A graph-embedded topic model enables characterization of diverse pain phenotypes among UK Biobank individuals

Yuening Wang¹, Rodrigo Benavides², Luda Diatchenko^{3,4,5}, Audrey V. Grant^{3,4,5,*} and Yue Li^{1,*}

¹School of Computer Science, McGill University

³Department of Anesthesia, McGill University

⁴Faculty of Dentistry, McGill University

⁵Alan Edwards Centre for Research on Pain, McGill University

²Department of Anesthesiology, Centro Nacional de Rehabilitación, San Jose, Costa Rica

*Correspondence to audrey.grant@mcgill.ca, yueli@cs.mcgill.ca

Abstract

Large biobank repositories of clinical conditions and medications data open opportunities to investigate the phenotypic disease network. To enable systematic investigation of entire structured phenomes, we present graph embedded topic model (GETM). We offer two main methodological contributions in GETM. First, to aid topic inference, we integrate existing biomedical knowledge graph information in the form of pre-trained graph embedding into the embedded topic model. Second, leveraging deep learning techniques, we developed a variational autoencoder framework to infer patient phenotypic mixture. For interpretability, we use a linear decoder to simultaneously infer the bi-modal distributions of the disease conditions and medications. We applied GETM to UK Biobank (UKB) self-reported clinical phenotype data, which contains conditions and medications for 457,461 individuals. Compared to existing methods, GETM demonstrates overall superior performance in imputing missing conditions and medications. Here, we focused on characterizing pain phenotypes recorded in the questionnaire of the UKB individuals. GETM accurately predicts the status of chronic musculoskeletal (CMK) pain, chronic pain by body-site, and non-specific chronic pain using past conditions and medications. Our analyses revealed not only the known pain-related topics but also the surprising predominance of medications and conditions in the cardiovascular category among the most predictive topics across chronic pain phenotypes.

1 Introduction

The advent of electronic health records (EHR) has started to transform the way healthcare data are recorded and used in clinical practice and in research settings. Besides free-form clinical text, most modern healthcare centers now routinely collect structured EHR data describing comprehensive aspects of care, including diagnosis, medications, treatments, laboratory test results, and other measures. Deriving coherent phenotypes from these EHR data is crucial in downstream phenome-wide association analyses (PheWAS) and may greatly improve the power of detecting genetic associations using the genome-wide association (GWA) approach [1]. Besides better characterizing known phenotypes (e.g. through co-morbidities or the demographics of age-of-onset), mining phenomic data also has the potential to reveal novel combinations of diseases and other variables of potential etiological interest. This will help identify specific strata of study subjects most at risk for disease or targeted for specific drug recommendations. Despite these promises, clinical phenotype data sources remain under-used [2]. As genotype and deep phenotype data become increasingly available through consortia or large government-funded cohorts such as the UK Biobank (UKB) [3] and the 100,000 Genomes Project [4], there is an urgent need for an automatic and accurate phenotyping tool to accelerate novel disease co-morbidity discoveries and improve the yield of GWA studies for complex phenotypes and diseases in humans.

Among many machine learning approaches, topic models [5] stand out as a particularly well-adapted framework for automatic phenotyping. They are extremely efficient at modeling sparse and discrete data such as text documents. Topic models were originally developed to discover patterns of word usages from corpuses of text documents by accomplishing two related tasks: (1) inferring a set of latent categorical distributions over the vocabulary (i.e., topics); (2) using these latent topic distributions to infer topic mixture memberships of each document, thereby connecting them under similar topical themes. In our context, we consider EHR as our documents and the diagnostic and medication codes as our vocabulary. Several topic methods were developed recently for effectively mining EHR data [6, 7]. However, existing topic models are unable to incorporate existing biomedical knowledge graphs, which manifest in several forms such as disease taxonomy and drug classification systems. Standard inference algorithms of topic models (e.g., Gibbs sampling or mean-field variational inference) also have limited flexibility to capture the non-linear connections between observed EHR codes and the underlying patient phenotypes.

In this paper, we present a graph-embedded topic model (GETM) for learning phenotypes from heterogeneous EHR data by leveraging biomedical graph information. As the main method contribution, GETM seamlessly integrates two existing models: the embedded topic model (ETM) [8] and node2vec [9]. Briefly, we first use node2vec to learn the condition and medication embeddings based on their taxonomic information; then, we incorporate these embeddings into the ETM, which tri-factorizes the individuals-by-conditions/medications matrix into individuals-by-topics, topics-by-embedding, and embedding-by-conditions/medications matrices. The distribution of the individuals-by-topics is approximated by the output of a feed-

forward neural network using the amortized variational inference technique and the embedding-by-conditions/medications matrices are fixed to the node2vec-learned node embedding from conditions/medications taxonomical graphs, respectively.

As a proof-of-concept, we applied GETM to UKB phenotype data, where 457,461 individuals of European descent from across the United Kingdom were deeply phenotyped through extensive self-report based questionnaires for about 443 well-defined phenotype conditions and 802 active ingredients of medications [3]. We then turned to a more focused analysis on predicting and characterizing different pain phenotypes. Chronic pain is the result of dysfunction of the nociceptive circuitry leading to sustained perception of pain. Chronic pain is highly prevalent in aging populations affecting up to 50% of older adults (>65y) [10] and decreases the overall mental and emotional health of affected individuals [11]. Using GETM, we provide a refreshing view of pain phenotypes, considered as CMK pain, chronic pain by body site, non-specific acute and chronic pain, and the transition from acute to chronic pain by making use of phenotype data from subsequent visits on a subset of the UKB study population. By correlating the inferred GETM-topics and pain phenotypes across the UKB subjects, we discover not only the known pain-related conditions and medications among the highly predictive topics but also novel combinations after removing labels of known pain-related conditions and analgesics.

2 Results

2.1 Graph embedded topic model overview

GETM models the distribution of medications and conditions as discrete clinical features for each individual. For a given subject of study, the expected rate of each feature is determined by both the logistic-normal latent patient's topic mixture and the point-estimate latent topic distributions over the features. The goal of GETM is to approximately infer the distributions of these latent variables. To this end, we carry out an amortized variational inference [12] in two steps (**Fig. 1**). In the first step, to infer the topic mixture of a given patient, we provide to a feedforward neural network (i.e., the encoder) the binary vector of the individual's observed discrete features. We then sample the topic mixture from a variational Gaussian distribution with the mean and standard deviation computed by the encoder. In the second step, we decode the sampled topic mixture back to the original conditions and medications. We used two linear decoders each with separate topic/feature embeddings for medications and conditions, respectively. Specifically, each decoder tri-factorizes the individuals-by-features matrix respectively into individuals-by-topics (θ), topics-by-embeddings ($\alpha^{(t)}$), and embeddings-by-features (medication or condition) ($\rho^{(t)}$) matrices ($t \in \{med, cond\}$). Notably, the two linear decoders share the patient-level topic mixture θ while $\alpha^{(t)}$ and $\rho^{(t)}$ are learned separately. Importantly, the medication embedding $\rho^{(med)}$ and condition embedding $\rho^{(cond)}$ are pre-trained by node2vec [9] from the taxonomic tree-structured graphs of conditions and medications while the two sets of topic embedding counterparts ($\alpha^{(med)}$, $\alpha^{(cond)}$) are directly learned from the UKB participant data by GETM. This tri-factorization design allows for exploring topics, study subjects and relationships

among conditions and medications in a highly interpretable way.

2.2 Topic quality evaluation

Using data from the UKB on 457,461 individuals of European descent from the baseline visit, we trained GETM to obtain the topic embedding and conditions/medications embedding (i.e., $\alpha^{(t)}$, $\rho^{(t)}$, respectively). As a qualitative exploratory analysis, we visualized these embeddings using UMAP (**Fig. 2**). For illustration purpose, we randomly picked 5 topics and observed that the top features under these topics belong to coherent categories of conditions or medications. For example, top medications atenolol, bisoprolol, metoprolol, nebivolol and carvedilol in topic 32 all belong to cardiovascular-system medications (**Fig. 2d**), while topic 32 is assigned to the cluster mainly composed of medications from same category (**Fig. 2b**). Moreover, the top 5 conditions from topic 8 are all from the musculoskeletal/trauma category while the top 5 medications from topic 8 are all from the dermatological category. In contrast, ETM without using the KG-informed embedding produced less interpretable topics each covering heterogeneous categories (**Fig. S1**). Thus, GETM allows for the identification of related features from sparse, heterogeneous data in a data-driven manner.

Given that we are leveraging embeddings of conditions and medications learned from their taxonomic graphs, we expected a higher quality of topics to be inferred by GETM compared to baseline models that do not use or use only partial graph embeddings. Topic quality was quantified as the product of topic coherence and topic diversity [8] (**Methods**). For ease of reference, we listed the 9 models that we compared in **Supplementary Table S1**. We first evaluated the quality of the medication-defined topic (**Table 1**). To evaluate the medication topic coherence unbiasedly, we used an external set of 59 expert-curated medication categories (**Methods** Section 6.9) that were not part of the medication taxonomy we used in training the graph embedding for our GETM model. Overall, we observed the highest topic quality (0.6026, **Table 1**) for the 50-topic GETM that used the graph embeddings for both the conditions and medications (i.e., GETM in **Supplementary Table S1**). Therefore, the results suggest the benefits of using knowledge-graph(KG)-informed medication embeddings in improving topic quality.

We also evaluated condition-derived topic quality. Due to the lack of external condition categories, we calculated topic coherence based on whether the top conditions from each topic were present in the same patients. One caveat of this approach is that a patient rarely exhibits multiple conditions from the same category and that many conditions can be mutually exclusive within the same patient. For example, asthma and COPD as shown in our GETM topic 60 in **Fig. 2b** are rarely observed in the same patient despite the common physiological connection between them. This is reflected in the low topic coherence among the methods. Nonetheless, we found that the GETM with pretrained embeddings of both medications and conditions (GETM, **Supplementary Table S1**) dominate other models among 3 out of 4 model settings in terms of the overall topic quality scores.

2.3 GETM reveals known or potentially novel condition-medication relations from UKB data

We compared the total number of unique known pairs between medications and conditions that were identified by five models (**Supplementary Table S1, Table 2**). The full GETM extracted most pairs of correlated conditions and medications illustrated by topic number 50 (161 pairs), 75 (175 pairs), and 100 (203 pairs). We examined some of the topics with high condition-medication associations based on pharmacological knowledge (**Fig. 2b,d**). For instance, in topic 32, the top medication bisoprolol is known to be used to treat the top condition heart failure under that topic. Under topic 60, the top medication salmeterol is often prescribed to treat asthma and chronic obstructive airways (COPD) [13], which are the top conditions under that topic. Interestingly, although the medication solifenacin in topic 59 is not known to be indicated for depression according to DrugBank, recent research shows potential for solifenacin, along with other muscarinic antagonists, in treating patients with depression through drug repurposing [14, 15].

2.4 GETM accurately imputes UKB conditions and medications

To further ascertain the overall utility of GETM, we used GETM to reconstruct randomly masked 50% of the medications or conditions of each test individual. As a baseline, we used the observed condition or medication prevalence from the training data to impute the masked feature in the test individuals. We have also experimented with the ablated models listed in **Supplementary Table S1**. Among all models, the full GETM conferred the lowest reconstruction errors for both conditions (**Table 3**) and medication (**Supplementary Table S5**) using 100 topics with ~5% improvement over the second best method. The improvements are largely attributed to the use of KG-informed embeddings compared to the ablated ETM model that learn these embeddings from the data.

