

Predicting Visit Cost of Obstructive Sleep Apnea Using Electronic Healthcare Records With Transformer

ZHAOYANG CHEN¹, LINA SILTALA-LI¹, MIKKO LASSILA¹, PEKKA MALO¹,
EEVA VILKKUMAA¹, TARJA SAARESANTA^{2,3}, AND ARHO VELI VIRKKI^{2,3}

¹Department of Information and Service Management, Aalto University, 02150 Espoo, Finland

²Division of Medicine, Department of Pulmonary Diseases, Turku University Hospital and Sleep Research Centre, University of Turku, 20014 Turku, Finland

³Department of Pulmonary Diseases and Clinical Allergy, University of Turku, 20014 Turku, Finland

(Zhaoyang Chen and Lina Siltala-Li contributed equally to this work.) CORRESPONDING AUTHOR: L. SILTALA-LI (lina.siltala-li@aalto.fi)

This work was supported in part by the Foundation for Economic Education under Grant 16-9442, in part by the Paulo Foundation, and in part by the Helsinki School of Economics Foundation (HSE).

This article has supplementary downloadable material available at <https://doi.org/10.1109/JTEHM.2023.3276943>, provided by the authors.

ABSTRACT Background: Obstructive sleep apnea (OSA) is growing increasingly prevalent in many countries as obesity rises. Sufficient, effective treatment of OSA entails high social and financial costs for healthcare. Objective: For treatment purposes, predicting OSA patients' visit expenses for the coming year is crucial. Reliable estimates enable healthcare decision-makers to perform careful fiscal management and budget well for effective distribution of resources to hospitals. The challenges created by scarcity of high-quality patient data are exacerbated by the fact that just a third of those data from OSA patients can be used to train analytics models: only OSA patients with more than 365 days of follow-up are relevant for predicting a year's expenditures. Methods and procedures: The authors propose a translational engineering method applying two Transformer models, one for augmenting the input via data from shorter visit histories and the other predicting the costs by considering both the material thus enriched and cases with more than a year's follow-up. This method effectively adapts state-of-the-art Transformer models to create practical cost prediction solutions that can be implemented in OSA management, potentially enhancing patient care and resource allocation. Results: The two-model solution permits putting the limited body of OSA patient data to productive use. Relative to a single-Transformer solution using only a third of the high-quality patient data, the solution with two models improved the prediction performance's R^2 from 88.8% to 97.5%. Even using baseline models with the model-augmented data improved the R^2 considerably, from 61.6% to 81.9%. Conclusion: The proposed method makes prediction with the most of the available high-quality data by carefully exploiting details, which are not directly relevant for answering the question of the next year's likely expenditure. Clinical and Translational Impact Statement: Public Health—Lack of high-quality source data hinders data-driven analytics-based research in healthcare. The paper presents a method that couples data augmentation and prediction in cases of scant healthcare data.

INDEX TERMS Cost prediction, healthcare data augmentation, obstructive sleep apnea, transformer.

I. INTRODUCTION

Obstructive sleep apnea (OSA) is a chronic respiratory disease in which the upper airway repeatedly collapses during sleep. There is no question that this results in poor sleep quality, thereby leading to increased daytime drowsiness, deterioration of cognitive abilities, various comorbidities, and

even high rates of traffic and workplace accidents [1], [2]. In addition, many studies attest to high morbidity and mortality associated with the disease [3], [4], [5]. The prevalence of clinically diagnosed OSA was 3.7% in the Finnish adult population [6]. In Finland alone, 1.46 million people are estimated to suffer from moderate to severe sleep apnea,

according to data presented in the Finnish Medical Journal [7]. This represents an astonishingly high percentage of the country's population of 5.54 million. In the wake of growing public awareness of the serious health consequences possible if OSA is left untreated, it is reported a significant increase in referrals connected with sleep apnea [8].

It goes without saying that sufficient resources must be made available to match. All OSA patients should receive treatment. This requires physicians and healthcare decision-makers to plan budgets accurately and distribute supplies efficiently, for better resource allocation. Hence, they need information about the coming year's potential costs for OSA-related visits. While electronic healthcare records (EHRs) are ideal for training predictive models with data on visits to physicians, laboratory tests, and therapies, OSA complicates the use of this rich source of data because it is a chronic disease that involves irregular check-up intervals, extensive follow-up, and highly individualized treatments using evolving technologies. Therefore, a dragon of chaos exists in mining EHRs here, brought in by inconsistent coding over the years, large quantities of missing data, human input or measurement error, and loss of follow-ups. To at least some extent, these issues frequently arise in data analytics involving EHRs, which is unsurprising when one considers the messy landscape wrought by the complexity of pathology and epidemiology. Irrespective of these difficulties, the burgeoning quantities of data collected in EHRs renders them one of the best resources for healthcare research, and data-driven studies need to grapple with them [9]. Researchers take many approaches to the problem of EHRs' "data chaos" with one of the most popular applying state-of-the-art deep learning models since these do not presume any particular stochastic distributions to the data [10], [11], [12], [13], [14]. Still, few studies address predicting the cost of healthcare visits in a way that accounts for both total costs and the visit type at each point in time, let alone focus on making the most of the limited body of data available for particularly complex diseases.

Our study represents five key contributions to the state of the art:

- i) We develop a data-augmentation algorithm that preserves the semantic invariance of discrete healthcare data.
- ii) We propose a method to augment the input via a subset of the high-quality healthcare data, material that cannot otherwise directly serve addressing the research question.
- iii) A multi-task loss function is designed for cost prediction that considers both the sum-total costs and the cost specific to each type of visit.
- iv) We combine two Transformer models (one for data augmentation and the other for cost prediction) to achieve better predictive performance while tackling the problem of insufficient data.
- v) Our research experiments with and hones the cost-prediction model by working with EHRs from Finland's Turku University Hospital. The study appears to be among very few projects of this type for OSA. The code from

this study is available via https://gitlab.com/lina.siltala/two_model_transformer_predict_cost.

II. RELATED WORK

The worldwide volume of clinical data exceeded 2,300 exabytes in 2020 [15]. This vast body of data holds tremendous potential for data-driven analytics to support decision-making in healthcare, assessment of pathology trajectories, public-health surveillance, precision medicine, and preventive treatment, since EHRs encompass data covering consultations with experts, lab tests, clinical notes, and medication records [16], [17], [18]. Since the sensitivity of the data held in EHRs makes them an obvious target for cyber-attacks and attractive for deliberate data leaks, strict regulations are in place for their use, such as the EU's General Data Protection Regulation (GDPR), the United States Health Insurance Portability and Accountability Act (or HIPAA), and rules for the Australian government's My Health Record system [9]. The ironic twist is that, through these, EHRs' treasure trove of data is not readily amenable to research. Scholars gain access to healthcare data only after a lengthy process for specified research questions, and the research must comply with ethics codes and rules – e.g, the GDPR's terms for purpose limitation (Article 5(1)(b)), "data minimization" (Article 5(1)(c)), storage restrictions (Article 5(1)(e)), and integrity and confidentiality (Article 5(1)(f)) [9]. It is, without doubt, imperative to protect individuals' privacy by means of standardization and strict ethics, yet this does bring challenges. For example, it is not easy to obtain the quantities of data needed for solid studies, especially with regard to particular diseases. This marks a stark contrast against natural language processing (NLP), for which Wikipedia, libraries, and social media offer ample material.

Several further factors contribute to the difficulties of scholarly use of EHRs. Firstly, the records are created primarily with physicians and administration in mind, not for research purposes [17]. Also, discrepancies arise, brought on by changes in technology, adjustments to diagnostic codes, and variations in practices between or even within healthcare facilities. Heterogeneous and free-form data create further difficulties for EHR analysis, as do complex intra-patient variations. A fourth important factor is that these records do not cover patients comprehensively: most people visit physicians only when unwell [9]. Of the many challenges bundled with EHRs' use in healthcare research, the two issues that we most needed to address for our study are the limited body of high-quality data available and the complicated characteristics of the data.

Data augmentation is one technique for solving the first of these problems. Scholars of computer vision have frequently employed it for such purposes as cultivating more image data or applying clipping, rotation, color changes, or blurring. It is easy to understand how such techniques could serve such applications even without domain knowledge, since we know that the image is not converted to something completely different. That is, the post-augmentation body of data has

retained the original's semantics [19]. It is not so straightforward to apply these techniques to healthcare data. One of the reasons is that EHRs include many discrete variables. With these, keeping the semantic information intact is far more challenging than with the continuous variables that images involve. Reference [20] Work in the NLP domain, in contrast, points to possible ways forward, in that NLP data feature discrete variables and many scholarly efforts in that domain have tackled the problem of insufficient data, with a broad range of methods: random deletion, replacement, or injection [19]; dependency-tree morphing [21]; back-translation [22]; the manifold mixup regularization method [23]; unsupervised data augmentation [24]; kernel methods [25]; semantic augmentation [26]; and others. Work specifically with data augmentation for EHRs has applied contrastive learning to find similarity patterns [27], examined particular types of data (such as images [27] and textual clinical notes for patient-outcome prediction [20]), and explored subfields such as skin-lesion analysis [28]. Scholars have discussed the potential for addressing the data-quantity issue with deep learning via knowledge distilling [29], patient representation [30], vector learning with non-negative restricted Boltzmann machines (eNRBMs) [31], and transfer learning in the EHR context [32]. Workable data augmentation should be easy to implement while still improving the model's performance for the primary goal. All the aforementioned tactics turned out to be challenging to implement for our goal of predicting the next year's expenditures. While none of the techniques were directly applicable for our study, they offered inspiration for our augmentation method.

