conference on Neural Information Processing Systems; 2017. p. 4768–77.
[25] Fan YJ. Autoencoder node saliency: selecting relevant latent representations. Pattern Recogn 2019;88:643–53.
[26] Charte D, Charte F, del Jesus MJ, Herrera F. An analysis on the use of autoencoders for representation learning: fundamentals, learning task case studies, explainability and challenges. Neurocomputing 2020;404:93–107.
[27] Kim J-Y, Cho S-B. Interpretable deep learning with hybrid autoencoders to predict electric energy consumption. In: International workshop on soft computing models in industrial and environmental applications. Springer; 2020. p. 133–43.
[28] Hughes JS, Averill RF, Eisenhandler J, Goldfield NI, Muldoon J, Neff JM, et al. Clinical risk groups (crgs): a classification system for risk-adjusted capitation-based payment and health care management. Med Care 2004:81–90.

[29] Pfister DG, Rubin DM, Elkin EB, Neill US, Duck E, Radzyner M, et al. Risk adjusting survival outcomes in hospitals that treat patients with cancer without information on cancer stage. JAMA Oncol 2015;1(9):1303–10.
[30] Mikalsen KØ, Soguero-Ruiz C, Bianchi FM, Jenssen R. Noisy multi-label semi-supervised dimensionality reduction. Pattern Recognit 2019;90:257–70.
[31] He H, Garcia EA. Learning from imbalanced data. In: IEEE Transactions on Knowledge and Data Engineering. 21 (9); 2009. p. 1263–84.
[32] Chushig-Muzo D, Soguero-Ruiz C, Engelbrecht A, Bohoyo PDM, Mora-Jiménez I. Data-driven visual characterization of patient health-status using electronic health records and self-organizing maps. IEEE Access 2020;8:137019–31.
[33] Soguero-Ruiz C, Mora-Jiménez I, Mohedano-Munoz MA, Rubio-Sanchez M, de Miguel-Bohoyo P, Sanchez A. Visually guided classification trees for analyzing chronic patients. BMC Bioinformatics 2020;21(2):1–19.
[34] American Medical Association. International classification of diseases, 9th revision, clinical modification: physician ICD-9-CM, 2005: volumes 1 and 2. In: Color-Coded, Illustrated. Vol. 1. Amer Medical Assn; 2004.
[35] World Health Organization, The anatomical therapeutic chemical classification system with defined daily doses (atc/ddd), Oslo: WHO.
[36] Bishop CM. Pattern recognition and machine learning (information science and statistics). Berlin, Heidelberg: Springer-Verlag; 2006.
[37] Soguero-Ruiz C, de Miguel-Bohoyo P, Mora-Jiménez I. A data-driven model based on support vector machine to identify chronic hypertensive and diabetic patients. In: Physiological Computing Systems. Springer; 2016. p. 110–29.
[38] Parzen E. On estimation of a probability density function and mode. Ann Math Stat 1962;33(3):1065–76.
[39] Kampffmeyer M, Løkse S, Bianchi FM, Jenssen R, Livi L. The deep kernelized autoencoder. Appl Soft Comput 2018;71:816–25.
[40] Silverman BW. Density estimation for statistics and data analysisvol. 26. CRC press; 1986.
[41] Chacón JE, Duong T. Multivariate kernel smoothing and its applications. CRC Press; 2018.
[42] Dempster AP, Laird NM, Rubin DB. Maximum likelihood from incomplete data via the em algorithm. J R Stat Soc B Methodol 1977;39(1):1–22.
[43] Schwarz G, et al. Estimating the dimension of a model. Ann Stat 1978;6(2):461–4.
[44] Rousseeuw PJ. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. J Comput Appl Math 1987;20:53–65.
[45] Khanmohammadi S, Chou C-A. A gaussian mixture model based discretization algorithm for associative classification of medical data. Expert Syst Appl 2016;58:119–29.
[46] Esfandiari N, Babavalian MR, Moghadam A-ME, Tabar VK. Knowledge discovery in medicine: current issue and future trend. Expert Syst Appl 2014;41(9):4434–63.
[47] Day WH, Edelsbrunner H. Efficient algorithms for agglomerative hierarchical clustering methods. J Classif 1984;1(1):7–24.
[48] Kullback S, Leibler RA. On information and sufficiency. Ann Math Stat 1951;22(1):79–86.
[49] Inselberg A, Dimsdale B. Parallel coordinates: a tool for visualizing multi-dimensional geometry. In: Proceedings of the first IEEE conference on visualization. IEEE; 1990. p. 361–78.
[50] Kingma DP, Ba J. Adam: A method for stochastic optimization. In: International conference on learning representations, San Diego, CA, USA, May 7–9; 2015.
[51] Maas AL, Hannun AY, Ng AY. Rectifier nonlinearities improve neural network acoustic models. Proc Mach Learn Res 2013;30:3.
[52] de Amorim RC, Hennig C. Recovering the number of clusters in data sets with noise features using feature rescaling factors. Inform Sci 2015;324:126–45.
[53] Atkinson MA, Eisenbarth GS, Michels AW. Type 1 diabetes. Lancet 2014;383(9911):69–82.
[54] Cho Y, Choe E, Lee Y-h, Seo JW, Choi Y, Yun Y, et al. Risk of diabetes in patients treated with hmg-coa reductase inhibitors. Metabolism 2015;64(4):482–8.
[55] Würtz M, Grove EL, Kristensen SD, Hvas A-M. The antiplatelet effect of aspirin is reduced by proton pump inhibitors in patients with coronary artery disease. Heart 2010;96(5):368–71.
[56] Kjeldsen SE. Hypertension and cardiovascular risk: general aspects. Pharmacol Res 2018;129:95–9.
[57] Axmon A, Kristensson J, Ahlström G, Midlöv P. Use of antipsychotics, benzodiazepine derivatives, and dementia medication among older people with intellectual disability and/or autism spectrum disorder and dementia. Res Dev Disabil 2017;62:50–7.
[58] Price HR, Collier AC. Analgesics in pregnancy: an update on use, safety and pharmacokinetic changes in drug disposition. Curr Pharm Des 2017;23(40):6098–114.
[59] Brown CM, Garovic VD. Drug treatment of hypertension in pregnancy. Drugs 2014;74(3):283–96.
[60] Bellamy L, Casas J-P, Hingorani AD, Williams D. Type 2 diabetes mellitus after gestational diabetes: a systematic review and meta-analysis. Lancet 2009;373(9677):1773–9.