Moreover, we also evaluated GETM in its ability to recapitulate the medication data using only the condition information by masking all medication information of the test individuals. The full GETM outperformed all other models and gave the lowest reconstruction error of 14.6125, the highest precision at top 5 predicted hits (precision@5) (0.2612), and the highest recall@5 (0.5787) (**Table 4,5**). The upper bounds, which were obtained from the GETM trained on unmasked test data achieved reconstruction error of 11.4936, precision@5 of 0.4226, and recall@5 of 0.8027. Interestingly, we found out that on average around 44% of the unobserved medications from the top 10 predicted medications by GETM in fact have a treatment effect based on condition-medication association information extracted from Comparative Toxicogenomics Database (CTD) <http://ctdbase.org/> and DrugBank <https://go.drugbank.com/>. In particular, we took a closer look at the 3 best predicted patients and 3 worse predicted patients by GETM (**Fig. 3**). As expected, among the best predicted patients, the top 10 predicted medications by GETM contain mostly observed medications, which are known to match the subset of the conditions of these 3 patients. Although most top predicted medications were

not observed for the worst 3 predicted patients, they were actually known to treat the observed conditions for these patients.

2.5 CMK pain prediction

We sought to investigate the predictive ability of reported past conditions and current medications on pain experience in the UKB using GETM. We first focused on CMK pain because this composite pain phenotype has the highest number of positive cases. Using logistic regression (LR), we evaluated the predictive accuracy of the inferred topic mixture on CMK pain in terms of area under the receiver operating characteristic (AUROC) and area under the precision-recall curve (AUPRC). As a baseline (i.e., the raw model), we trained another LR model that directly used all of the 443 conditions and 802 medications. We also sought to investigate the relative improvements of GETM over standard topic modeling after removing obvious condition or medication features for CMK pain. We iterated through 6 different ways of filtering out features based on ORs or expert knowledge (**Methods**).

Predictive performance using original features (i.e., unfiltered) and filtered features is summarized in **Fig. 4**. From these results, we observed that (1) The GETM topic mixture (GETM, **Supplementary Table S1**) achieved larger AUROC and larger AUPRC across all feature filtering regimes (i.e., no filter plus the 6 filtering rules; **Supplementary Table S2**); (2) As we removed more pain-related signature conditions and medications, the performance of using raw features dropped more drastically than that using the patient topic mixtures. In other words, the relative improvement of using patient topic mixtures over using raw data increased as we removed more indicative conditions and medications. In particular, GETM conferred a larger than 40% improvement over the baseline when using the fewest conditions and medications (i.e., m579c379: filtered feature set 7) in contrast to less than 5% improvement over the baseline when using all features (m802c443: filtered feature set 1; **Fig. 4**; **Supplementary Table S2**). These results echo the superior imputation performance of GETM we observed above. This is because GETM uses the pre-trained graph embedding and inferred topic mixture to compensate for the information loss from the feature removals.

2.6 CMK pain-related conditions and medications

We then investigated the most pain-related conditions and medications based on LR coefficients and calculated overlapping proportions with physician-curated lists (**Fig. 5**). In comparison with the overlapping proportions from ETM and odds ratio calculation, GETM identified a much greater proportion of known pain-related conditions among the top conditions and medications under these topics (36.7% from top 10 conditions, 33.3% from top 30 conditions and 30.0% from top 50 conditions) and medications (60.0% from top 10 medications, 33.3% from top 30 medications and 32.0% from top 50 medications) in the provided lists. This suggests that GETM improves the ability to extract otherwise hidden associations, which could identify pain-related comorbidities. Therefore, GETM does not only confer superior performance but

higher model explainability from its inferred topics.

We also had a close look at the three most positively associated topics and three most negatively associated topics to CMK pain based on learned ω (Figure 6a). This analysis showed that topics 56, 34, and 51 are strongly positively associated with CMK pain and topics 73, 68, 89 are strongly negatively associated with CMK pain, respectively. For each topic, we examined semantics according to the top conditions and top medications (Figure 6b,c). Particularly, topics 56 and 34 contained musculoskeletal system conditions and medications, which are clinically meaningful since they are highly related to CMK pain. In particular, prolapsed disc or slipped disc as a condition is painful, and ibuprofen is an analgesic in the NSAID (non-steroidal anti-inflammatory drug) class.

Topic 51 is in the cardiovascular category, and the top medication acetylsalicylic acid (aspirin), also an NSAID, is prescribed more frequently and typically at lower doses for its cardioprotective properties in prevention of stroke and heart attacks; it acts as a “blood thinner”. Dipyridamole inhibits blood clot formation and therefore prevents potential consequences of blood clotting. The top condition under topic 51 is high cholesterol, which is a known and common risk factor for atherosclerosis. Atherosclerosis is a process of deposition of fatty material in the walls of arteries, and this thickening leads to an increase in stroke and heart attack risk. Thus, although the two medications are not directly used as a cure for the condition high cholesterol, by way of atherosclerosis, high cholesterol leads to higher risk for other cardiovascular outcomes and the medications are used prevent those outcomes [16].

Contrary to the risk topics, the protective topics identified combinations of medications and conditions that did not yield as straightforward pairings or links. Topic 73 is one of the top 3 negatively predictive topics of CMK pain. Its top medications identified were ramipril, lisinopril and enalapril, which are ACE inhibitors, used to treat high blood pressure and may be used in response to heart failure or heart attack. Incidentally, all conditions under topic 73 may be categorized as allergic or atopic with hayfever contributing the most.

This finding suggests that a particular subset of individuals suffering from allergic conditions, whose immunity is therefore skewed, and who are also undergoing cardiovascular treatment, are at lower risk for CMK pain. The implication of immune mechanisms in chronic pain manifestation is well known, where proinflammatory states are associated with chronic pain [17]. Also, high blood pressure, the most common indication for ACE inhibitors [18] has known analgesic effects that persist with treatment [19, 20]. Thus, hypertension related analgesia would help explain this negative predictive topic.

Another negatively predictive topic of CMK pain, topic 89, is limited to women given the top medication, conjugated estrogens. The top condition, hypertension, as seen above, is known to have analgesic effects. Sex hormones including estrogens modulate pain, and estrogens have been found to have neuroprotective, antinociceptive properties. These could explain the protective association observed here on CMK pain. Indeed, estrogens are known to influence chronic pain conditions [21]. These combinations of medications and conditions within topics could lead to new CMK pain etiology hypotheses for further exploration.

2.7 Characterizing diverse pain types by topic analysis

We extended our analysis from CMK to predicting several other definitions of pain phenotypes, including acute pain, the acute to chronic pain transition for a subset of UKB individuals, chronic pain at specific body sites (neck/shoulder, hip, back, stomachache/abdominal, knee, headache, face) and chronic pain all over body. We used three different feature filtering regimes to progressively remove obvious pain-type predictors (**Methods** 6.16). Logistic regression using the GETM's topic mixture conferred larger AUROC and larger AUPRC values across all three feature filtering regimes (**Fig. 7; Supplementary Table S1**). We then examined the relative contributions of the medication/condition categories for each pain type (**Supplementary Fig. S3; Methods** Section 6.5). The results we presented below came from at least one of 3 feature filtering regimes as highlighted in **Supplementary Fig. S3**. In addition to defined categories, we used the ATC classification system to consider the contribution of the medication sub-class analgesics, which has a direct relationship with pain. By considering the implication of condition and medication categories in predicting different pain phenotypes, we were able to identify expected trends as well as consider and interpret unexpected findings.

Overall, the acute and acute-to-chronic transition displayed similar trends. Here the acute-to-chronic transition is defined such that non-site specific MSK acute pain observed at the first visit turned into chronic pain as diagnosed at the following visit of the UKB individuals. As an example, low positive prediction from analgesics and high negative contribution were seen for acute pain (**Fig. 8**).

As expected, the category from the literature that displays highest levels of chronic pain co-morbidities, neurology/eye/psychiatry (conditions) and nervous system (medications) contributed the most to prediction of chronic pain, particularly headache and face pain (**Supplementary Fig. S3a-I, b-I**). Endocrine/diabetes (conditions) exhibited highly prominent contribution for the acute phenotype classification (**Supplementary Fig. S3a-II**). Immunological/systemic disorders showed a strong contribution for protection from acute pain (**Supplementary Fig. S3c-I**). Gastrointestinal/abdominal (conditions) and alimentary tract and metabolism (medications) showed an expected higher positive proportion of contribution towards stomach/abdominal pain predication compared to other chronic pain types (**Supplementary Fig. S3a-III**).

Knee pain compared to face pain and headache exhibited a sharp contrast in predictive compositions (**Supplementary Fig. S4a, b**). Interestingly, obesity-related medication categories such as systemic hormonal preparations excluding sex hormones and insulins exhibited the highest protective contribution for chronic knee pain (**Supplementary Fig. S5d-I**). Chronic knee pain also had the highest positive contribution toward cardiovascular condition prediction and no negative contribution (**Supplementary Fig. S4a-I, c-I**). A similar but less striking pattern was observed for hip pain (**Supplementary Fig. S4a-II, c-II**). For headache and face pain, this trend was reversed with the highest protective contribution and lowest positive contribution by cardiovascular conditions (**Supplementary Fig. S4a-III, c-III**).

Intriguingly, the cardiovascular medications category exhibited strong negative prediction of

0.7 across all chronic pain phenotypes but not acute pain (**Fig. 8**). This suggests that topics containing cardiovascular medications, other than analgesics, had a protective effect against chronic pain. Headache, and more specifically migraine, has a known vascular etiological component. Nonetheless, the patterns seen across all chronic pain phenotypes is similar, indicating a more general chronic pain protective effect for cardiovascular medications.

3 Discussion

Large biobanks with clinical data such as the UKB are valuable resources enabling greater understanding of factors impacting the manifestation and treatment regimens of complex diseases such as chronic pain. However, the sparsity, heterogeneity and sheer size of these data pose challenges when restricted to a classical statistical toolbox. The typical biomedical researcher is hindered from taking full advantage of these invaluable data resources. As a result, most investigations focus on one or a small subset of related diseases [22–24]. In the present study, we developed GETM to model all self-reported conditions and medications among 457,461 UKB subjects. By introducing the knowledge graph and simultaneously training different types of features, GETM was able to infer more coherent topics compared to the ablated models without using the KG embedding (**Table 1** and **2**). In contrast to the topic modeling without the graph embedding, this allows for better interpretation of the disease topics and any findings related to certain topics with clearer clinical grounds.

As demonstrated using UKB data, GETM achieved superior performance in imputing 50% of the observed conditions (**Table 3**) and 50% of the observed medications over all test individuals (**Table 4**) as well as in imputing the entire medication records based only on the conditions records of test individuals (**Table 5**). Many top unobserved medications predicted by GETM also exhibit known connections with the observed conditions of the subject (**Fig. 3**). Importantly, GETM offers excellent model explainability and can be used to discover meaningful disease comorbidities and disease-medication combinations via topic analysis (**Fig. 2**).

In a focused analysis to predict chronic musculoskeletal (CMK) pain, the same logistic regression (LR) classifier using GETM-inferred topic mixtures conferred robustly and consistently higher AUROC and AUPRC compared to using the raw features (**Fig. 4**). The top conditions and medications under the most predictive topics are enriched for elements from the physician-curated condition list and medication list, respectively (**Fig. 5**). Also, the most predictive CMK pain topics contained top conditions and medications that were strongly associated with CMK pain (**Fig. 6**), implying their potential usage as phenotyping markers. Additionally, the GETM topic mixtures are robust to the removal of clinically obvious conditions and medications (**Fig. 4**). This is reflected in the larger relative improvements over the LR classifier that operates on the raw (filtered) features. This implies its utility in predicting predisposition for CMK pain for those individuals with no reported pain symptoms. We observed consistent quantitative performance when extending to other pain types (**Fig. 7**).

Several existing methods utilize topic models to find meaningful latent topics from electronic

health record (EHR) data using structured administrative data such as the ICD codes [6, 7]. In contrast, we demonstrated the utility of GETM on the less structured and more sparse self-reported questionnaire information from the UKB including 443 conditions and 802 medications. We expect that GETM would work equally well if not better on more structured EHR data and leave that to future exploration. Indeed, although we used UKB data and focused on pain phenotypes as a case study, our approach is a generalizable and highly efficient method that can be used to characterize other phenotypes in the UKB or from other similarly scaled biobanks.