For each patient, the data records in our study are sequential and linked to specific visits, which vary in length. They are very similar to the many-to-many sequence-to-sequence (seq2seq) conditions in NLP [33]. Among the methods traditionally applied for seq2seq modeling are hidden Markov models (HMMs), latent semantic analysis (LSA), latent Dirichlet allocation (LDA), bag-of-words (BOW), skip-gram, words2vec, and global-vector representation [34], [35], [36], [37], [38], [39], [40], [41]. The renaissance of rapidly developing deep learning has channeled current approaches to sequential data mainly into the associated stream, though [42], with special attention surrounding the recurrent neural network and such variants as long short-term memory (LSTM) [43]. Among the state-of-the-art methods applied specifically in NLP-related work are variational autoencoders (VAEs) [44], [45], [46], [47], generative adversarial nets (GANs) [48], [49], adversarial learning for dialogue generation [50], text generation with reinforcement learning [51], [52], Transformer models [53], bidirectional encoder representation from Transformers (BERT) [54], [55], and individual solutions such as Generative Pre-trained Transformer 3 (GPT-2) [56] and ChatGPT [57].

Following in the wake of the deep learning revolution in the NLP field, much research with EHRs has considered scalable

deep learning in light of the two domains' parallels [12], [13], [58], [59], [60]. Many studies subject EHR data to deep learning for risk and disease prediction, data-privacy work, phenotyping, and disease classification, in a shift from labor-intensive feature engineering and other expert-driven methods. The target is data-driven approaches to representing complicated data in lower-dimensional space [61]. Several of these have demonstrated success in applying deep learning models in conjunction with EHRs. Among the tools produced are Deepr, using a conventional neural network (CNN) for deep extraction [62]; DoctorAI, which utilizes a recurrent neural network (RNN) for disease prediction; and the DeepCare system, incorporating LSTM for predicting medicine quantities [63]. Scholars have studied deep learning for numerous aspects of healthcare, such as predicting obesity [64], assessing the likelihood of readmission in cases of congestive heart failure [65], and providing diagnostic decision support via BERT [66]. Also, recent work has directed attention to the general issue of explainability, by means of RNNs and graphing of temporal data [10], [11], [67].

Encoder–decoder was the most suitable state-of-the-art deep learning technique for our research. Encoders provide an embedding that can successfully learn the latent patient representation while converting the multivariates to a lower-dimensional space. Decoder architecture uses the latent representation discovered in the encoder phase as the context information to learn autoregressively about the following visit. Thus, the model takes age, gender, and other temporal multivariable features of historical patient data as inputs, while its output is temporally univariate (visit costs only). A Transformer model accommodates several inductive biases for sequential forecasting [68]. This model type offers one of the most powerful encoder–decoder architectures because it has multi-head attention and self-attention [53].

Transformers are frequently compared with CNNs and RNNs. A CNN induces the inductive biases of invariance and locality with kernel functions, while an RNN handles temporal-invariance and locality-related inductive biases via its Markovian structure. In contrast against both of these, Transformers do not demand any strong assumption as to the data's structural nature [68]. We refer the readers to the original paper of Transformer for more details [53]. Although many studies apply Transformers accordingly [68], [69], [70], few of them have produced EHR-based cost prediction that consider not only the sum of all costs but also the type of cost associated with each visit instance. Filling this gap, we took inspiration from the context-learning functionality of encoders that can retain semantic and syntactic information. Our literature-informed approach makes sure that the context of the patients is learned with a similar structure so that the patients' semantics do not change significantly during data augmentation. Applying two Transformers with the same encoder structure made it easy to implement data augmentation from a subset of the data.

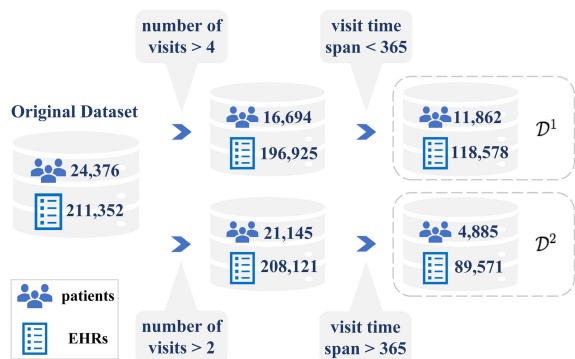


FIGURE 1. For analysis, the data were filtered by the number of visits and total follow-up duration. \mathcal{D}^1 is the set of patients with fewer than 365 days of follow-up, and \mathcal{D}^2 contains those patients with more than 365 days’ follow-up.

III. MATERIAL AND METHODOLOGY

Our data processing and methods are detailed below. The overall aim is to predict visit expenditures with a combination of original and augmented data. We address the data-augmentation and the cost-prediction element separately in relation to both the data and the model.

A. DATA

The filtering and preparation of data for this study are presented in Fig. 1, 2, and 3. The data included the years from 2002 to 2019. Its procedures were approved under research permit T164/2019 from Turku University Hospital. On account of patients’ irregular visit intervals and differences in follow-up duration, we would have had to contend with large quantities of “missing data” had we processed variables for all patients at the same time points, as studies often do [10], [69]. To circumvent the issue, we applied a data-processing “trick” from survival analysis [71]: we set the date of the patient’s first OSA diagnosis (identified as G47.3 in EHRs) as the start time of the study, then calculated the number of days between that and each visit for the variable `diff_dgn`, for “difference from diagnosis.” Other variables were recorded for each visit instance alongside `diff_dgn`. These captured both static and time-varying information, such as age, gender, the type of visit, and specialist type, as shown in Fig. 2 and 4. More detailed data descriptions and filtering of OSA cohort are presented in our supplementary material.

Formally, each patient i is covered via sequential multivariate data until time t_i , where the value of t at the maximum number of days from diagnosis may differ freely between patients on account of the differences in follow-up duration. All `diff_dgn` values for patient i form a set $T_i = \{0, \dots, t_i, t_i \in \mathbb{N}\}$. The elements of T_i correspond to the patient’s visits, which can be represented by set $S_i = \{1, 2, 3, \dots, s_i, s_i \in \mathbb{N}\}$, where s_i is the total number of visits for patient i (see Fig. 3). The patients’ data are represented as $\mathcal{D} = \{\mathbf{d}_{ik} : i = 1, \dots, N, k \in S_i\}$, where N is the number of patients and \mathbf{d}_{ik} is a data vector with M variables

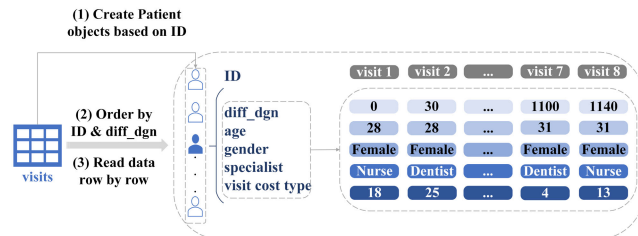


FIGURE 2. Variables for OSA patients. There are five patient-specific variables: days from diagnosis (`diff_dgn`), age, gender, specialist type, and visit cost type.

corresponding to patient i and visit k . As Fig. 1 and 3 indicate, we split the patients into two groups: those with under 365 days in all before the last visit (\mathcal{D}^1) and the other with more than 365 days’ follow-up duration (\mathcal{D}^2).

1) THE DATA FOR AUGMENTATION \mathcal{D}^1 .

The first set covers 11,862 patients, with 118,578 visit records. Because these patients’ consultations spanned less than 365 days, they could not be included in the cost prediction for the next year, which requires more than one year of visits. Hence, only set \mathcal{D}^2 was available for cost prediction. That set contains only 4,885 patients. The low number of patients having data for more than 365 days is explained by the two facts: 1) about 30 % of patients discontinue the treatment within the first year and 2) the number of patients has been much lower before the year 2019. Therefore, predicting a year’s costs from so few patient data is highly challenging. It is natural to turn to data augmentation for a possible solution. In NLP settings, models that have learned the similarity of words can substitute another word for one that carries similar meaning in the context (e.g., “cat” in place of “dog”), or some grammatical (semantic) information may be extracted such that deleting or inserting words yields a new sentence without distorting the meaning. Healthcare-specific data augmentation, in contrast, is problematic, because such patterns of word similarity or grammar have not been found yet, especially with regard to certain diseases; therefore, there is no clear standard of what one can delete, replace, or inject for data augmentation that preserves the patient records’ semantic information. Although we could not use \mathcal{D}^1 directly for the cost prediction, those sleep-apnea patients were treated at the same hospital as members of \mathcal{D}^2 . Thanks to the associated similarities in sequence patterns and other characteristics, extracting information from \mathcal{D}^1 represented a feasible route for data augmentation to ameliorate the issue of the restricted pool of data for our prediction.

Data augmentation is designed to expand the input dataset in a manner that fills the material out with noise alongside the semantic information preserved, thereby improving the performance of the model [22]. The augmentation in our case entailed changing only the details of one or two visits, so as to keep the characteristic sequence pattern of sleep-apnea patients intact. Hence, when we trained Transformer model

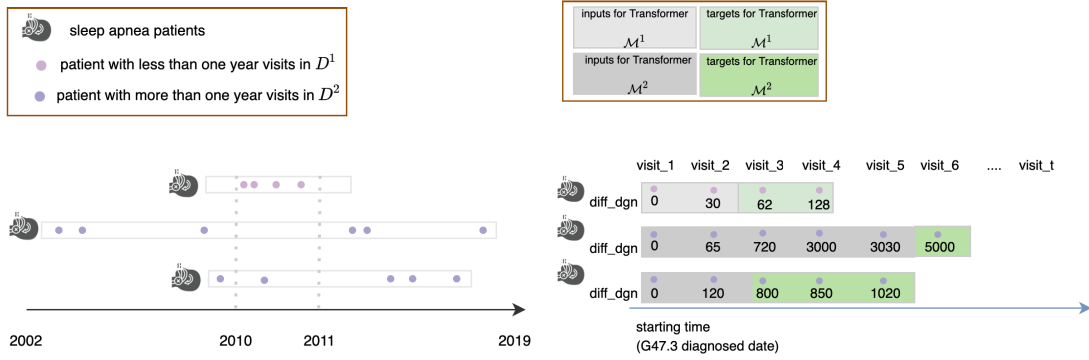


FIGURE 3. Processing of the study's data. The figure illustrates the data-processing via three hypothetical patients, with (as shown at left) unique start dates, follow-up durations, inter-visit intervals, and visit frequency. Our processing used the date of G47.3 diagnosis as the start time. Each visit was processed as a discrete record with corresponding `diff_dgn` values (shown in the right-hand pane). Patients with less than a year of follow-up were assigned to \mathcal{D}^1 while we placed the rest, with longer follow-up, in \mathcal{D}^2 . The study used \mathcal{D}^1 to train Transformer \mathcal{M}^1 for data augmentation, while \mathcal{D}^2 was employed for visit cost prediction with Transformer \mathcal{M}^2 .

\mathcal{M}^1 with \mathcal{D}^1 , we processed the data as inputs and target in the following way: the patient history data $\mathcal{D}_{\text{inp}}^1 = \{\mathbf{d}_{ik} : i = 1, \dots, N_1, 0 \leq k \leq s_i - 2, k \in S_i\}$, constitute the input, and the target (the patient's visit cost type) is $\mathcal{V}_{\text{out}}^1 = \{\mathbf{v}_{ik} : i = 1, \dots, N_1, s_i - 2 < k \leq s_i, k \in S_i\}$ as shown in Fig. 3, where N_1 is the number of patients in \mathcal{D}^1 , and $\mathbf{v}_{ik} = [v_{ikc}] \in \{0, 1\}^C$ is a binary vector, where C is the number of unique visit cost types, and $v_{ikc} = 1$ if and only if the cost type of visit k for patient i is c .

2) THE DATA FOR COST PREDICTION \mathcal{D}^2 .

The 4,885-patient dataset contains 89,571 visit records. These data were processed as inputs and targets (per Fig. 3 and 4) to predict, in our research setting, the visit costs over the last year (ignore leap year and only assume it has 365 days for simplicity). Since patients' visit history and intervals may be different in length, their time indices for the allocation of inputs and targets vary accordingly. We set $t'_i = \max\{x \in T_i : 0 \leq x < t_i - 364\}$ as the last time for the inputs and $t''_i = \min\{x \in T_i : t_i - 364 \leq x \leq t_i\}$ as the first time for the targets. The visits corresponding to t'_i and t''_i are s'_i and s''_i . Then, the inputs get expressed as $\mathcal{D}_{\text{inp}}^2 = \{\mathbf{d}_{ik} : i = 1, \dots, N_2, 0 \leq k \leq s'_i, k \in S_i\}$, and the targets (the patient's visit cost type) are $\mathcal{V}_{\text{out}}^2 = \{\mathbf{v}_{ik} : i = 1, \dots, N_2, s''_i \leq k \leq s_i, k \in S_i\}$. For simplicity, we denote the times corresponding to the inputs and targets as set $T_{i_{\text{inp}}} = \{0, \dots, t'_i\}$ and $T_{i_{\text{out}}} = \{t''_i, \dots, t_i\}$.

Each unique visit cost type has an associated cost value. The patients' visit costs can be represented as cost vectors, $\mathcal{C}_{\text{out}}^2 = \{c_{ik} : i = 1, \dots, N_2, s''_i \leq k \leq s_i, k \in S_i\}$, where the elements of $c_{ik} \in \mathbb{R}^C$ represent the costs attributed to different cost types. If the visit is not attributed to a cost type $c \in \{1, \dots, C\}$, then c_{ikc} is set equal to zero. When feeding the data to the neural networks, special tokens [CLS] and [SEP] are inserted to indicate the first and final visit by the patient, respectively.

It is worth noting that, because \mathcal{D}^1 contains many more patients than \mathcal{D}^2 , and, consequently, also many more visit

cost types (91 cost types in \mathcal{D}^1 versus 50 cost types in \mathcal{D}^2). To address the variation in the length of target sequences, all sequences are padded to have the same length.

B. THE MODEL ARCHITECTURE

Encoder–decoder models are very popular for seq2seq prediction problems, and Transformers are among the most powerful tools in this class [70]. Our choice to develop two Transformer models for enabling the efficient use of patient data in solid prediction of coming costs led to the architecture depicted in Fig. 5. This design addresses the two main differences between our work and application of seq2seq in NLP. 1) Rather than mostly univariate data (words in the NLP case), we had to factor in the multivariate nature of healthcare data (with variables such as age, gender, and specialists visited). For predicting next year's visit costs, historical cost information is not the only relevant variable. Others too are important, because they reflect between-patient differences and within-patient variance during follow-up – demographic information plays important roles in visit patterns and trajectories. For our design to consider all of the most influential variables, we applied multivariable embedding in the encoder for \mathcal{M}^1 and \mathcal{M}^2 , as shown in Fig. 4. 2) Because of the small quantity of data for training and the challenges created for transfer learning by the complexity of pathology and epidemiology, we needed a solution for efficiently putting data to use in healthcare studies. The trained model \mathcal{M}^1 with \mathcal{D}^1 for data augmentation held promise to solve this two-horned problem via more data and an alternative solution for transfer learning.

1) THE MODEL FOR DATA AUGMENTATION: \mathcal{M}^1 .

As the diagram in Fig. 5 indicates, \mathcal{M}^1 must be trained and evaluated before it gets employed for data augmentation. Although, as Fig. 6 shows, \mathcal{M}^1 and \mathcal{M}^2 have identical encoder portions, the decoder in \mathcal{M}^1 uses only a fixed visit horizon (for augmentation inferring the next one or two visits) while the decoder part of \mathcal{M}^2 accounts for differences in

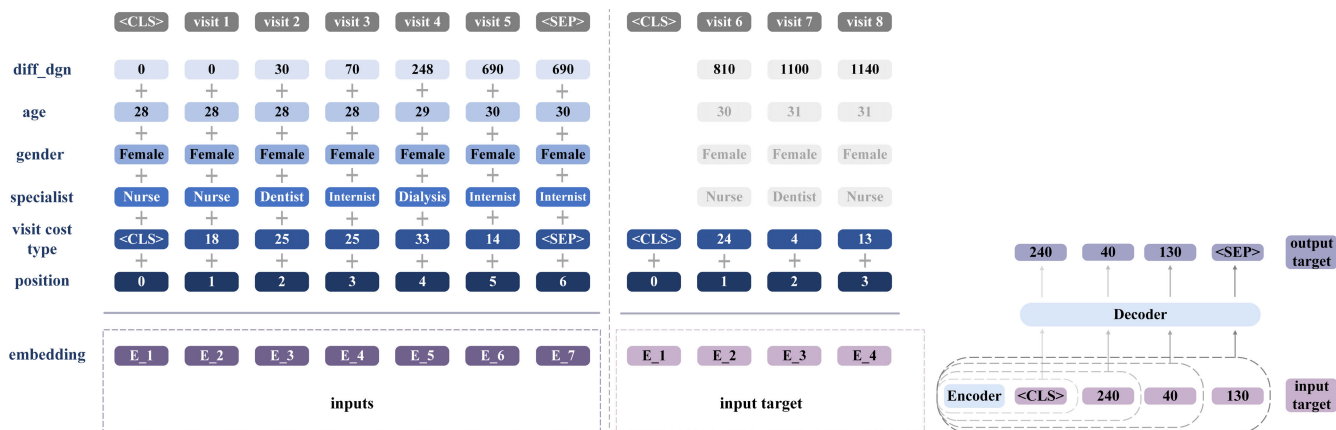


FIGURE 4. Split the data into inputs and target, and inputs' embedding. The patient represented made eight visits, and the $diff_dgn$ set is $\{0, 30, 70, 248, 690, 810, 1100, 1140\}$. We have records extending until 1,140 days after this patient's sleep-apnea diagnosis. The goal is to predict the visit cost over the last year ($1140 - 364 = 776$) on the basis of the previous visits' records (from the date of G47.3 diagnosis to day 776). Therefore, the time indices for the inputs and target are the maximum between 0 and 776 and the minimum between 776 and the last visit point, day 1,140. In line with the definition $t'_i = \max\{x \in T_i : 0 \leq x < t_i - 364\}$ and $t''_i = \min\{x \in T_i : t_i - 364 \leq x \leq t_i\}$, the inputs for this patient are the records from day 0 to day 690 on diagnosis, and the targets are the visit costs from day 810 to 1,140. The autoregressive prediction mechanism of Transformer is shown at right.