In further developing GETM methodology, we will allow for integration of longitudinal health-care information such as age-of-onset across conditions and multiple follow-up timepoints. While this is a key direction, in our present context, we only have baseline data for most the UKB subjects. As longitudinal data become increasingly available, we will extend our model to a dynamic topic model that accounts for the evolution of subjects' health status over time. We will also improve embedding learning by considering a more expressive graph embedding approach than node2vec. For instance, we will consider a graph convolutional neural network (GCN) [25] to produce the embedding used by the ETM model. Furthermore, we will explore a multi-relational graph approach [26] to model the relationships within and among conditions and medications. Since both the GCN and the ETM models share the same objective function (i.e., the evidence lower bound), we are able to perform end-to-end training instead of the pipeline approach presented here. To produce competitive performance with this approach, however, careful model fine-tuning of both the GCN and ETM are needed at the expense of computational resources.

We now discuss the epidemiological implications of our results on the pain analysis in the context of existing studies. First of all, the UKB as a cross-sectional study creates challenges in interpreting the role of temporality and identifying putative causal effects. The pain questionnaire was administered at baseline, and medications reported were taken regularly, also at baseline. Conditions were considered to occur at any time in the subjects' past, through a record of age-of-onset for each condition. Thus, the impact of medications on reported conditions do not allow for direct investigations of medication efficacy. In classifying chronic pain, the pain questionnaire referred to a time including at least the prior three months in the past from baseline, but this could extend to any length in a given subject's history. By considering the subset of individuals without baseline chronic pain, but who developed chronic pain by the time of a subsequent visit, it would be possible to approximately quantify the total length of time of pain chronicity. We may also consider chronification as temporally subsequent to the appearance of conditions and taking of medications from prior to baseline. This is only true for a subset of UKB subjects, however. Thus, the UKB study design has limited potential to evaluate chronic pain causality. Rather, our analyses produced hypotheses that might best be tested in the context of other cohorts that are truly longitudinal in design.

Our findings identified intriguing links between cardiovascular medications and their protective effect for chronic (but not acute) pain that stood out as particularly predictive compared with other medication or condition categories (**Fig. 8**). Also, cardiovascular conditions were negatively predictive for headache and face pain particularly. Recent meta-analyses have at-

tempted to quantify the relationship between cardiovascular conditions and chronic pain [27, 28], but their relevance to our findings is low. None of these, nor the studies that were included in the meta-analyses explicitly quantified the effects of taking medications. Also, the direction of causality focused on chronic pain as a risk factor for cardiovascular outcomes, including mortality, and did not consider reverse causality; authors cited diversity in outcomes and in chronic pain taxonomy making meta-analysis results generally inconclusive [27]. The more recent meta-analysis was more limited in scope, estimating that people with CMK pain were 1.91 times more likely to report having a cardiovascular disease compared with those without CMK pain with statistical significance [28].

The dominance of the predictive effect of cardiovascular medications across chronic pain phenotypes leads to several avenues for further exploration. The specifics of what medications and what conditions are found in individually highly predictive topics would need to be explored in further detail. Also, further analyses are needed to identify whether more specific sub-categories of cardiovascular medications drove the strong cardiovascular medications predictive component. The extent of similarity across specific medications within and across topics, such as common mechanisms of action, would be of particular interest. Several medications united by a common mechanism would have a higher potential for generalizability beyond the limited population subsets represented by individual items or topics. This was illustrated by our focused analysis on CMK pain, where the predictive medications identified in the top protective topic (**Fig. 6b,c**, Topic 73) were all ACE inhibitors. Given that the most common indication for this medication class is hypertension, perhaps one of the protective cardiovascular medication effects at play was hypertension-associated hypoalgesia, especially since the hypoalgesia effect is maintained even if the hypertension is treated [19, 29]. Baroreceptor sensitivity is the most studied specific mechanism [19, 29]. Hypertension-associated hypoalgesia manifests through a range of phenotypes from pain sensitivity to pain chronification, with higher blood pressure associated with lower pain sensitivity or lesser chronification. This has been demonstrated in both animal experimental and human observational studies [19, 30]. Nonetheless, hypertension-associated hypoalgesia is not well recognized.

Perhaps the present report will lead to findings supporting greater clinical and public health importance for hypertension-associated hypoalgesia in preventing pain chronification than previously thought. Beyond ACE inhibitors, cardiovascular medications classified as beta blockers are used as prophylaxis for migraine headaches [31]. They have been shown to have analgesic effects for other types of chronic pain [32, 33]. There may also be further sub-categories of cardiovascular medications of interest. Nonetheless, there are variable time frames between development of cardiovascular conditions, taking of cardiovascular medications and development of chronic pain in the present study. Therefore, other longitudinal studies that control these sources of variability would be needed to better understand the protective effect suggested by the present analysis.

4 Acknowledgements

Y.L. is supported by Natural Sciences and Engineering Research Council (NSERC) Discovery Grant (RGPIN-2019-0621), Fonds de recherche Nature et technologies (FRQNT) New Career (NC-268592), and Canada First Research Excellence Fund Healthy Brains for Healthy Life (HBHL) initiative New Investigator start-up award (G249591). AVG and YW were partly funded by the Canadian Excellence Research Chairs (CERC09). LD was supported by the Canadian Excellence Research Chairs fund (CERC09), a Pfizer Canada Professorship in Pain Research, and CIHR (SCA-145102) for Health Research's Strategy for Patient-Oriented Research (SPOR) in Chronic Pain. The current study was conducted under UK Biobank application 20802.

5 Author contributions

These authors jointly supervised this work: Audrey V. Grant and Yue Li. Y.L. and A.V.G. conceived the study. Y.L. and Y.W. developed the methodology. Y.W. created the computational software and ran the analyses. Y.W., Y.L., and A.V.G. analyzed the results and wrote the initial manuscript. L.D. and R.B. participated in initial discussions and subsequently in critical writing related to pain phenotypes and medications. All of the authors wrote the final version.

6 Methods

6.1 UKB data processing

For conditions data, we used UKB datafield 20002, which records self-reported non-cancer diseases for each subject. This was collected by questionnaire during participant interviews. The participants were asked whether or not they had been diagnosed with certain conditions (heart attack, angina, stroke, high blood pressure, blood clot in leg, blood clot in lung, emphysema/chronic bronchitis, asthma or diabetes), were asked to add any other conditions and to provide a date of diagnosis for each when possible [34]. Medication usage data was similarly collected during participant interviews prompted by a question on regular prescription medication use, resulting in datafield 20003 which contains treatment/medication codes [35].

We kept 457,461 individuals of European descent to reduce potential confounding by ethnic group. In total, 802 active ingredients were kept as medications and 443 conditions were extracted. Here we encoded the medications and conditions as binary variables. Only baseline visit data was included.

6.2 UKB pain-related phenotype labels extractions

We sought to associate pain phenotypes with comorbid conditions and medications. To this end, we used the pain-related phenotypes as the label data and patient-dependent topic mix-

ture derived from the 443 conditions and 802 medications as the input features in a post-hoc supervised topic analysis as detailed in Section 2.5.

The pain phenotypes were collected through questionnaire. We drew on datafield 6159; participants were asked “In the last month have you experienced any of the following that interfered with your usual activities? (You can select more than one answer).” If they answer “yes” to pain at any site, for example, back pain, they were further asked “Have you had back pains for more than 3 months?” (datafield 3571), and so on. In total, we collected 9 pain-related labels with data field identifiers listed as follows:

1. pain type(s) experienced in last month: 6159
2. headaches for 3+ months: 3799
3. facial pains for 3+ months: 4067
4. neck/shoulder pain for 3+ months: 3404
5. back pain for 3+ months: 3571
6. stomach/abdominal pain for 3+ months: 3741
7. hip pains for 3+ months: 3414
8. knee pains for 3+ months: 3773
9. general pain for 3+ months: 2956

We can also query the UKB website to obtain relevant information for each field by replacing x in <https://biobank.ndph.ox.ac.uk/ukb/field.cgi?id=x> with the above data field identifiers.

6.3 Graph embedded topic model (GETM) generative process

We formulated the problem of modelling discrete patient healthcare data into a topic modeling problem. In particular, we treat each of the D patient health records as a document and each observed feature in the record as a word sampled from a defined vocabulary. In our case, we have two vocabularies covering $M = 802$ medications and $C = 443$ conditions, respectively. With this analogy, each patient record is represented as a mixture of latent topics. In the original latent Dirichlet allocation (LDA) model [36], a topic distribution over a vocabulary is defined as an independent Dirichlet prior $\beta_k \sim \text{Dirichlet}(\tau_\beta)$, where $\sum_{w=1}^V \beta_{k,w} = 1$ over V words. Inspired by a more recent work called embedded topic model (ETM) [8], we decomposed the K topic distributions over medications $\beta^{(med)}$ into a medication-defined topic embedding $\alpha^{(med)} \in \mathbb{R}^{K \times L_1}$, and a medication embedding $\rho^{(med)} \in \mathbb{R}^{L_1 \times M}$, where L_1 denotes the medication embedding dimension and M denotes the number of unique medications. Similarly, the topic distribution over conditions $\beta^{(cond)}$ is proportional to the inner product of the condition-defined topic embedding $\alpha^{(cond)} \in \mathbb{R}^{K \times L_2}$, and condition embedding $\rho^{(cond)} \in \mathbb{R}^{L_2 \times C}$, where L_2 denotes the condition embedding dimension and C denotes the number of unique conditions.

484 For the d^{th} patient record ($d \in \{1, \dots, D\}$), the generative process starts by drawing the topic
485 mixture θ_d from a logistic normal distribution $\theta_d \sim \mathcal{LN}(0, \mathbf{I})$:

$$\eta_d \sim \mathcal{N}(0, \mathbf{I}); \quad \theta_d = \text{softmax}(\eta_d) = \frac{\exp(\eta_d)}{\sum_{k=1}^K \exp(\eta_{d,k})} \quad (1)$$

We then draw the i^{th} medication token or the j^{th} condition token from two respective categorical distributions:

$$w_{i,d}^{(med)} = \prod_m r_{m,d}^{[w_{i,d}^{(med)}=m]}, \quad w_{j,d}^{(cond)} = \prod_c r_{c,d}^{[w_{j,d}^{(cond)}=c]} \quad (2)$$

where

$$r_{m,d} = \sum_k \theta_{d,k} \beta_{k,m}^{(med)} = \theta_{d,\cdot} \beta_{\cdot,m}, \quad r_{c,d} = \sum_k \theta_{d,k} \beta_{k,c}^{(cond)} = \theta_{d,\cdot} \beta_{\cdot,c} \quad (3)$$

486 Here we use notation $\theta_{d,\cdot}$ to denote a $1 \times K$ row vector for the d^{th} patient record and $\beta_{\cdot,m}$ (and
487 $\beta_{\cdot,c}$) to denote a $K \times 1$ vector for the m^{th} medication (and the c^{th} condition).

The k^{th} topic distribution for the m^{th} medication (or the c^{th} condition) is defined as the product of the corresponding topic embedding and medication (or condition) embedding followed by softmax normalization over the corresponding vocabularies:

$$\beta_{k,m}^{(med)} = \frac{\exp(\alpha_{k,\cdot}^{(med)} \rho_{\cdot,m}^{(med)})}{\sum_{m'} \exp(\alpha_{k,\cdot}^{(med)} \rho_{m'}^{(med)})}, \quad \beta_{k,c}^{(cond)} = \frac{\exp(\alpha_{k,\cdot}^{(cond)} \rho_{\cdot,c}^{(cond)})}{\sum_{c'} \exp(\alpha_{k,\cdot}^{(cond)} \rho_{c'}^{(cond)})} \quad (4)$$

488 Here we treat both the topic embeddings $\alpha^{(med)}$ and $\alpha^{(cond)}$ and condition and medication em-
489 beddings $\rho^{(med)}$ and $\rho^{(cond)}$ as point estimates without imposing any prior distribution.