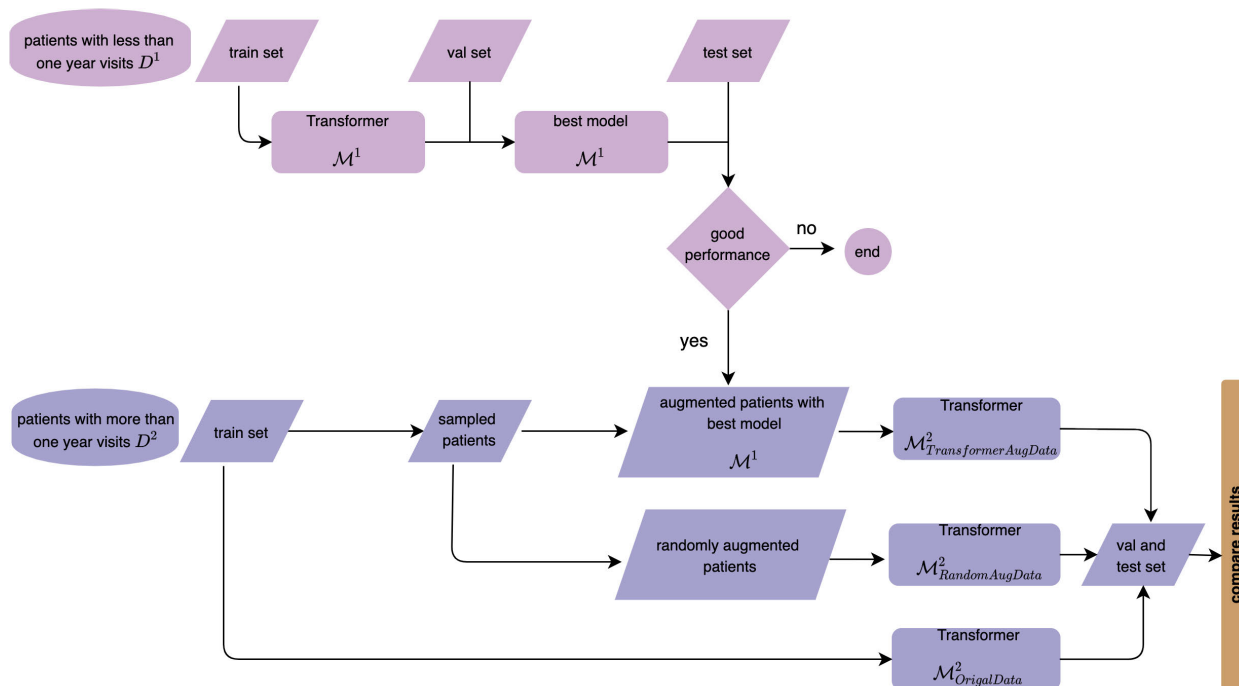


FIGURE 5. The model architecture. One of the two components is for data augmentation (applied with sub-sample D^1), and the other is for cost prediction (utilizing augmented data and D^2 in combination). The paper presents results from comparisons with the raw data, randomly augmented data, and Transformer \mathcal{M}^1 augmented data, to clarify the better prediction performance.

visit length, since patients' visit frequencies in the next year will differ. The loss function in \mathcal{M}^1 measures cross-entropy, which is often used for multi-label classification.

The algorithm developed for data augmentation includes deletion, replacement, and insertion, all of which are common tactics, especially for augmentation in situations with discrete variables [22]. Considering a function \mathcal{F} that embodies the data augmentation process, \mathcal{F} accepts a subset of the original data \mathcal{D}^2 as input and produces the augmented data $\widetilde{\mathcal{D}}^2$ as output. The concept of semantic invariance, as it pertains

to our research, signifies that the probability distribution of subsequent visits \mathcal{D}^2 closely approximates the distribution found in \mathcal{D}^2 . Given that the patients in \mathcal{D}^1 and \mathcal{D}^2 predominantly share the same ethnicity, are treated at the same regional hospital, and all have OSA, the clinical relevance of the augmented data $\widetilde{\mathcal{D}}^2$ relies on the similarity between these two datasets with respect to domain-specific information. Additionally, the high predictive performance of the trained model \mathcal{M}^1 , featuring a constrained parameter range of visit types and limited to only two subsequent visit predictions,

further reduces the noise of data alleviation and supports the clinical relevance of the augmented data.

After training and evaluating \mathcal{M}^1 with \mathcal{D}^1 , our method augments the data with randomly sampled patients \mathcal{D}_s^2 from \mathcal{D}^2 (line 1 in algorithm 1). For our study, \mathcal{M}^1 was trained with \mathcal{D}^1 , after which \mathcal{M}^1 is better suited to predicting the next visits for those patient having fewer than 365 days' follow-up as input data. For \mathcal{D}^2 , with more than 365 days of follow-up, the algorithm specifies time indices T_s^2 that aid in extracting visits that are less than 365 days from patients in \mathcal{D}_s^2 (line 2), to guarantee input data suitable for inferences with \mathcal{M}^1 . Visits up to T_s^2 are taken as inputs, and \mathcal{M}^1 infers the next one or two visits, predicting them from these inputs (lines 7-8). The next one or two visits predicted for these sample patients in light of their time indices are stored for later data augmentation. The data is stored for augmentation only if the prediction by \mathcal{M}^1 for the first visit is such that it is found in \mathcal{V}^2 , which contains the true visit cost type vectors associated with patients in \mathcal{D}^2 . If the predicted cost type is found only in \mathcal{D}^1 but not in \mathcal{D}^2 , then the observation is omitted from the sample (lines 13-14).

The augmentation of predicted cost types depends on the accuracy of the prediction produced by \mathcal{M}^1 . If the model is able to predict the first visit correctly, then the prediction for the second visit is augmented to the sample as a new observation (lines 16-17). However, if the prediction is incorrect for the first visit, but the visit cost type is still something that is found in \mathcal{V}^2 , then the original visit cost type in the sample is still replaced with the predicted value even though the prediction is known to differ from the true value (lines 17-18). This is done to ensure that the entire augmented data sample is consistent with the prediction produced by \mathcal{M}^1 . The predicted cost type for the second visit is omitted.

In our experiment, the training set had 3,910 patients from \mathcal{D}^2 and the sampling ratio was set to 55%, so the resulting set \mathcal{D}_s^2 included 2,150 patients. We iterated algorithm 1 three times and aggregated all results with the original training set, which gave us a total of 10,360 patients as $\mathcal{M}_{TransformerAugData}^2$ for \mathcal{M}^2 (as shown in Fig. 5) for cost prediction.

2) THE MODEL FOR COST PREDICTION: \mathcal{M}^2 .

Taking an approach similar to that in NLP, we compute the conditional probability of a patient's visit cost data $\mathcal{C}_i \in \mathcal{C}_{out}^2$ for each patient i given the corresponding input data $\mathcal{D}_i \in \mathcal{D}_{inp}^2$ and previous visit costs

$$P(\mathcal{C}_i | \mathcal{D}_i) = P(\mathbf{c}_{i,s_i'} | \mathcal{D}_i) \prod_{k=s_i'+1}^{s_i} P(\mathbf{c}_{i,k} | \mathbf{c}_{i,s_i'}, \dots, \mathbf{c}_{i,k-1}, \mathcal{D}_i). \quad (1)$$

3) THE LOSS FUNCTION FOR \mathcal{M}^2 .

Although our primary objective is to predict the sum-total visit costs for the next year on the basis of demographic

Algorithm 1 – Using Transformer \mathcal{M}^1 for Data Augmentation

Require: $\mathcal{M}_1, \mathcal{D}^2$, sample ratio

Ensure: augmented data

Step 1: Preparing the sample

- 1: Sample patients \mathcal{D}_s^2 from \mathcal{D}^2 with sample ratio
- 2: Define T_s^2 as a set of visit indices in \mathcal{D}_s^2 that have taken place in less than 365 days
- 3: Define \mathcal{V}^2 as the set of unique visit cost type vectors in \mathcal{D}^2
- 4: Define \mathcal{V}_s^2 as the visit cost type vectors for the patients included in sample \mathcal{D}_s^2

Step 2: Predicting cost types of next visits in \mathcal{D}_s^2

- 5: **for** $\mathbf{d}_{ik} \in \mathcal{D}_s^2$ **do**
- 6: # Check that visit index can be used as input
- 7: **if** $k \in T_s^2$ **then**
- 8: # Predict the cost types of next visits
- 9: $\hat{\mathbf{v}}_{ik}, \hat{\mathbf{v}}_{i(k+1)} \leftarrow \mathcal{M}^1.\text{predict}(\mathbf{d}_{ik})$
- 10: **else** go to line 4
- 11: **end if**
- 12: # Check that predicted visit cost is supported in \mathcal{D}^2
- 13: # Note: model \mathcal{M}^1 supports visit cost types that are not found in \mathcal{D}^2
- 14: **if** $\hat{\mathbf{v}}_{ik} \notin \mathcal{V}^2$ **then**
- 15: delete \mathbf{d}_{ik} from \mathcal{D}_s^2
- 16: go to line 4
- 17: **end if**
- 18: # Augment predicted cost types to \mathcal{D}_s^2
- 19: **if** $(\hat{\mathbf{v}}_{ik} == \mathbf{v}_{ik})$ and $(\hat{\mathbf{v}}_{i(k+1)} \neq \text{null})$ **then**
- 20: # Case 1: next visit is correctly predicted by \mathcal{M}^1
- 21: inject the 2nd predicted visit $\hat{\mathbf{v}}_{i(k+1)}$ in \mathcal{D}_s^2
- 22: go to line 4
- 23: **else**
- 24: # Case 2: next visit is predicted incorrectly but found
- 25: # as a visit cost type supported by \mathcal{D}^2
- 26: $\hat{\mathbf{p}}_{ik} \leftarrow \mathcal{M}^1.\text{prob}(\mathbf{d}_{ik})$
- 27: $c \leftarrow \text{argmax}_{c \in \{1, \dots, C\}} \{\hat{p}_{ikc} \exists \mathbf{v}_{ik} \in \mathcal{V}^2, v_{ikc} = 1\}$
- 28: define \mathbf{v} such that $v_c = 1$ and $v_j = 0$ if $j \neq c$.
- 29: replace \mathbf{v}_{ik} with \mathbf{v} in \mathcal{D}_s^2
- 30: **end if**
- 31: **end for**

details and information on visits (which occur sequentially during the last year of follow-up), it is important also to

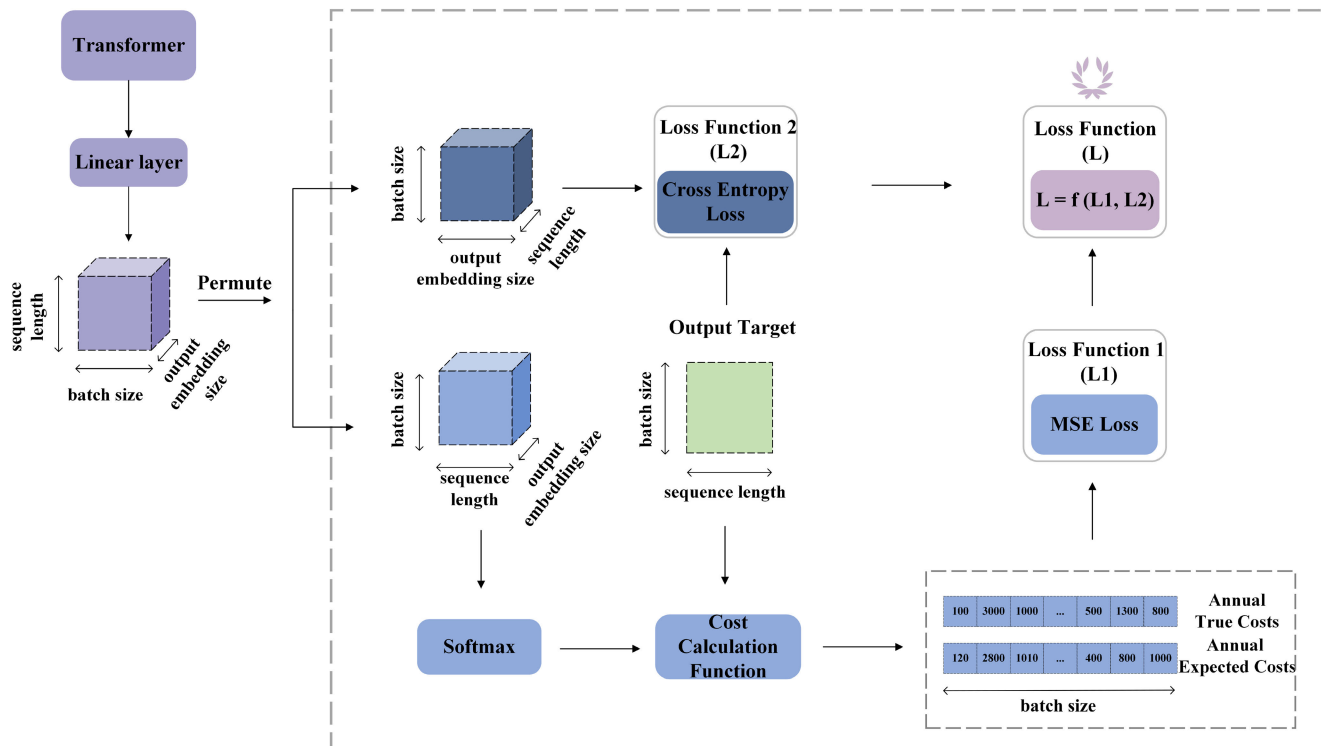


FIGURE 6. Loss function. To compute \mathcal{L}_1 , we permute the dimension of the output of the linear layer, i.e., [sequence length, batch size, output embedding size], into [batch size, sequence length, output embedding size]. Then the output of the new dimension is passed to the Softmax layer to calculate the possibilities of the dimension of the output embedding size. Next, the cost calculation function helps calculate the labeled and predicted annual costs based on the output target and the Softmax output. The process is much simpler for calculating \mathcal{L}_2 . The output of the linear layer is permuted into the dimension of [batch size, output embedding size, sequence length]. Then, the `cross_entropy()` function from PyTorch takes the permuted output and the output target to calculate the cross-entropy loss. Finally, since there is a large difference in magnitude between \mathcal{L}_1 and \mathcal{L}_2 , we choose the common logarithm (\log_{10}) to scale down \mathcal{L}_1 and then add it to \mathcal{L}_2 , i.e., $\mathcal{L} = \log_{10}(\mathcal{L}_1) + \mathcal{L}_2$.

predict the type of visit cost at each time point. This kind of prediction is necessary for two reasons: 1) From a practical perspective, it is highly informative for healthcare decision-makers. With this information, they not only can calculate annual costs accurately for budget purposes but also can efficiently allocate specific resources to individual departments in accordance with the predicted visit cost types. 2) From the computation standpoint, predicting individual visits serves to regularize the cost prediction such that its performance can be improved by means of a regularizer when it is forced to consider both the total cost and distinct cost types at each time point. Therefore, the total-loss function of \mathcal{M}^2 is defined as $f(\mathcal{L}_1, \mathcal{L}_2)$, which is a combination of regression cost \mathcal{L}_1 and multi-label cost \mathcal{L}_2 . In our study, we have experimented with three different functions $f(\mathcal{L}_1, \mathcal{L}_2)$. Their results are presented in the supplementary material. Based on the results, we have selected $f(\mathcal{L}_1, \mathcal{L}_2) = \log_{10}(\mathcal{L}_1) + \mathcal{L}_2$ as loss function.

Let $\mathcal{C}_b \subset \mathcal{C}_{\text{out}}^2$ be a minibatch of actual cost data for N_b patients'. Let $\mathcal{D}_b \subset \mathcal{D}_{\text{inp}}^2$ be the corresponding sample from the input data set. The regression cost is given by

$$\mathcal{L}_1(\mathcal{C}_b, \mathcal{D}_b) = \frac{1}{N_b} \sum_{i=1}^{N_b} \left(\sum_{k=1}^{K_i} \sum_{c=1}^C (\hat{p}_{ikc} e_c - c_{ikc}) \right)^2 \quad (2)$$

as the mean-squared error of the annual true cost and the annual predicted cost, where K_i is the number of visits for patient i , \hat{p}_{ikc} is the predicted probability that the cost type of the visit k is c , and e_c is the cost for a visit with type c . Since the model \mathcal{M}^2 does not directly predict the costs, we use the estimated probabilities $\hat{\mathbf{p}}_{ik} = \mathcal{M}^2.\text{prob}(\mathbf{d}_{ik})$, where \mathbf{d}_{ik} is the data of patient i associated with the visit k , together with expected visit type costs to approximate the total costs for a visit. Here, the actual cost of visit k by patient i is c_{ikc} , when the visit is known to be of cost type c . The predicted probabilities are calculated using soft-max function.

The multi-label cost \mathcal{L}_2 is defined as the cross entropy loss

$$\begin{aligned} \mathcal{L}_2(\mathcal{C}_b, \mathcal{D}_b) &= -\frac{1}{N_b} \sum_{i=1}^{N_b} \left(\frac{1}{K_i} \sum_{k=1}^{K_i} \sum_{c=1}^C \log \frac{\exp(o_{ikc})}{\sum_{c=1}^C \exp(o_{ikc})} \right), \quad (3) \end{aligned}$$

where o_{ikc} are the unnormalized logits produced by the last linear layer of the model, $\mathbf{o}_{ik} = \mathcal{M}^2.\text{logits}(\mathbf{d}_{ik})$.

C. EVALUATION METRICS

Because our objectives encompassed predicting two distinct major elements – the sum total of costs in the next year and each visit's cost type – we used two sets of metrics for model

evaluation: 1) for the regression modeling, root mean-squared error (RMSE) and R^2 to measure cost-prediction performance and 2) top- k accuracy indicators ($k=3,5,10$) for evaluation of the classification performance.

1) METRICS FOR CLASSIFICATION (\mathcal{M}^1 and \mathcal{M}^2).

Recommendation systems' ability to find the best options is often judged in terms of top- k accuracy [72]. healthcare analytics work has often followed design philosophy with such a "best bet" concept because it reflects the mindset of physicians performing diagnosis as they assess which diseases could be considered and whether, upon examination, the culprit might indeed be one of the candidates found [10]. In calculation of top- k accuracy, the prediction is deemed correct if the true label is among the model's k prediction with the highest predicted likelihood [72].

The formula is presented in Eq. 4. To avoid confusion with our earlier notation, where k is used as an index for the patient's visits, we will use h instead of k to denote the rank of the prediction, and we will use H as the number of prediction allowed for every true label. Let \hat{v}_{ij}^h denote the predicted cost type vector for patient i at visit number j , where the cost type is selected based on the h -th highest predicted likelihood. Let v_{ij} be the corresponding true label of visit cost type for patient i at visit number j . Indicator function $\mathbb{1}(\hat{v}_{ij}^h = v_{ij})$ has a value of 1 if $\hat{v}_{ij}^h = v_{ij}$; otherwise, the value is 0. In our study, we took the values 3, 5, and 10 as h for prediction of any single cost.

$$\begin{aligned} & \text{Top-}k \text{ accuracy}(v, \hat{f}) \\ &= \frac{1}{\sum_{i=1}^N (s_i - s'_i)} \sum_{i=1}^N \sum_{j=s'_i}^{s_i} \sum_{h=1}^H \mathbb{1}(\hat{v}_{ij}^h = v_{ij}) \end{aligned} \quad (4)$$

2) METRICS FOR REGRESSION (\mathcal{M}^2).