For the ease of mathematical expression below, we denote $x_{m,d}^{(med)}$ (and $x_{c,d}^{(cond)}$) as the frequency of the m^{th} medication (and the c^{th} condition) for the d^{th} patient record. Formally, modeling the likelihood of the count frequency simply requires reformulating the above multinomial likelihood:

$$p(x_{m,d}^{(med)} | \theta_d, \Theta) = \prod_i [r_{m,d}^{(med)}]^{[w_{i,d}^{(med)}=m]} = [r_{m,d}^{(med)}]^{\sum_i [w_{i,d}^{(med)}=m]} = [r_{m,d}^{(med)}]^{x_{m,d}^{(med)}} \quad (5)$$

$$p(x_{c,d}^{(cond)} | \theta_d, \Theta) = \prod_j [r_{c,d}^{(cond)}]^{[w_{j,d}^{(cond)}=c]} = [r_{c,d}^{(cond)}]^{\sum_j [w_{j,d}^{(cond)}=c]} = [r_{c,d}^{(cond)}]^{x_{c,d}^{(cond)}} \quad (6)$$

490 where θ_d is the patient topic mixture of patient d and $\Theta = \{\alpha^{(med)}, \rho^{(med)}, \alpha^{(cond)}, \rho^{(cond)}\}$.

The vector $\mathbf{x}_{\cdot,d}^{(med)}$ and $\mathbf{x}_{\cdot,d}^{(cond)}$ denote the frequency of all medications and all conditions for the d^{th} patient record, respectively. The entire data can then be represented as a $M \times D$ matrix

$\mathbf{X}^{(cond)}$ and a $C \times D$ matrix $\mathbf{X}^{(cond)}$ and modelled by Eq (5) and Eq (6):

$$p(\mathbf{X}^{(med)}|\boldsymbol{\theta}, \Theta) = \prod_d \prod_m p(x_{m,d}^{(med)}|\boldsymbol{\theta}_d, \Theta) = \prod_d \prod_m [r_{m,d}^{(med)}]^{x_{m,d}^{(med)}} \quad (7)$$

$$p(\mathbf{X}_{c,d}^{(cond)}|\boldsymbol{\theta}, \Theta) = \prod_d \prod_c p(x_{c,d}^{(cond)}|\boldsymbol{\theta}_d, \Theta) = \prod_d \prod_c [r_{c,d}^{(cond)}]^{x_{c,d}^{(cond)}} \quad (8)$$

491 Inference of $\boldsymbol{\theta}_d$'s and learning of Θ are described next.

492 6.4 Model inference and estimation

To train GETM, we want to maximize the marginal likelihood of the individuals with respect to $\Theta = \{\boldsymbol{\alpha}^{(med)}, \boldsymbol{\rho}^{(med)}, \boldsymbol{\alpha}^{(cond)}, \boldsymbol{\rho}^{(cond)}\}$:

$$\boldsymbol{\alpha}^{(med)}, \boldsymbol{\rho}^{(med)}, \boldsymbol{\alpha}^{(cond)}, \boldsymbol{\rho}^{(cond)} \leftarrow \max_{\Theta} \mathcal{L}(\Theta) = \max_{\Theta} \sum_{d=1}^D \log p(\mathbf{x}_d^{(\cdot)}|\Theta) \quad (9)$$

where $\mathbf{x}_d^{(\cdot)}$ is word frequency vector (i.e., the bag of words of medications and conditions for individual d). This marginal likelihood is intractable to compute since it involves integrating out the logistic normal topic mixture variable $\boldsymbol{\theta}_d$, $\forall d \in \{1, \dots, D\}$:

$$\log p(\mathbf{x}_d^{(\cdot)}|\Theta) = \log \int p(\boldsymbol{\eta}_d) \prod_{m=1}^M p(x_{d,n}^{(med)}|\boldsymbol{\eta}_d, \boldsymbol{\alpha}^{(med)}, \boldsymbol{\rho}^{(med)}) \prod_{c=1}^C p(x_{d,n}^{(cond)}|\boldsymbol{\eta}_d, \boldsymbol{\alpha}^{(cond)}, \boldsymbol{\rho}^{(cond)}) d\boldsymbol{\eta}_d \quad (10)$$

$$= \log \int p(\boldsymbol{\eta}_d) \prod_{m=1}^M (f(\boldsymbol{\eta}_{d,\cdot}) f(\boldsymbol{\alpha}^{(med)} \boldsymbol{\rho}_{\cdot,m}^{(med)}))^{x_{m,d}^{(med)}} \prod_{c=1}^C (f(\boldsymbol{\eta}_{d,\cdot}) f(\boldsymbol{\alpha}^{(cond)} \boldsymbol{\rho}_{\cdot,c}^{(cond)}))^{x_{c,d}^{(cond)}} d\boldsymbol{\eta}_d \quad (11)$$

$$\equiv \log \int p(\boldsymbol{\eta}_d) \prod_{m=1}^M (\boldsymbol{\theta}_{d,\cdot} \boldsymbol{\beta}_{\cdot,m}^{(med)})^{x_{m,d}^{(med)}} \prod_{c=1}^C (\boldsymbol{\theta}_{d,\cdot} \boldsymbol{\beta}_{\cdot,c}^{(cond)})^{x_{c,d}^{(cond)}} d\boldsymbol{\eta}_d \quad (12)$$

493 where $f(\cdot)$ denotes the softmax function and $\boldsymbol{\theta}_{d,\cdot} \boldsymbol{\beta}_{\cdot,m}^{(med)}$ is the inner matrix product of the $1 \times K$
494 vector $\boldsymbol{\theta}_{d,\cdot}$ and the $K \times 1$ vector $\boldsymbol{\beta}_{\cdot,m}^{(med)}$. Same for conditions.

495 Therefore, we took an *Amortized Variational Inference* (AVI) approach to approximate the
496 above intractable integral. This is quite similar to the one described by Kingman & Welling
497 2014 [12] and by Dieng et al 2019 [8]. For the sake of completeness, we describe it below.

To approximate true posterior $p(\boldsymbol{\eta}_d|\mathbf{x}_d)$, we define the proposed distribution as a Gaussian with mean and variance produced by a feedforward neural network with parameters $\mathbf{W}_{\theta}$:

$$q(\boldsymbol{\eta}_d|\mathbf{x}_d, \mathbf{W}_{\theta}) = \mathcal{N}(\boldsymbol{\mu}_d, \text{diag}(\boldsymbol{\sigma}_d)\mathbf{I}), \quad [\boldsymbol{\mu}_d, \log \boldsymbol{\sigma}_d^2] = \text{NNET}(\mathbf{x}_d; \mathbf{W}_{\theta}) \quad (13)$$

The evidence lower bound (ELBO) of the above log-likelihood is approximated as follows:

$$\log p(\mathbf{x}_d^{(\cdot)}|\Theta) = \log \int q(\boldsymbol{\eta}_d) \frac{p(\boldsymbol{\eta}_d)}{q(\boldsymbol{\eta}_d)} \prod_{m=1}^M (\boldsymbol{\theta}_{d,\cdot} \boldsymbol{\beta}_{\cdot,m}^{(med)})^{x_{m,d}^{(med)}} \prod_{c=1}^C (\boldsymbol{\theta}_{d,\cdot} \boldsymbol{\beta}_{\cdot,c}^{(cond)})^{x_{m,d}^{(cond)}} d\boldsymbol{\eta}_d \quad (14)$$

$$\geq \int q(\boldsymbol{\eta}_d) \log p(\boldsymbol{\eta}_d) d\boldsymbol{\eta}_d \quad (15)$$

$$+ \int q(\boldsymbol{\eta}_d) \sum_{m=1}^M x_{m,d}^{(med)} \log (\boldsymbol{\theta}_{d,\cdot} \boldsymbol{\beta}_{\cdot,m}^{(med)}) d\boldsymbol{\eta}_d \quad (16)$$

$$+ \int q(\boldsymbol{\eta}_d) \sum_{c=1}^C x_{m,d}^{(cond)} \log (\boldsymbol{\theta}_{d,\cdot} \boldsymbol{\beta}_{\cdot,c}^{(cond)}) d\boldsymbol{\eta}_d - \int q(\boldsymbol{\eta}_d) \log q(\boldsymbol{\eta}_d) d\boldsymbol{\eta}_d$$

$$\begin{aligned} &= \mathbb{E}_{q(\boldsymbol{\eta})} [\log p(\boldsymbol{\eta}_d)] \\ &\quad + \mathbb{E}_{q(\boldsymbol{\eta})} [\log p(\mathbf{x}_d^{(med)} | \boldsymbol{\eta}_d, \boldsymbol{\beta}^{(med)})] \\ &\quad + \mathbb{E}_{q(\boldsymbol{\eta})} [\log p(\mathbf{x}_d^{(cond)} | \boldsymbol{\eta}_d, \boldsymbol{\beta}^{(cond)})] \\ &\quad - \mathbb{E}_{q(\boldsymbol{\eta}_d)} [\log q(\boldsymbol{\eta}_d)] \end{aligned} \quad (17)$$

$$= \sum_{t \in \{med, cond\}} \mathbb{E}_{q(\boldsymbol{\eta})} [\log p(\mathbf{x}_d^{(t)} | \boldsymbol{\eta}_d, \boldsymbol{\beta}^{(t)})] - KL[q(\boldsymbol{\eta}_d) || p(\boldsymbol{\eta}_d)] \quad (\text{ELBO}) \quad (18)$$

$$\approx \sum_{t \in \{med, cond\}} \log p(\mathbf{x}_d^{(t)} | \hat{\boldsymbol{\eta}}_d, \boldsymbol{\beta}^{(t)}) - KL[q(\hat{\boldsymbol{\eta}}_d) || p(\hat{\boldsymbol{\eta}}_d)] \quad (19)$$

where

$$\log p(\mathbf{x}_d^{(t)} | \hat{\boldsymbol{\eta}}_d, \boldsymbol{\beta}^{(t)}) = \sum_{m=1}^M x_{m,d}^{(med)} \log (\hat{\boldsymbol{\theta}}_{d,\cdot} \boldsymbol{\beta}_{\cdot,m}^{(med)}) + \sum_{c=1}^C x_{c,d}^{(cond)} \log (\hat{\boldsymbol{\theta}}_{d,\cdot} \boldsymbol{\beta}_{\cdot,c}^{(cond)})$$

$$\mathbb{E}_{q(\boldsymbol{\eta}_d | \mathbf{x}_d)} [\log p(\hat{\boldsymbol{\eta}}_d)] = \mathbb{E}_{q(\boldsymbol{\eta}_d | \mathbf{x}_d)} [\log \mathcal{N}(\hat{\boldsymbol{\eta}}_d; 0, 1)] = -\frac{1}{2}(\log(2\pi) + \boldsymbol{\mu}_d^2 + \boldsymbol{\sigma}_d^2)$$

$$\mathbb{E}_{q(\boldsymbol{\eta}_d | \mathbf{x}_d)} [\log q(\hat{\boldsymbol{\eta}}_d)] = -H(\hat{\boldsymbol{\eta}}) = -\frac{1}{2}(\log \boldsymbol{\sigma}_d^2 + 1 + \log(2\pi))$$

$$\mathbb{E}_{q(\boldsymbol{\eta}_d | \mathbf{x}_d)} [\log q(\hat{\boldsymbol{\eta}}_d) - \log p(\hat{\boldsymbol{\eta}}_d)] \equiv KL[q(\hat{\boldsymbol{\eta}}_d) || p(\hat{\boldsymbol{\eta}}_d)] = -\frac{1}{2}(\log \boldsymbol{\sigma}_d^2 + 1 - \boldsymbol{\mu}_d^2 - \boldsymbol{\sigma}_d^2)$$

Here Eq (19) uses the sampled $\hat{\boldsymbol{\eta}}_d$ from the feedforward network encoder as a surrogate to approximate the intractable variational expectation w.r.t. $\hat{\boldsymbol{\eta}}_d$. We re-parameterized the Gaussian distribution in Eq (13) to enable stochastic gradient calculation of the ELBO w.r.t. $\mathbf{W}_\theta$ and AVI training of the encoder network *NNET*:

$$q(\boldsymbol{\eta}_d | \mathbf{x}_d, \mathbf{W}_\theta) = \boldsymbol{\mu}_d + \text{diag}(\boldsymbol{\sigma}_d \mathcal{N}(0, \mathbf{I}), \quad [\boldsymbol{\mu}_d, \log \boldsymbol{\sigma}_d^2] = \text{NNET}(\mathbf{x}_d; \mathbf{W}_\theta) \quad (20)$$

$$\frac{\partial \text{ELBO}}{\partial \mathbf{W}_\theta} = \frac{\text{NNET}(\mathbf{x}_d; \mathbf{W}_\theta)}{\partial \mathbf{W}_\theta} \quad (21)$$

6.5 Knowledge graph construction for learning the embedding of UKB conditions and medications

In the original ETM [8], the word embedding ρ can be either learned from the documents or pre-trained on a separate large data corpus such as Wikipedia using the word2vec approach. The latter approach allows leveraging the contextual data to improve the topic modeling. As one of our main contributions, we sought to develop a simple framework that exploits ETM's ability to incorporate a pre-trained medication embedding $\rho^{(med)}$ and condition embedding $\rho^{(cond)}$ from their taxonomic knowledge graphs (KG). Leveraging the structural KG information and the internal relational information among medications and conditions in a principled way is beneficial to modeling the UKB phenotype data and other EHR data in general, which are often sparse, noisy, and bias.