For measuring how close the predicted cost $\hat{C}_{i,\text{out}}^2$ is to the actual cost $C_{i,\text{out}}^2$, we chose two commonly used metrics suited to evaluating regression models [73]. In this evaluation, N was the number of patients.

Firstly, RMSE (the square root of the mean-squared error) gave us the expected value of the squared error or loss and it was computed as the square root of \mathcal{L}_1 . It enjoys widespread use because it is expressed in the same units as the response variable [73].

R^2 , in turn, expresses the proportion of the variance explained by the independent variables in the model [74]. Via the proportion of the variance explained, it shows how well the model can predict the unseen data. The best possible R^2 value is 1.0, while a value of 0.0 indicates that the model does not aid in explanation (i.e., it predicts the average value \bar{c}_{ikc}). [73] Eq. 5 presents the calculation of R^2 , where $\hat{c}_{ikc} = \hat{p}_{ikc}e_c$:

$$R^2 = 1 - \frac{\sum_{i=1}^{N_b} \left(\sum_{k=1}^{K_i} \sum_{c=1}^C (\hat{c}_{ikc} - c_{ikc}) \right)^2}{\sum_{i=1}^{N_b} \left(\sum_{k=1}^{K_i} \sum_{c=1}^C (\bar{c}_{ikc} - c_{ikc}) \right)^2} \quad (5)$$

TABLE 1. Top- k performance of \mathcal{M}^1 .

Set	Top3	Top5	Top10
Train	95.92%	97.70%	99.08%
Val	87.15%	90.77%	94.68%
Test	91.26%	94.04%	96.64%

IV. RESULTS

Our design split the dataset into training, validation, and testing sets. The purpose of validation is to prevent overfitting during training and to guarantee that the model gets evaluated in relation to an entirely unseen set of data [75]. The batch size and the learning rate were set as 64 and 0.0001 respectively. More information on experiments of hyperparameters can be found in the supplementary material. The performance of model \mathcal{M}^1 is presented in Table 1. Since all the top- k values for the test set exceed 91%, we apply the train \mathcal{M}^1 model for data augmentation based on the model design architecture shown in Fig. 5.

For the baseline models required as a reference for judging our two-model design, we have chosen four seq2seq models based on LSTM because LSTM is commonly used for seq2seq as one of the most popular encoder-decoder models. These four baseline models are: 1) a general LSTM-based encoder-decoder model, 2) an attention-oriented one, 3) a bidirectional one, and 4) an attention-based bidirectional one. Taking a combination approach, the last of these encoder-decoder models employs a bidirectional LSTM encoder to encode the historical time series while another LSTM decoder produces the future time series, with an attention mechanism implemented for coordinating the input and output time series and dynamically selecting the most pertinent contextual data for prediction purposes [76]. The other three baseline models are variants of the attention-based bidirectional LSTM (BiLSTM) encoder-decoder model. All four baseline models can learn from historical records and, thereby, create representative material to inform prediction.

We compared these models with our Transformer-based prediction, assessing the performance of each with the original raw data, randomly augmented data, and data augmented via Transformer model \mathcal{M}^1 . This section details the results, presented concisely in Table 2. Since our goal entailed giving the highest priority to predicting regression cost (\mathcal{L}_1), special emphasis is placed on regression metrics (R^2).

As the table attests, models without an attention mechanism display excellent top- k accuracy but have negative R^2 values, suggesting that these models fail to predict the year's total visit costs. In contrast, LSTM with attention and our \mathcal{M}^2 achieves an R^2 of more than 0.88. Transformer \mathcal{M}^2 stands out from all the other models for every indicator. As for performance with randomly augmented data for training and evaluation, the LSTM models do worse, BiLSTM models function slightly better, and – surprisingly – \mathcal{M}^2 has a higher R^2 value: 0.944. Finally, when compared to the original data, data augmented via the same Transformer model (\mathcal{M}^2) afford

TABLE 2. Models performance with original data and augmented data with transformer.

Model	Set	Original Data					Randomly Augmented Data					Augmented Data with Transformer, \mathcal{M}^1				
		Top3	Top5	Top10	RMSE	R^2	Top3	Top5	Top10	RMSE	R^2	Top3	Top5	Top10	RMSE	R^2
LSTM without attention	Train	57.87%	73.32%	80.32%	403.63	-0.002	63.90%	73.75%	80.55%	183.18	0.590	61.62%	73.34%	80.13%	347.08	0.317
	Val	57.92%	72.82%	80.04%	419.22	-0.001	62.90%	73.00%	80.09%	354.92	0.124	62.35%	74.30%	80.74%	273.83	0.152
	Test	60.46%	75.48%	80.85%	297.77	-0.011	66.00%	76.09%	81.22%	312.43	-0.098	62.08%	73.65%	80.28%	353.81	0.036
LSTM with attention	Train	60.72%	74.02%	80.33%	121.19	0.910	64.19%	76.44%	80.80%	93.86	0.943	60.95%	76.32%	80.92%	103.04	0.941
	Val	60.96%	73.36%	80.10%	222.66	0.728	65.43%	75.88%	80.09%	175.20	0.833	64.50%	76.54%	80.70%	96.62	0.890
	Test	60.46%	75.48%	80.85%	195.85	0.884	64.70%	76.34%	80.55%	161.64	0.697	63.49%	75.20%	79.85%	133.59	0.858
BiLSTM without attention	Train	62.95%	73.40%	80.10%	183.18	0.794	66.98%	76.54%	81.08%	135.95	0.881	66.65%	75.92%	80.87%	94.51	0.951
	Val	59.07%	72.53%	79.57%	418.73	0.036	62.94%	75.38%	79.96%	340.36	0.094	62.93%	76.38%	80.76%	294.24	-0.011
	Test	60.54%	75.74%	80.89%	311.16	-0.074	64.73%	76.82%	81.39%	290.28	-0.020	62.77%	75.67%	79.78%	363.77	0.011
BiLSTM with attention	Train	61.26%	74.10%	80.40%	118.44	0.914	65.34%	77.24%	81.31%	88.53	0.950	65.73%	77.11%	81.26%	92.21	0.953
	Val	59.74%	72.15%	79.85%	242.76	0.678	64.16%	76.21%	80.79%	223.40	0.732	65.35%	77.84%	81.00%	100.02	0.878
	Test	63.08%	76.29%	80.94%	175.97	0.616	64.56%	77.56%	81.34%	183.75	0.624	64.33%	76.58%	80.84%	151.35	0.819
Transformer \mathcal{M}^2	Train	80.62%	86.30%	92.46%	29.63	0.994	78.72%	85.41%	93.60%	8.96	0.999	78.22%	85.06%	92.88%	46.15	0.988
	Val	81.36%	86.27%	93.31%	33.84	0.995	80.61%	87.62%	93.56%	2.29	1.000	79.85%	87.25%	93.22%	21.64	0.995
	Test	83.64%	87.76%	93.45%	116.89	0.888	78.00%	83.58%	92.98%	94.51	0.944	80.25%	86.12%	93.14%	57.89	0.975

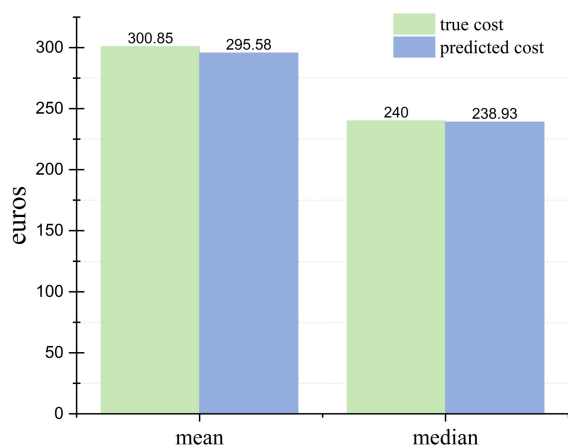


FIGURE 7. Comparison of mean and median values of true costs and predicted costs per patient per year.

better prediction of cost for all regression-model conditions except LSTM with attention (though the latter still is able to reach an R^2 of 0.858). The R^2 of Transformer \mathcal{M}^2 rises to 0.975 in this condition, and that of BiLSTM increases from 0.616 to 0.819.

Our designed model performs the best based on the results shown in Table 2. Therefore, we only apply this model to predict visit costs and compare them with true visit costs. The sum of the true total cost for all patients in the test set is €146,815. In comparison, the prediction by the trained model with Transformer augmented data is €144,244. Fig. 7 illustrates the comparison of mean and median values.