Among the graph representational learning methods, we chose node2vec [9] as it is an unsupervised approach that can directly learn from the KG without linking to a prediction task as in other methods. Specifically, we applied node2vec to separately learn the node embedding of the 443 conditions or 802 medications from their hierarchical trees (**Fig. 1**).

For the KG of the 443 conditions, we constructed a tree graph using coding taxonomy designed by the UKB team (i.e., datafield 20002). The tree describes the topology of the conditions with 473 nodes and 4 hierarchical levels.

For the KG of the 802 medications, we constructed a KG based on the Anatomical Therapeutic Chemical (ATC) classification system [37]. The entire tree is composed of 5 levels. We first kept the top 4 levels of ATC, of which the first level contains main anatomical or pharmacological groups; the second level includes pharmacological or therapeutic subgroups; and in the third and fourth levels are chemical, pharmacological or therapeutic subgroups. To harmonize the UKB medications with the ATC graph, we mapped the names of active ingredients from UKB datafield 20003 to the fifth level codes of ATC, which are chemical substances. Some UKB medication codes were mapped to multiple ATC fifth level codes as they could belong to different subgroups depending on their usages. In that case, we replaced all of the mapped ATC fifth level codes with the same corresponding UKB medication code of the active ingredient in the second step. The final medication graph contains 2561 nodes in total.

Both the condition graph and the medication graph were treated as undirected graphs as input to the node2vec model (**Fig. 1**).

6.6 Overall learning algorithm of GETM

GETM learning algorithm is summarized in Algorithm 1.

6.7 Ablation experiments

To gain a good understanding of how each component of our method contribute to the overall improvements, we performed ablation experiments. In particular, we compared GETM with 8

Algorithm 1: Topic modelling with GETM

Input : UKB phenotype matrices $\mathbf{X}^{(cond)}$ and $\mathbf{X}^{(med)}$, node2vec-pretrained embedding for conditions and medications $\rho^{(cond)}$ and $\rho^{(med)}$

Initialize : condition-defined topic embedding $\alpha^{(cond)}$, medication-defined topic embedding $\alpha^{(med)}$, encoder neural network $\mathbf{W}_\theta$

Output : Learned topic embeddings $\hat{\alpha}^{(cond)}$, $\hat{\alpha}^{(med)}$ and encoder network parameters $\hat{\mathbf{W}}_\theta$ for patient topic mixtures

```

1 for epoch  $i = 1, 2, \dots$  do
2   for each batch  $\mathcal{B}$  do
3     Get normalized bag-of-words  $\tilde{\mathbf{X}}_{\mathcal{B}} = [\tilde{\mathbf{X}}_{\mathcal{B}}^{(cond)}, \tilde{\mathbf{X}}_{\mathcal{B}}^{(med)}]$ 
4     Sample  $\hat{\theta}_{\mathcal{B}}$  using  $NNET(\tilde{\mathbf{X}}_{\mathcal{B}}; \mathbf{W}_\theta)$  by Eq (21)
5     Compute approximate ELBO( $\hat{\theta}_{\mathcal{B}}, \alpha^{(med)}, \alpha^{(cond)}$ ) using Eq (19)
6     Update parameters  $\alpha^{(cond)}, \alpha^{(med)}, \mathbf{W}_\theta$  by backpropagation to maximize ELBO
7   end for
8 end for

```

ablated baseline models (**Supplementary Table S1**):

6.8 Model specifications

The encoders for GETM, partial GETMs and ETMs are 2-layer neural networks with hidden sizes of 128, ReLU activations [38] and 1D batch normalization [39]. The medication embedding dimension and condition embedding dimension were both 128 and they were fixed during training of the models with KG-informed pre-trained embeddings. The models were optimized with Adam optimizer at 0.01 learning rate. We trained each model with batch size of 100. For inferring individual-topic mixtures used by prediction tasks, the topic number was set to be 128.

6.9 Topic quality evaluation

For medication, the topic coherence was calculated as:

$$TC_{med} = \frac{1}{K} \sum_{k=1}^K \frac{m_k}{n} \quad (22)$$

where n is the number of top medications, $m_k \leq n$ is the maximum number of medications that are from the same category across all categories, and K is the number of topics. We set $n = 5$ because of the softmax-normalization of each topic, where only the top 5 conditions or medications under each topic has non-negligible probabilities.

To avoid overestimating the topic quality, the categories we used to evaluate the topic coherence were not processed from the ATC graph since it was involved in the pre-trained model. Instead, we employed 59 categories which are physician-curated and pain-focused (**Supplementary Table S6**).

For conditions, we did not have any external gold-standard reference to compare against. Therefore, we calculated the topic coherence for conditions based on the co-occurrence of the top conditions under the same topic observed from the same patient record. Specifically, for any top two conditions under the same topic, we calculated the average pointwise mutual information:

$$TC_{cond} = \frac{1}{K} \sum_{k=1}^K \frac{1}{n(n-1)} \sum_{i=1}^n \sum_{j=i+1}^n f(c_i^k, c_j^k) \quad (23)$$

where c_i^k is the i^{th} top most likely condition in topic k , and $f(\cdot, \cdot)$ is normalized pointwise mutual information:

$$f(c_i, c_j) = \frac{\log \frac{P(c_i, c_j)}{P(c_i)P(c_j)}}{-\log P(c_i, c_j)} \quad (24)$$

where $P(c_i, c_j)$ is the probability of condition i and condition j co-occurring in one individual and $P(c_i)$ is the marginal probability of condition i . The probabilities were approximated by empirical counts.

Topic diversity is the percentage of total unique features of the top 5 features of all topics. We first chose top 5 features from K topics. Among $5 \times K$ features there are U unique words. The topic diversity is calculated as $TQ = \frac{U}{5 \times K}$. Therefore, the higher topic diversity the more diverse the topics are.

6.10 Study of medication and condition relations

In GETM, since the encoder takes both medication and condition information as input (**Fig. 1**), the top medications and conditions from the same topic (i.e. the same index) are related. To quantitatively measure the ability of the model to capture known condition-medication relations, we first obtained all of the combinations from top 3 conditions and top 3 medications for each topic. We then counted the number of condition-medication pairs which are known to be related. The reference of known pairs was extracted from CTD <http://ctdbase.org/> and DrugBank <https://go.drugbank.com/>. We eventually mapped 222 conditions and 529 medications from UKB to these two databases. As a result, we had 2444 positive pairs of which the medication has treatment effects on the condition and 3231 negative pairs of which the condition belongs to the adverse effects of the medication. The number of known pairs discovered by our model and baselines were compared using different number of topic numbers (**Table 2**).

6.11 Data imputation

Using the generative model, we can (1) impute missing entries including conditions or medications and (2) impute medications based on only clinical conditions. Accordingly, two sets of experiments were performed to evaluate how well our model could complete these two tasks. To simulate missing entries, we randomly masked 50% of medications and conditions for test individuals. Then we calculated reconstruction error as the negative log-likelihood of the held-out

data:

$$-\log p(\mathbf{X}|\boldsymbol{\theta}_d, \boldsymbol{\beta}) = \frac{1}{D'} \sum_{d=1}^{D'} -\mathbf{x}_d \log(\hat{\mathbf{x}}_d^T) \quad (25)$$

where D' is the number of test individuals and $\hat{\mathbf{x}}_d = \boldsymbol{\theta}_d \boldsymbol{\beta}$

For medication imputation experiment, we masked the entire medication data of the test individuals and then reconstructing the medication matrix. This experiment mimics the scenario that we recommend relevant medications based only on individuals' condition history.

In addition to reconstruction error, we also calculated recall and precision rates at the top 5 predictions: recall@5 and precision@5. Specifically, we sorted the predicted probabilities of medications and chose top 5 medications, of which we calculated recall and precision with respect to medications that the individual took as true labels. The recall and precision are calculated as: $Recall = \frac{True\ Positives}{True\ Positives + False\ Negatives}$ and $Precision = \frac{True\ Positives}{True\ Positives + False\ Positives}$, respectively.

6.12 CMK pain prediction

CMK pain was defined as musculoskeletal pain lasting at least three months localized to any of five body sites: knee, back, neck/shoulder, face or hip. The label information was obtained from the UKB pain questionnaire as described in **Methods**.

We used individual topic mixture $\boldsymbol{\theta}_d$'s to predict whether a individual d has chronic musculoskeletal pain at the time of the visit as the binary outcome y_d . GETM was first trained with training data to get $\boldsymbol{\theta}^{(train)}$. Then we fit logistic regression model $\text{logit}(y_d) \sim \boldsymbol{\theta}_d^{(train)} \boldsymbol{\omega}$, where $\boldsymbol{\omega}$ is a $K \times 1$ vector of the linear coefficients corresponding to the predictive weights of the K topics.

To evaluate model performance, we first obtained $\boldsymbol{\theta}^{(test)}$ using the trained GETM (i.e., the encoder network parameters $\mathbf{W}_\theta$) and then predicted CMK pain status using the trained logistic regression coefficients $\hat{\mathbf{y}}^{(test)} = \boldsymbol{\theta}^{(test)} \hat{\boldsymbol{\omega}}$.

6.13 Odds ratios of CMK-related conditions or medications

We calculated odds ratios of each condition and medication with respect to CMK pain as the outcome (**Supplementary Fig. S2**). The odds ratios were used for two purposes: (1) obtaining conditions and medications that are obvious for chronic musculoskeletal pain prediction and to be removed for some experiments; (2) constructing a baseline reference to construct a curated list of pain-related conditions and medication by the physician in our team. For each feature (condition or medication), we formed a contingency table as below. The odds ratios were calculated as $OR = \frac{ad}{bc}$.

Feature \ Pain	Pain	
	Yes	No
Yes	a	b
No	c	d

6.14 Feature filtering schemes in removing obvious CMK-related conditions and medications

Let $\mathcal{C}$ and $\mathcal{M}$ be the complete sets of 443 conditions and 802 medications, respectively. Based on the odds ratios, we identified 50 conditions ($\mathcal{C}_1$) and 150 medications ($\mathcal{M}_1$) that are significantly associated with the CMK pain (**Supplementary Fig. S2**). A physician also provides a general pain related condition list ($\mathcal{C}_2$), a musculoskeletal pain related condition list ($\mathcal{C}_3$) and a general pain related medication set ($\mathcal{M}_2$). We created three filtered condition sets: (1) $\mathcal{C} - \mathcal{C}_2$; (2) $\mathcal{C} - \mathcal{C}_3$; (3) $\mathcal{C} - (\mathcal{C}_1 \cup \mathcal{C}_2)$, where using the setting notations $\mathcal{C} - \mathcal{C}_1$ denotes condition set without $\mathcal{C}_1$ and $\mathcal{C}_1 \cup \mathcal{C}_2$ denotes the union of $\mathcal{C}_1$ and $\mathcal{C}_2$.

We also created two filtered medication sets: (1) $\mathcal{M} - \mathcal{M}_1$; (2) $\mathcal{M} - (\mathcal{M}_1 \cup \mathcal{M}_2)$. These filtered condition sets and filtered medication sets formed six experiment settings plus the full set (**Supplementary Table S2**), which we used to explore CMK prediction performance as a function of reduced feature sets (**Fig. 4**).

6.15 Calculating importance of conditions and medications for CMK pain

Using the $K \times 1$ coefficients from logistic regression $\hat{\omega}$ and the topic embeddings, we calculated the relevance of medications and conditions to CMK pain as follows (**Fig. 5a,b**):

$$\mathbf{v} = \hat{\omega}^\top \hat{\beta} \quad (26)$$

We selected the top $N \in \{10, 30, 50\}$ medications and conditions from $\mathbf{v}^{(med)}$ and $\mathbf{v}^{(cond)}$. We then calculated the proportions of these top N medications or conditions overlapping with the pain-related lists created by a physician. The results are visualized as barplots in **Fig. 5c** and **Fig. 5d**.

6.16 Characterizing pain types from 7 body sites and pain all over the body

We used similar approaches as described for the CMK pain predictions to characterize other pain types. For all experiments, we used three different filtered feature sets of conditions and medications, which are (1) all: non-filtered conditions and medications, (2) general: physician-curated general pain-related conditions and medications were filtered, and (3) least: physician-curated / top odds ratio calculation conditions and medications were filtered.

For each pain type p , we calculated the importance weight that a medication or condition positively or negatively related to the pain using logistic regression coefficient for predicting that pain type $\hat{\omega}_p$ and GETM topic-feature mixture β of certain feature filtering regime (**Fig. 5(b)**).

Specifically, the way to calculate importance weight $w_{c,p}$ for a positive relation between a condition c and the pain type p is:

$$\hat{w}_{c,p}^{(+)} = \begin{cases} 0 & \text{if } \hat{w}_{c,p} < 0 \\ \hat{w}_{c,p} & \text{otherwise} \end{cases}, \quad w_{c,p}^{(+)} = \frac{[\hat{\mathbf{w}}_p^{(+)}]^\top \boldsymbol{\beta}_{c,\cdot}^{(cond)}}{\sum_c \hat{w}_{c,p}^{(+)}} \quad (27)$$

Similarly, we calculated importance weight for negative relationship between a condition c and pain type p :

$$\hat{w}_{c,p}^{(-)} = \begin{cases} 0 & \text{if } \hat{w}_{c,p} > 0 \\ \hat{w}_{c,p} & \text{otherwise} \end{cases}, \quad w_{c,p}^{(-)} = \frac{[\hat{\mathbf{w}}_p^{(-)}]^\top \boldsymbol{\beta}_{c,\cdot}^{(cond)}}{\sum_c \hat{w}_{c,p}^{(-)}} \quad (28)$$

641 Calculation for the positive and negative relations between every medication and every pain
642 type is the same.

7 Figures

Figure 1: Overview of Graph-embedded topic model and its application on UKB phenotype data. The UKB data consists of 443 conditions and 802 medications for 457,461 individuals. We developed GETM to model these data although our GETM can be applied to other datasets as well. **(a)** GETM training. GETM is a variational autoencoder (VAE) model. The neural network encoder takes individuals' condition and medication information as input and produces the variational mean μ and variance σ^2 for the patient topic mixtures θ . The decoder is linear and consists of two tri-factorizations. One learns medication-defined topic embedding $\alpha^{(med)}$ and medication embedding $\rho^{(med)}$. The other learns condition-specific topic embedding $\alpha^{(cond)}$ and the condition embedding $\rho^{(cond)}$. We separately pre-train **(b)** the embedding of medications $\rho^{(med)}$ and **(c)** the embedding of conditions $\rho^{(cond)}$ using node2vec [9] based on their structural meta-information. This is done learning the node embedding that maximizes the likelihood of the tree-structured relational graphs of conditions and medications.

Figure 2: Topic quality visualization. The analysis was based on results from GETM model with both condition and medication embedding that are KG-informed using 75 topics (**Supplementary Table S1**). We applied UMAP to visualize (a) the condition embedding $\rho^{(cond)}$ and topic embedding $\alpha^{(cond)}$ and (b) medication embedding $\rho^{(med)}$ and topic embedding $\alpha^{(med)}$. The points are colored by corresponding category of condition and medication or annotated by its topic number. (c) & (d) Heatmap visualization of select topics. We randomly chose 5 as labelled in panel (a) and (b). We then generated heatmaps for the top conditions and the top medications with the highest probability $\beta_k^{(cond)}$ and $\beta_k^{(med)}$ under each of the five topics. The color intensity display proportionally the probabilities. The color bars on the right indicate the categories of the conditions and medications. All four panels share the same color legend of the categories.

(a)

(b)

Figure 3: Illustration of GETM for medication imputation through examples of study subjects. (a) The three best-matched study subjects displayed as 3 barplot panels on the left. We chose 3 best-matched individuals for whom our imputed medications matched well with the observed medications. (b) The three worst-matched individuals displayed as 3 barplot panels on the right. Y-axis represent the predicted probability of the medications on the x-axis. The numbers on the top of each bar represents the patient's observed conditions that are known to be treated by the corresponding medication under the bar. Inset in each panel lists the conditions names corresponding to the numbers on top of each bar. Blue bars indicate observed medications in the patient, and red bars indicate unobserved medications.

Figure 4: **Performance of logistic regression (LR) for CMK pain.** We trained LR models using θ obtained from GETM (GETM, **Supplementary Table S1**) with 128 topics as input to predict CMK pain. The baseline LR model directly used the raw condition and medication data as input. We iterated through seven feature filtering schemes with different filtered condition sets and medication sets as indicated by x-axis (details in **Supplementary Table S2** and described in **Methods** Section 6.12). Barplot displays the AUROC and AUPRC across these experiments.

Figure 5: Analysis of CMK-pain-related conditions and medications. (a) Logistic regression using patient topic mixture $\boldsymbol{\theta} \in \mathbb{R}^{D \times K}$ of D individuals and K topics to predict CMK pain as a binary outcome. (b) Importance score computation for medications and conditions. Taking the inner product of the regression coefficients $\boldsymbol{\omega}^T \in \mathbb{R}^{1 \times K}$ and $\boldsymbol{\beta}^{(med)} \in \mathbb{R}^{K \times M}$ from GETM (GETM, **Supplementary Table S2**), we obtained the importance scores of predicting CMK pain as $\mathbf{v}^{(med)} \in \mathbb{R}^{1 \times M}$ and $\mathbf{v}^{(cond)} \in \mathbb{R}^{1 \times C}$ for medications and conditions, respectively. (c) Proportions of known CMK-related conditions based on a physician-compiled list. (d) Proportions of known CMK-related medications based on a physician-compiled list of analgesic medications. Comparisons relate to two baselines: (1) Using ETM which treated conditions and medications as the same type of feature (i.e., ETM-both-flat, **Supplementary Table S1**). For this baseline, we selected the top medications and conditions from the resulting $\mathbf{v} \in \mathbb{R}^{1 \times (M+C)}$. (2) Odds Ratio (**Methods** Section 6.13).

Figure 6: Topic analysis for CMK pain. (a) The most predictive CMK pain topics. Based on the logistic regression coefficients of predicting CMK pain ω (Fig. 5a), we chose three topics with the highest coefficients and three topics with the most negative coefficients. Each bar is composed of condition and medication category. (b) & (c) The top conditions and medications under the 6 most predictive topics. Same as in Fig. 2, the color bars on the right indicate the categories of the conditions and medications.

Figure 7: Prediction performance of 10 pain types. Logistic regression was trained to predict 10 pain types as indicated in the x-axis. We experimented with three different filtered feature sets of conditions and medications: (1) all (m802c443): non-filtered conditions and medications, (2) general (m680c351): physician-curated general pain-related conditions and medications were filtered, and (3) least: physician-curated conditions and medications based on with odds ratios were filtered. Details of the data filtering were described in **Supplementary Table S2**.

Figure 8: **Contributions of cardiovascular-related conditions and medications to 12 pain types.** The importance weights of cardiovascular conditions and medications in predicting the 12 pain types were calculated by the sum of their topic probabilities weighted by the linear regression coefficients across all topics (**Methods**). Descriptions of the names of these pain types are in **Supplementary Table S3**. The results for cariovascular category were based on the filtered sets, where the known pain-related conditions and medications based on a physician curated list were removed. The results of analgesic category (anlgsc_med) were based on the non-filtered set, as only full set contains analgesic medications. All of the importance weights from other filtering schemes were displayed in **Supplementary Fig. S3**.

8 Tables

Topic #	15	50	75	100	15	50	75	100	15	50	75	100
Algorithm	Topic Coherence				Topic Diversity				Topic Quality			
ETM-medOnly	0.2800	0.332	0.3173	0.346	0.8800	0.432	0.2533	0.298	0.2464	0.1434	0.0804	0.1031
GETM-medOnly	0.4933	0.7160	0.7040	0.780	1.0000	0.832	0.7120	0.574	0.4933	0.5957	0.5012	0.4477
ETM-both-multi	0.3200	0.300	0.3494	0.346	0.8533	0.504	0.2720	0.236	0.2731	0.1512	0.0950	0.0817
GETM-conEmbOnly	0.3733	0.452	0.3307	0.406	0.9600	0.552	0.4507	0.374	0.3584	0.2495	0.1490	0.1518
GETM-medEmbOnly	0.4933	0.7000	0.7707	0.7860	0.9867	0.772	0.6053	0.488	0.4867	0.5404	0.4665	0.3836
GETM	0.5200	0.7040	0.7200	0.722	0.9600	0.856	0.6400	0.550	0.4992	0.6026	0.4608	0.3971

Table 1: **Medication-specific topic evaluation.** We have calculated topic coherence, topic diversity and topic quality in terms of medications for 6 algorithms using different topic numbers (model details in **Supplementary Table S1**. Condition-specific topic evaluation is in **Supplementary Table S4**.

Topic #	15	50	75	100
Algorithm	Number of matched pairs			
ETM-both-flat	53	84	119	132
ETM-both-multi	63	86	105	126
GETM-condEmbOnly	62	118	159	162
GETM-medEmbOnly	65	135	171	178
GETM	61	161	175	203

Table 2: **Number of known pairs between conditions and medications.** We mapped our medications and conditions identifiers to CTD and DrugBank databases to obtain known connections between them. For each topic inferred, we generated 9 condition-medication pairs from the top 3 conditions and the top 3 medications. Among these pairs, we calculated the number of known pairs (**Supplementary Table S1**). Higher numbers imply more clinically meaningful topics inferred by the algorithm.