V. DISCUSSION AND CONCLUSION

In this paper, we have developed a method for predicting (1) the cost types for individual OSA patients’ future visits and (2) the total costs resulting from these visits using EHR data. This method utilizes a versatile Transformer-based architecture which helps make the most of the limited EHR data. While Transformers are applied in many fields – computer

vision, NLP, and also modern healthcare [70], our approach presents advances relative to existing work by using two Transformers (\mathcal{M}^1 , \mathcal{M}^2). The two-model system is easy to implement in this context since the encoder portion is identical and enables retaining the sequential patterns of the OSA patients who live in the same region and receive treatment at the same hospital. Model \mathcal{M}^1 is designed to augment the input by using material that is not suitable as-is for our prediction task. The second component, Transformer \mathcal{M}^2 , outputs not only the prediction of the next year’s total costs but also itemization by visit cost type for each visit instance, with the aid of two subsidiary loss functions. Our system outperforms all the baseline models covered in Table 2 both in top- k accuracy and by regression metrics. We empirically demonstrated prediction improvements arising from our model-informed data augmentation, which enriched the input relative to the original longer-follow-up a third of the high-quality data. We found also that summation of embeddings and utilization of special tokens can serve as an effective way to deal with multivariate sequences of healthcare data. The design sheds new light on approaches to tackle the problem of small bodies of relevant healthcare data and offers a different perspective on cost prediction for better decision-making.

Data augmentation is one of the best ways to make additional data available for research, and many studies already attest to the effectiveness of a corresponding strategy, via empirical evidence and design insight. However, the difficulty of assessing any given technique in the absence of a large-scale study creates obstacles: quantitative measurement of how well it fits the data is rare, and there is little research into why it works. Further issues often arise from a lack of variety in the augmented body of data. Frequently, the augmentation is accomplished under supervision, which may result in overfitting or bias in the prediction task [22]. Although a seq2seq model’s ability to retain long-term relationships enables it to handle tasks involving lengthy sequences quite well, it is unable to store contextual data. In contrast, the Transformer-based approach preserves data related to context. While it performs all tasks better than

seq2seq does, one must bear in mind Transformers' proneness to overfitting with small bodies of data in general, not just in cases of data augmentation [77].

Our future work will consider such issues. We recommend exploring other data-augmentation techniques and see how they affect prediction accuracy with datasets of various sizes. Further studies could also extend beyond the healthcare decision-making framework, delving into prediction outcomes at other levels and probing/cultivating direct linkages throughout the systems involved. For example, we should study integration of the cost-prediction model into healthcare's treatment-optimization process. For data specific to OSA, multitasking with EHRs could be handled more effectively. In one example, one might implement mortality/survival evaluation and cost prediction jointly in such a way that not only the financial element but also OSA patients' quality of life would be considered. Since abundant research has shown the power of Bayesian optimization for improved modelling, its use should be considered for informing prediction projects, although the method may prove time-consuming as the models grow larger and more complicated [78]. Our study exclusively emphasizes complete data, while ignoring variables containing missing data, such as medical markers including BMI and blood pressure. However, in our future research, we propose to address the issue of missing data and utilize the data more effectively.

ACKNOWLEDGMENT

The authors would like to thank Anna Shefl for proofreading the article.

REFERENCES

- [1] H. K. Walia, "Beyond heart health: Consequences of obstructive sleep apnea," *Cleveland Clinic J. Med.*, vol. 86, pp. 19–25, Sep. 2019.
- [2] S. Garbarino, O. Guglielmi, A. Sanna, G. L. Mancardi, and N. Magnavita, "Risk of occupational accidents in workers with obstructive sleep apnea: Systematic review and meta-analysis," *Sleep*, vol. 39, no. 6, pp. 1211–1218, Jun. 2016.
- [3] M. R. Bonsignore, P. Baiamonte, E. Mazzuca, A. Castrogiovanni, and O. Marrone, "Obstructive sleep apnea and comorbidities: A dangerous liaison," *Multidisciplinary Respiratory Med.*, vol. 14, no. 1, pp. 1–12, Dec. 2019.
- [4] H.-P. Chang, Y.-F. Chen, and J.-K. Du, "Obstructive sleep apnea treatment in adults," *Kaohsiung J. Med. Sci.*, vol. 36, no. 1, pp. 7–12, 2020.
- [5] L. A. Salman, R. Shulman, and J. B. Cohen, "Obstructive sleep apnea, hypertension, and cardiovascular risk: Epidemiology, pathophysiology, and management," *Current Cardiol. Rep.*, vol. 22, no. 2, pp. 1–9, Feb. 2020.
- [6] M. Palomäki, T. Saaresranta, U. Anttalainen, M. Partinen, J. Keto, and M. Linna, "Multimorbidity and overall comorbidity of sleep apnoea: A Finnish nationwide study," *ERJ Open Res.*, vol. 8, no. 2, pp. 3–5, Apr. 2022.
- [7] A. Bachour and H. Avellan-Hietanen, "Obstruktiivinen uniapnea aikuisilla [obstructive sleep apnea in adults]," *Suomen Lääkärilehti*, vol. 76, pp. 865–870, Apr. 2021.
- [8] T. Saaresranta and U. Anttalainen, "Uniapneaepidemia—Mitä hoidolla saavutetaan? [sleep apnea epidemics—What do we achieve with treatment?]," *Duodecim*, vol. 138, pp. 377–379, May 2022.
- [9] S. M. Shah and R. A. Khan, "Secondary use of electronic health record: Opportunities and challenges," *IEEE Access*, vol. 8, pp. 136947–136965, 2020.
- [10] E. Choi, M. T. Bahadori, A. Schuetz, W. F. Stewart, and J. Sun, "Doctor AI: Predicting clinical events via recurrent neural networks," in *Proc. Mach. Learn. Healthcare Conf.*, 2016, pp. 301–318.
- [11] E. Choi, M. T. Bahadori, J. Sun, J. Kulas, A. Schuetz, and W. Stewart, "RETAIN: An interpretable predictive model for healthcare using reverse time attention mechanism," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 29, 2016, pp. 1–9.
- [12] B. Shickel, P. J. Tighe, A. Bihorac, and P. Rashidi, "Deep EHR: A survey of recent advances in deep learning techniques for electronic health record (EHR) analysis," *IEEE J. Biomed. Health Informat.*, vol. 22, no. 5, pp. 1589–1604, Sep. 2018.
- [13] J. R. A. Solares et al., "Deep learning for electronic health records: A comparative review of multiple deep neural architectures," *J. Biomed. Informat.*, vol. 101, Jan. 2020, Art. no. 103337.
- [14] G. Kennedy, M. Dras, and B. Gallego, "Augmentation of electronic medical record data for deep learning," *Stud. Health Technol. Informat.*, vol. 290, pp. 582–586, Jan. 2022.
- [15] P. K. D. Pramanik, S. Pal, and M. Mukhopadhyay, "Healthcare big data: A comprehensive overview," in *Research Anthology on Big Data Analytics, Architectures, and Applications*. Hershey, PA, USA: IGI Global, 2022, pp. 119–147.
- [16] P. J. O'Connor et al., "Impact of electronic health record clinical decision support on diabetes care: A randomized trial," *Ann. Family Med.*, vol. 9, no. 1, pp. 12–21, Jan. 2011.
- [17] N. Menachemi and T. H. Collum, "Benefits and drawbacks of electronic health record systems," *Risk Manag. Healthcare Policy*, vol. 4, p. 47, May 2011.
- [18] G. S. Birkhead, M. Klompas, and N. R. Shah, "Uses of electronic health records for public health surveillance to advance public health," *Annu. Rev. Public Health*, vol. 36, no. 1, pp. 345–359, Mar. 2015.
- [19] J. Wei and K. Zou, "EDA: Easy data augmentation techniques for boosting performance on text classification tasks," 2019, *arXiv:1901.11196*.
- [20] Q. Lu, D. Dou, and T. H. Nguyen, "Textual data augmentation for patient outcomes prediction," in *Proc. IEEE Int. Conf. Bioinf. Biomed. (BIBM)*, Dec. 2021, pp. 2817–2821.
- [21] G. G. Sahin and M. Steedman, "Data augmentation via dependency tree morphing for low-resource languages," 2019, *arXiv:1903.09460*.
- [22] S. Y. Feng et al., "A survey of data augmentation approaches for NLP," 2021, *arXiv:2105.03075*.
- [23] V. Verma et al., "Manifold mixup: Better representations by interpolating hidden states," in *Proc. Int. Conf. Mach. Learn.*, 2019, pp. 6438–6447.
- [24] Q. Xie, Z. Dai, E. Hovy, M.-T. Luong, and Q. V. Le, "Unsupervised data augmentation for consistency training," in *Proc. NIPS*, vol. 33, 2020, pp. 6256–6268.
- [25] T. Dao, A. Gu, A. Ratner, V. Smith, C. De Sa, and C. Re, "A kernel theory of modern data augmentation," in *Proc. Int. Conf. Mach. Learn.*, 2019, pp. 1528–1537.
- [26] Y. Nie, Y. Tian, X. Wan, Y. Song, and B. Dai, "Named entity recognition for social media texts with semantic augmentation," 2020, *arXiv:2010.15458*.
- [27] T. Wanyan, J. Zhang, Y. Ding, A. Azad, Z. Wang, and B. S. Glicksberg, "Bootstrapping your own positive sample: Contrastive learning with electronic health record data," 2021, *arXiv:2104.02932*.
- [28] F. Perez, C. Vasconcelos, S. Avila, and E. Valle, "Data augmentation for skin lesion analysis," in *OR 2.0 Context-Aware Operating Theaters, Computer Assisted Robotic Endoscopy, Clinical Image-Based Procedures, and Skin Image Analysis*. Berlin, Germany: Springer, 2018, pp. 303–311.
- [29] Z. Che, S. Purushotham, R. Khemani, and Y. Liu, "Distilling knowledge from deep networks with applications to healthcare domain," 2015, *arXiv:1512.03542*.
- [30] E. Choi et al., "Multi-layer representation learning for medical concepts," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, Aug. 2016, pp. 1495–1504.
- [31] T. Tran, T. D. Nguyen, D. Phung, and S. Venkatesh, "Learning vector representation of medical objects via EMR-driven nonnegative restricted Boltzmann machines (eNRBM)," *J. Biomed. Informat.*, vol. 54, pp. 96–105, Apr. 2015.
- [32] S. Dubois, N. Romano, K. Jung, N. Shah, and D. C. Kale, "The effectiveness of transfer learning in electronic health records data," in *Proc. ICLR Workshop*, 2017, pp. 1–4.
- [33] I. Sutskever, O. Vinyals, and Q. V. Le, "Sequence to sequence learning with neural networks," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 27, 2014, pp. 1–9.
- [34] S. R. Eddy, "Hidden Markov models," *Current Opinion Struct. Biol.*, vol. 6, no. 6, pp. 361–365, 1996.
- [35] J. Lafferty, A. McCallum, and F. C. Pereira, "Conditional random fields: Probabilistic models for segmenting and labeling sequence data," in *Proc. 18th Int. Conf. Mach. Learn. (ICML)*. San Francisco, CA, USA: Morgan Kaufmann, 2001, pp. 282–289.