Topic #	15	50	75	100
Algorithm	Reconstruction Error			
ETM-condOnly	5.89	5.56	5.39	5.80
GETM-condOnly	5.76	5.33	5.05	4.93
ETM-both_multi	5.08	5.09	4.80	4.96
GETM-condEmbOnly	5.12	4.84	4.55	4.40
GETM-medEmbOnly	5.06	4.86	4.75	4.58
GETM	4.69	4.84	4.45	4.34
Proportion	8.69			

Table 3: **Reconstructing 50% masked conditions.** The 50% of test data was randomly masked such that each test patient would have 50% of their conditions and medications observed. We then reconstructed the matrix with learned θ , α and ρ . The reconstruction error (i.e. negative log-likelihood) was calculated for the held-out data. Same condition data was used for all algorithms. For ETM-both-multi, GETM-condEmbOnly, GETM-medEmbOnly and GETM, the same models and data were used as for the medication reconstruction **Supplementary Table S5**. Description of the algorithm names are in **Supplementary Table S1**. As another baseline method, we evaluated the performance of filling in masked conditions based on their overall proportion over all of the UKB population.

Topic #	15	50	75	100
Algorithm	Medication Recovery Error			
Upper bound	12.25	11.24	11.50	11.50
ETM-both-flat	18.51	19.85	19.81	20.06
GETM-both-multi	14.99	14.88	14.94	15.05
GETM-condEmbOnly	15.24	14.95	14.95	14.96
GETM-medEmbOnly	15.18	14.88	14.90	14.88
GETM	14.82	14.65	14.61	14.65

Table 4: **Medication imputation.** The medication data was masked for test individuals. In contrast to the medication reconstruction experiments (**Supplementary Table S5**), we masked the entire medications for the test patients. We imputed their medication data using the inferred θ from their condition data only and the learned embedding $\alpha^{(med)}$ and $\rho^{(med)}$. The reconstruction error (i.e. negative log-likelihood) was calculated. The upper bounds were obtained using reconstruction errors calculated from unmasked test data using GETM (GETM, **Supplementary Table S1**). Description of the algorithm names are in **Supplementary Table S1**.

Algorithm \ Topic #	15	50	75	100	15	50	75	100
	recall@5				precision@5			
Upper bound	0.6672	0.7943	0.8027	0.7862	0.3342	0.4084	0.4226	0.4049
ETM-both-flat	0.4397	0.4486	0.4667	0.4614	0.1901	0.2109	0.2111	0.2162
GETM	0.5543	0.5639	0.5568	0.5388	0.2479	0.2516	0.2504	0.2417
GETM-condEmbOnly	0.5479	0.5664	0.5670	0.5647	0.2373	0.2524	0.2493	0.2486
GETM-medEmbOnly	0.5519	0.5722	0.5668	0.5716	0.2440	0.2533	0.2521	0.2532
GETM	0.5692	0.5787	0.5732	0.5753	0.2504	0.2612	0.2606	0.2578

Table 5: **Medication imputation accuracy.** The imputation procedures were the same as in Table 4. We calculated the precision and recall at the top 5 imputed medications for each individual. The upper bounds were calculated from unmasked test data using GETM (GETM, Supplementary Table S1).

References

1. Thomas H McCoy, Victor M Castro, Leslie A Snapper, Kamber L Hart, and Roy H Perlis. Efficient genome-wide association in biobanks using topic modeling identifies multiple novel disease loci. *Molecular Medicine*, 23(1):285–294, 2017.
2. Peter B Jensen, Lars J Jensen, and Søren Brunak. Mining electronic health records: towards better research applications and clinical care. *Nature Reviews Genetics*, 13(6):395–405, 2012.
3. Clare Bycroft, Colin Freeman, Desislava Petkova, Gavin Band, Lloyd T Elliott, Kevin Sharp, Allan Motyer, Damjan Vukcevic, Olivier Delaneau, Jared O’Connell, Adrian Cortes, Samantha Welsh, Alan Young, Mark Effingham, Gil McVean, Stephen Leslie, Naomi Allen, Peter Donnelly, and Jonathan Marchini. The uk biobank resource with deep phenotyping and genomic data. *Nature*, 562(7726):203–209, 10 2018.
4. Clare Turnbull, Richard H Scott, Ellen Thomas, Louise Jones, Nirupa Murugaesu, Freya Boardman Pretty, Dina Halai, Emma Baple, Clare Craig, Angela Hamblin, et al. The 100 000 genomes project: bringing whole genome sequencing to the nhs. *Bmj*, 361, 2018.
5. David M. Blei, Andrew Y. Ng, and Michael I. Jordan. Latent dirichlet allocation. *J. Mach. Learn. Res.*, 3(null):993–1022, March 2003.
6. Y. Li, Pratheeksha Nair, Xing Han Lu, Zhi Wen, Yuening Wang, Amir Ardalani Kalantari Dehaghi, Yanchun Miao, Weiqi Liu, T. Ordog, J. Biernacka, E. Ryu, J. Olson, M. Frye, Aihua Liu, Liming Guo, A. Marelli, Y. Ahuja, J. Davila-Velderrain, and Manolis Kellis. Inferring multimodal latent topics from electronic health records. *Nature Communications*, 11, 2020.
7. Ziyang Song, Xavier Sumba Toral, Yixin Xu, Aihua Liu, Liming Guo, Guido Powell, Aman Verma, David Buckeridge, Ariane Marelli, and Yue Li. Supervised multi-specialist topic model with applications on large-scale electronic health record data. In *Proceedings of the 12th ACM Conference on Bioinformatics, Computational Biology, and Health Informatics*, BCB ’21, New York, NY, USA, 2021. Association for Computing Machinery.
8. Adji B. Dieng, Francisco J. R. Ruiz, and David M. Blei. Topic modeling in embedding spaces. *CoRR*, abs/1907.04907, 2019.
9. Aditya Grover and Jure Leskovec. node2vec: Scalable feature learning for networks. *CoRR*, abs/1607.00653, 2016.
10. A Fayaz, P Croft, R M Langford, L J Donaldson, and G T Jones. Prevalence of chronic pain in the uk: a systematic review and meta-analysis of population studies. *BMJ Open*, 6(6), 2016.

11. M. Dueñas, B. Ojeda, A. Salazar, J. Micó, and I. Failde. A review of chronic pain impact on patients, their social environment and the health care system. *Journal of Pain Research*, 9:457 – 467, 2016.
12. Diederik P Kingma and Max Welling. Auto-encoding variational bayes, 2014.
13. DRUGBANK kernel description. <https://go.drugbank.com/>. Accessed: 2021-07-31.
14. Andrzej Wróbel, Anna Serefko, Andrzej Woźniak, Jacek Kociszewski, Aleksandra Szopa, Radosław Wiśniewski, and Ewa Poleszak. Duloxetine reverses the symptoms of overactive bladder co-existing with depression via the central pathways. *Pharmacology Biochemistry and Behavior*, 189:172842, 2020.
15. Crystian B. Oliveira, Chris G. Maher, Marcia R. Franco, Steven J. Kamper, Christopher M. Williams, Fernanda G. Silva, and Rafael Z. Pinto. Co-occurrence of Chronic Musculoskeletal Pain and Cardiovascular Diseases: A Systematic Review with Meta-analysis. *Pain Medicine (Malden, Mass.)*, 21(6):1106–1121, June 2020.
16. Sameer Al-Ghamdi, Mamdouh M. Shubair, Ashraf El-Metwally, Majid Alsalamah, Saeed Mastour Alshahrani, Badr F Al-Khateeb, Salwa Bahkali, Sara M. Aloudah, Jamaan Al-Zahrani, Turkey H. Almigbal, and Khaled K. Aldossari. The relationship between chronic pain, prehypertension, and hypertension. a population-based cross-sectional survey in al-kharj, saudi arabia. *Postgraduate Medicine*, 133(3):345–350, 2021. PMID: 33317375.
17. Pankaj Baral, Swalpa Udit, and I. Chiu. Pain and immunity: implications for host defence. *Nature Reviews Immunology*, pages 1–15, 2019.
18. Balraj S. Heran, Michelle My Wong, Inderjit K. Heran, and James M. Wright. Blood pressure lowering efficacy of angiotensin converting enzyme (ACE) inhibitors for primary hypertension. *The Cochrane Database of Systematic Reviews*, (4):CD003823, October 2008.
19. S. Ghione. Hypertension-associated hypalgesia. Evidence in experimental animals and humans, pathophysiological mechanisms, and potential clinical consequences. *Hypertension (Dallas, Tex.: 1979)*, 28(3):494–504, September 1996.
20. Elena Makovac, Giuseppina Porciello, Daniela Palomba, Barbara Basile, and Cristina Ottaviani. Blood pressure-related hypoalgesia: a systematic review and meta-analysis. *Journal of Hypertension*, 38(8):1420–1435, August 2020.
21. Qing Chen, Wenxin Zhang, Neeti Sadana, and Xinzhong Chen. Estrogen receptors in pain modulation: cellular signaling. *Biology of Sex Differences*, 12(1):22, February 2021.

22. R. Groen, O. Ryan, J. Wigman, H. Riese, B. Penninx, E. Giltay, M. Wichers, and C. Hartman. Comorbidity between depression and anxiety: assessing the role of bridge mental states in dynamic psychological networks. *BMC Medicine*, 18, 2020.
23. Jia Song, M. Zeng, Hai Wang, C. Qin, H. Hou, Zi-Yong Sun, San-Peng Xu, Guo ping Wang, Cuilian Guo, Yi ke Deng, Zhi chao Wang, J. Ma, L. Pan, B. Liao, Zhi-Hui Du, Q. Feng, Y. Liu, Jun-Gang Xie, and Z. Liu. Distinct effects of asthma and copd comorbidity on disease expression and outcome in patients with covid-19. *Allergy*, 76:483 – 496, 2020.
24. Umesh Singh, Victoria Wangia-Anderson, and J. Bernstein. Chronic rhinitis is a high-risk comorbidity for 30-day hospital readmission of patients with asthma and chronic obstructive pulmonary disease. *The journal of allergy and clinical immunology. In practice*, 7 1:279–285.e6, 2019.
25. Thomas N Kipf and Max Welling. Semi-Supervised Classification with Graph Convolutional Networks. *arXiv.org*, cs.LG, 09 2016.
26. Michael Schlichtkrull, Thomas N Kipf, Peter Bloem, Rianne Van Den Berg, Ivan Titov, and Max Welling. Modeling relational data with graph convolutional networks. In *European semantic web conference*, pages 593–607. Springer, 2018.
27. Alan Fayaz, S. A. Ayis, Sukhmeet Panesar, Richard Langford, and Liam Donaldson. Assessing the relationship between chronic pain and cardiovascular disease: A systematic review and meta-analysis. *Scandinavian Journal of Pain*, 13:76–90, 10 2016.
28. Crystian B Oliveira, Chris G Maher, Marcia R Franco, Steven J Kamper, Christopher M Williams, Fernanda G Silva, and Rafael Z Pinto. Co-occurrence of chronic musculoskeletal pain and cardiovascular diseases: A systematic review with meta-analysis. *Pain Medicine*, 21(6):1106–1121, oct 2019.
29. Heberto Suarez-Roca, Rebecca Y. Klinger, Mihai V. Podgoreanu, Ru-Rong Ji, Martin I. Sigurdsson, Nathan Waldron, Joseph P. Mathew, and William Maixner. Contribution of Baroreceptor Function to Pain Perception and Perioperative Outcomes. *Anesthesiology*, 130(4):634–650, April 2019.
30. Stephen Bruehl and Ok Yung Chung. Interactions between the cardiovascular and pain regulatory systems: an updated review of mechanisms and possible alterations in chronic pain. *Neuroscience and Biobehavioral Reviews*, 28(4):395–414, July 2004.
31. Jeffrey L. Jackson, Akira Kuriyama, Yachiyo Kuwatsuka, Sarah Nickoloff, Derek Storch, Wilkins Jackson, Zhi-Jiang Zhang, and Yasuaki Hayashino. Beta-blockers for the prevention of headache in adults, a systematic review and meta-analysis. *PloS One*, 14(3):e0212785, 2019.