- [36] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent Dirichlet allocation," *J. Mach. Learn. Res.*, vol. 3, pp. 993–1022, Mar. 2003.
- [37] S. T. Dumais, "Latent semantic analysis," *Annu. Rev. Inf. Sci. Technol.*, vol. 38, no. 1, pp. 188–230, 2004.
- [38] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean, "Distributed representations of words and phrases and their compositionality," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 26, 2013, pp. 1–9.
- [39] T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient estimation of word representations in vector space," 2013, *arXiv:1301.3781*.
- [40] P. Bojanowski, E. Grave, A. Joulin, and T. Mikolov, "Enriching word vectors with subword information," *Trans. Assoc. Comput. Linguistics*, vol. 5, pp. 135–146, Dec. 2017.
- [41] J. Pennington, R. Socher, and C. Manning, "GloVe: Global vectors for word representation," in *Proc. Conf. Empirical Methods Natural Lang. Process. (EMNLP)*, 2014, pp. 1532–1543.
- [42] T. Iqbal and S. Qureshi, "The survey: Text generation models in deep learning," *J. King Saud Univ.-Comput. Inf. Sci.*, vol. 34, no. 6, pp. 2515–2528, Jun. 2022.
- [43] H. Salehinejad, S. Sankar, J. Barfett, E. Colak, and S. Valaee, "Recent advances in recurrent neural networks," 2017, *arXiv:1801.01078*.
- [44] D. P. Kingma and M. Welling, "Auto-encoding variational Bayes," 2013, *arXiv:1312.6114*.
- [45] S. Semeniuta, A. Severyn, and E. Barth, "A hybrid convolutional variational autoencoder for text generation," 2017, *arXiv:1702.02390*.
- [46] Z. Yang, Z. Hu, R. Salakhutdinov, and T. Berg-Kirkpatrick, "Improved variational autoencoders for text modeling using dilated convolutions," in *Proc. Int. Conf. Mach. Learn.*, 2017, pp. 3881–3890.
- [47] D. P. Kingma and M. Welling, "An introduction to variational autoencoders," *Found. Trends Mach. Learn.*, vol. 12, no. 4, pp. 307–392, 2019.
- [48] I. Goodfellow et al., "Generative adversarial networks," *Commun. ACM*, vol. 63, no. 11, pp. 139–144, 2020.
- [49] S. R. Bowman, L. Vilnis, O. Vinyals, A. M. Dai, R. Jozefowicz, and S. Bengio, "Generating sentences from a continuous space," 2015, *arXiv:1511.06349*.
- [50] J. Li, W. Monroe, T. Shi, S. Jean, A. Ritter, and D. Jurafsky, "Adversarial learning for neural dialogue generation," 2017, *arXiv:1701.06547*.
- [51] J. Li, W. Monroe, A. Ritter, M. Galley, J. Gao, and D. Jurafsky, "Deep reinforcement learning for dialogue generation," 2016, *arXiv:1606.01541*.
- [52] Z. Shi, X. Chen, X. Qiu, and X. Huang, "Toward diverse text generation with inverse reinforcement learning," 2018, *arXiv:1804.11258*.
- [53] A. Vaswani et al., "Attention is all you need," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 30, 2017, pp. 1–11.
- [54] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," 2018, *arXiv:1810.04805*.
- [55] S. Rothe, S. Narayan, and A. Severyn, "Leveraging pre-trained checkpoints for sequence generation tasks," *Trans. Assoc. Comput. Linguistics*, vol. 8, pp. 264–280, Dec. 2020.
- [56] L. Floridi and M. Chiriatti, "GPT-3: Its nature, scope, limits, and consequences," *Minds Mach.*, vol. 30, no. 4, pp. 681–694, Dec. 2020.
- [57] *ChatGPT: Optimizing Language Models for Dialogue*. OpenAI, San Francisco, CA, USA, Nov. 2022.
- [58] A. Rajkomar et al., "Scalable and accurate deep learning with electronic health records," *NPJ Digit. Med.*, vol. 1, no. 1, pp. 1–10, 2018.
- [59] A. Esteva et al., "A guide to deep learning in healthcare," *Nature Med.*, vol. 25, no. 1, pp. 24–29, Jan. 2019.
- [60] R. Miotto, F. Wang, S. Wang, X. Jiang, and J. T. Dudley, "Deep learning for healthcare: Review, opportunities and challenges," *Briefings Bioinf.*, vol. 19, no. 6, pp. 1236–1246, Nov. 2018.
- [61] C. Xiao, E. Choi, and J. Sun, "Opportunities and challenges in developing deep learning models using electronic health records data: A systematic review," *J. Amer. Med. Inform. Assoc.*, vol. 25, no. 10, pp. 1419–1428, Oct. 2018.
- [62] P. Nguyen, T. Tran, N. Wickramasinghe, and S. Venkatesh, "DeepR: A convolutional net for medical records," *IEEE J. Biomed. Health Informat.*, vol. 21, no. 1, pp. 22–30, Jan. 2017.
- [63] T. Pham, T. Tran, D. Phung, and S. Venkatesh, "DeepCare: A deep dynamic memory model for predictive medicine," in *Proc. Pacific-Asia Conf. Knowl. Discovery Data Mining*. Cham, Switzerland: Springer, 2016, pp. 30–41.
- [64] M. Gupta, T.-L.-T. Phan, H. T. Bunnell, and R. Beheshti, "Obesity prediction with EHR data: A deep learning approach with interpretable elements," *ACM Trans. Comput. Healthcare*, vol. 3, no. 3, pp. 1–19, Jul. 2022.
- [65] A. Ashfaq, A. Sant'Anna, M. Lingman, and S. Nowaczyk, "Readmission prediction using deep learning on electronic health records," *J. Biomed. Informat.*, vol. 97, Sep. 2019, Art. no. 103256.
- [66] R. Tang et al., "Embedding electronic health records to learn BERT-based models for diagnostic decision support," in *Proc. IEEE 9th Int. Conf. Healthcare Informat. (ICHI)*, Aug. 2021, pp. 311–319.
- [67] E. Choi, M. T. Bahadori, L. Song, W. F. Stewart, and J. Sun, "GRAM: Graph-based attention model for healthcare representation learning," in *Proc. 23rd ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, Aug. 2017, pp. 787–795.
- [68] Y. Tay, M. Dehghani, D. Bahri, and D. Metzler, "Efficient transformers: A survey," *ACM Comput. Surv.*, vol. 55, no. 6, pp. 1–28, Jul. 2023.
- [69] Y. Li et al., "BEHRT: Transformer for electronic health records," *Sci. Rep.*, vol. 10, no. 1, pp. 1–12, Apr. 2020.
- [70] T. Lin, Y. Wang, X. Liu, and X. Qiu, "A survey of transformers," *AI Open*, vol. 3, pp. 111–132, Oct. 2022.
- [71] R. G. Miller, *Survival Analysis*. Hoboken, NJ, USA: Wiley, 2011.
- [72] P. Cremonesi, Y. Koren, and R. Turrin, "Performance of recommender algorithms on top-N recommendation tasks," in *Proc. 4th ACM Conf. Recommender Syst.*, Sep. 2010, pp. 39–46.
- [73] N. R. Draper and H. Smith, *Applied Regression Analysis*, vol. 326. Hoboken, NJ, USA: Wiley, 1998.
- [74] A. Botchkarev, "Performance metrics (error measures) in machine learning regression, forecasting and prognostics: Properties and typology," 2018, *arXiv:1809.03006*.
- [75] M. A. Ottom, H. A. Rahman, and I. D. Dinov, "ZNet: Deep learning approach for 2D MRI brain tumor segmentation," *IEEE J. Translational Eng. Health Med.*, vol. 10, pp. 1–8, 2022.
- [76] Y. Yuan et al., "Using an attention-based LSTM encoder–decoder network for near real-time disturbance detection," *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 13, pp. 1819–1832, 2020.
- [77] P. W. Battaglia et al., "Relational inductive biases, deep learning, and graph networks," 2018, *arXiv:1806.01261*.
- [78] A. I. Cowen-Rivers et al., "HEBO pushing the limits of sample-efficient hyperparameter optimisation," 2020, *arXiv:2012.03826*.

• • •