32. Inna E. Tchivileva, Pei Feng Lim, Shad B. Smith, Gary D. Slade, Luda Diatchenko, Samuel A. McLean, and William Maixner. Effect of catechol-O-methyltransferase polymorphism on response to propranolol therapy in chronic musculoskeletal pain: a randomized, double-blind, placebo-controlled, crossover pilot study. *Pharmacogenetics and Genomics*, 20(4):239–248, April 2010.
33. Inna E. Tchivileva, Holly Hadgraft, Pei Feng Lim, Massimiliano Di Giosia, Margarete Ribeiro-Dasilva, John H. Campbell, Janet Willis, Robert James, Marcus Herman-Giddens, Roger B. Fillingim, Richard Ohrbach, Samuel J. Arbes, and Gary D. Slade. Efficacy and safety of propranolol for treatment of temporomandibular disorder pain: a randomized, placebo-controlled clinical trial. *Pain*, 161(8):1755–1767, August 2020.
34. Uk biobank data-field 20002: non-cancer illness code, self-reported. <https://biobank.ctsu.ox.ac.uk/crystal/field.cgi?id=20002>. Accessed: 2021-07-31.
35. Uk biobank data-field 20003: treatment/medication code. <https://biobank.ctsu.ox.ac.uk/crystal/field.cgi?id=20003>. Accessed: 2021-07-31.
36. David M. Blei, Michael I. Jordan, Thomas L. Griffiths, and Joshua B. Tenenbaum. Hierarchical topic models and the nested chinese restaurant process. In *Proceedings of the 16th International Conference on Neural Information Processing Systems*, NIPS'03, page 17–24, Cambridge, MA, USA, 2003. MIT Press.
37. WHO Collaborating Centre for Drug Statistics Methodology kernel description. <https://www.whocc.no/>. Accessed: 2021-07-31.
38. Xavier Glorot, Antoine Bordes, and Yoshua Bengio. Deep sparse rectifier neural networks. In *AISTATS*, 2011.
39. Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *CoRR*, abs/1502.03167, 2015.

A graph-embedded topic model enables characterization of diverse pain phenotypes among UK Biobank individuals

Supplementary Information

**Yuening Wang¹, Rodrigo Benavides², Luda Diatchenko^{3,4,5}, Audrey V. Grant^{3,4,5,*},
Yue Li^{1,*}**

¹School of Computer Science, McGill University, ²Department of Anesthesiology,
Centro Nacional de Rehabilitación, San Jose, Costa Rica, ³Department of
Anesthesia, McGill University, ⁴Faculty of Dentistry, McGill University, ⁵Department of
Anesthesia, McGill University

*Correspondence: audrey.grant@mcgill.ca, yueli@cs.mcgill.ca

1 Supplementary Figures

Figure S3: Importance weights computed using non-filtered conditions and medications over all pain types. For each pain type p , we calculated the importance weight that a medication or condition positively or negatively related to the pain using logistic regression coefficients for predicting that pain type $\hat{\omega}_p$ and GETM topic-feature mixture β (**Methods** 6.16). The physician-curated general pain-related conditions and medications were filtered.

Figure S4: Importance weights computed using physician-curated general pain-related filtered conditions and medications over all pain types. For each pain type p , we calculated the importance weight that a medication or condition positively or negatively related to the pain using logistic regression coefficients for predicting that pain type $\hat{\omega}_p$ and GETM topic-feature mixture β (Methods 6.16).

2 Supplementary Tables

Model Name	Model Description
ETM-condOnly	ETM modeling conditions only without node2vec pre-trained embedding
ETM-medOnly	ETM modeling medications only without pre-trained embedding
ETM-both-flat	Baseline ETM that takes both conditions and medications but treats them as the same type of features (i.e., single-view topics).
ETM-both-multi	ETM without using node2vec embedding but infers medication-specific and condition-specific topics (i.e., multi-view topics)
GETM-condOnly	GETM modeling only UKB conditions using the pre-trained node2vec embedding
GETM-medOnly	GETM modeling only UKB medications using the pretrained node2vec embedding
GETM-conEmbOnly	GETM using node2vec-embedding on condition but not on medications
GETM-medEmbOnly	GETM using node2vec-embedding on medication but not on conditions
GETM	The proposed full GETM model with node2vec-embeddings on both the conditions and medications

Table S1: **The list of ablation models.** Abbreviation and description of the ablated models. All of the GETM models use the node2vec-embeddings pretrained on conditions or medications graphs. All of the ETM models do not use the node2vec-embeddings.

Set #	Set Name	Set Description
1	m802c443	containing all 802 medications; containing all 443 conditions
2	m680c389	removing physician-curated general pain-related medications; removing physician-curated musculoskeletal pain-related conditions
3	m680c351	removing physician-curated general pain-related medications; removing physician-curated general pain-related conditions
4	m680c322	removing physician-curated general pain-related medications; removing physician-curated general pain-related conditions and top conditions correlated with chronic musculoskeletal pain from odds ratio calculation
5	m579c389	removing physician-curated general pain-related medications and top medications correlated with chronic musculoskeletal pain from odds ratio calculation; removing physician-curated musculoskeletal pain-related conditions
6	m579c351	removing physician-curated general pain-related medications and top medications correlated with chronic musculoskeletal pain from odds ratio calculation; removing physician-curated general pain-related conditions
7	m579c322	removing physician-curated general pain-related medications and top medications correlated with chronic musculoskeletal pain from odds ratio calculation; removing physician-curated general pain-related conditions and top conditions correlated with chronic musculoskeletal pain from odds ratio calculation
8	m666c340	removing physician-curated general pain-related medications and top medications correlated with chronic neck/shoulder pain from odds ratio calculation; removing physician-curated general pain-related conditions and top conditions correlated with chronic neck/shoulder pain from odds ratio calculation
9	m669c343	removing physician-curated general pain-related medications and top medications correlated with chronic hip pain from odds ratio calculation; removing physician-curated general pain-related conditions and top conditions correlated with chronic hip pain from odds ratio calculation
10	m674c341	removing physician-curated general pain-related medications and top medications correlated with chronic back pain from odds ratio calculation; removing physician-curated general pain-related conditions and top conditions correlated with chronic back pain from odds ratio calculation
11	m661c336	removing physician-curated general pain-related medications and top medications correlated with chronic stomachache/abdominal pain from odds ratio calculation; removing physician-curated general pain-related conditions and top conditions correlated with chronic stomachache/abdominal pain from odds ratio calculation
12	m671342	removing physician-curated general pain-related medications and top medications correlated with chronic knee pain from odds ratio calculation; removing physician-curated general pain-related conditions and top conditions correlated with chronic knee pain from odds ratio calculation
13	m670c337	removing physician-curated general pain-related medications and top medications correlated with chronic headache from odds ratio calculation; removing physician-curated general pain-related conditions and top conditions correlated with chronic headache from odds ratio calculation
14	m662c338	removing physician-curated general pain-related medications and top medications correlated with chronic facial pain from odds ratio calculation; removing physician-curated general pain-related conditions and top conditions correlated with chronic facial pain from odds ratio calculation
15	m665c348	removing physician-curated general pain-related medications and top medications correlated with chronic all-body pain from odds ratio calculation; removing physician-curated general pain-related conditions and top conditions correlated with chronic all-body pain from odds ratio calculation
16	m605c284	removing physician-curated general pain-related medications and top medications correlated with chronic pain* from odds ratio calculation; removing physician-curated general pain-related conditions and top conditions correlated with chronic pain* from odds ratio calculation
17	m640c332	removing physician-curated general pain-related medications and top medications correlated with acute pain from odds ratio calculation; removing physician-curated general pain-related conditions and top conditions correlated with acute pain from odds ratio calculation

Table S2: List of runs conducted among UKB European descent study subjects including features from selected combinations of subsets of conditions (from a total of 443) and medications (from a total of 802).

Label Name	Description
acute	acute pain on any of body sites
acute2chronic	acute pain on any of musculoskeletal body site at first visit and became chronic pain diagnosed at the following visits
chronic	chronic pain on any of body sites
back	chronic back pain
neck_shlder	chronic neck/shoulder pain
hip	chronic hip pain
knee	chronic knee pain
face	chronic facial pain
musk_pain	chronic musculoskeletal pain
headch	chronic headache
stom_abdmnl	chronic stomachache or abdominal pain
all	chronic pain on all over the body (all_over_body)

Table S3: **Pain type label description.** The description for pain label name. Musculoskeletal pain sites includes face, knee, back, neck/shoulder and hip.

Topic #	15	50	75	100	15	50	75	100	15	50	75	100
Algorithm	Topic Coherence				Topic Diversity				Topic Quality			
ETM-condOnly	0.0035	0.0030	0.0019	0.0076	0.6267	0.332	0.2160	0.154	0.0022	0.0010	0.0004	0.0012
GETM-condOnly	0.0253	0.0153	0.0172	0.0170	0.9067	0.656	0.5707	0.476	0.0229	0.0100	0.0098	0.0081
ETM-both-multi	0.0125	0.0105	0.0044	0.0046	0.9200	0.420	0.2400	0.224	0.0115	0.0044	0.0011	0.0010
GETM-conEmbOnly	0.0213	0.0109	0.0193	0.0169	0.9467	0.728	0.5760	0.472	0.0202	0.0079	0.0111	0.0080
GETM-medEmbOnly	0.0144	0.0070	0.0087	0.0087	0.8267	0.576	0.4640	0.368	0.0119	0.0040	0.0040	0.0032
GETM	0.0206	0.0188	0.0186	0.0186	0.9333	0.760	0.6293	0.496	0.0192	0.0143	0.0117	0.0092

Table S4: **Condition-defined topic quality.** We have calculated topic coherence, topic diversity and topic quality in terms of conditions for 6 algorithms using different topic numbers. The description of algorithm name is in **Supplementary Table S1**.

Topic #	15	50	75	100
Algorithm	Reconstruction Error			
ETM-medOnly	8.54	8.24	8.10	7.56
GETM-medOnly	7.84	7.46	6.13	6.99
ETM-both-multi	7.35	7.40	7.63	6.66
GETM-condEmbOnly	6.70	6.85	6.83	6.74
GETM-medEmbOnly	7.71	7.02	6.98	6.28
GETM	6.34	5.70	5.69	5.53
Proportion	12.23			

Table S5: **Imputation error of masked medications.** The 50% of test data was randomly masked. Then we reconstructed the matrix with learned θ , α and ρ . The reconstruction error (i.e. negative log-likelihood) was calculated for the held-out data. Same medication data was used for all algorithms. Same models and same data were used as in **Table 3** for ETM-both-multi, GETM-condEmbOnly, GETM-medEmbOnly and GETM. Description of algorithm names are in **Supplementary Table S1**. As another baseline method, we evaluated the performance of filling in masked conditions based on their overall proportion over all of the UKB population.

Column 1	Column 2
evening primrose oil	multivitamins minerals and combinations
antiadrenergic agents	opioid - non-opioid combinations
dopaminergic agents	oral hypoglycemics
insulins and analogues	anxiolytics hypnotics and sedatives
antifungals	vasodilators
antispasmodics	decongestants and antiallergics
antivertigo preparations	chondroitin
topical products for joint and muscular pain	cardiac glycosides
laxatives emollients and protectives	hormones and related agents
antimigraine preparations	antiepileptics
beta blocking agents	fish oil
diuretics	antipsoriatics
other topical preparations	nsaid
parasympathomimetics	ginkgo biloba
anticholinergic agents	calcium
glucosamine	paracetamol
calcium channel blockers	antidepressants
bisphosphonates	all other therapeutic products
iron and other antianemic preparations	antigout preparations
thyroid preparations	antacids drugs for peptic ulcer and gerd
angiotensin ii receptor blockers	corticosteroids
antithrombotic agents	other psychotherapeutic agents
antiarrhythmics - i and iii	antimalarials
intestinal antiinflammatory agents	antibacterials
estrogens progestogens and combinations	ace inhibitors
immunosuppressants	antivirals
urologicals	opioids
lipid modifying agents	salicylic acid and derivatives
ophthalmologicals	adrenergic agents
antihistamines	

Table S6: **External medication categories** All 59 external medication categories are listed.