

Deep Learning for Medical Prediction in Electronic Health Records

Xinlu Zhang

Department of Computer Science
University of California, Santa Barbara
xinluzhang@ucsb.edu

Abstract

The widespread adoption of electronic health records (EHRs) has opened up new opportunities for using deep neural networks to enhance healthcare. However, modeling EHR data can be challenging due to its complex properties, such as missing values, data scarcity in multi-hospital systems, and multimodal irregularity. How to tackle various issues in EHRs for improving medical prediction is challenging and under exploration. I separately illustrate my works to address these issues in EHRs and discuss potential future directions.

Missing Value

Different methods are available to address missing value in EHRs, such as mean value imputation (Acuna and Rodriguez 2004), multivariate imputation by chained equations (MICE) (Buuren and Groothuis-Oudshoorn 2010) and generative adversarial imputation nets (GAIN) (Yoon, Jordon, and Van Der Schaar 2018). However, these methods are performed in the data preprocessing stage to obtain imputed data which is feed into a classifier for downstream tasks, ignoring the underlying connections between missing value and downstream tasks. To fill this gap, we propose a generative adversarial network (GAN) (Goodfellow et al. 2020) variant to impute missing value in EHRs with classifier-guided for improving downstream medical predictions. In training, the generator imputes missing values and obtains imputed data; the classifier takes the generator’s outcomes, models the relationship between imputed data and labels by joint training with the generator, and outputs estimated labels. The discriminator attempts to identify which value is observed by taking imputed data from the generator and predicted label from the classifier. We conducted experiments in two EHRs and show that our proposed method outperforms GAIN and MICE consistently for downstream predictions in different missing ratio scenarios and evaluation metrics.

Data Scarcity in Multi-hospital Systems

Although applying deep neural networks in EHRs has received growing attention, data scarcity is prevalent in different hospitals and sub-medical domains (Desautels et al.

2017). Domain adaptation (DA) is a subcategory of transfer learning, utilizing knowledge from a different but related large source domain to improve model performance for a small target domain (Ben-David et al. 2010). DA has made remarkable progress in computer vision (Csurka 2017) and natural language processing (Ziser and Reichart 2016). Prior works (Sun et al. 2019; Khoshnevisan and Chi 2020) have applied DA to clinical predictions across multiple hospital systems. However, these works only use overlapping features in both domains, and ignore information provided by distinct features in the target and source (*feature disparity*), which contain the properties of different EHRs and can be important for prediction performance. In this work, we aim to fully utilize DA to improve medical prediction on a small target dataset by leveraging domain-invariant knowledge from another large source dataset with feature disparity. Specifically, we introduce the private encoding technique to map domains from different feature spaces into the same hidden representation space, and utilize pairing sampling techniques (Motiian et al. 2017a,b) to pair each target data point to abundant source data points for applying different DA approaches. We conduct experiments of different DA techniques on early-stage mortality prediction for trauma patients in multi-hospital system EHRs with feature disparity. Our experimental results show that DA techniques with our proposed methodology improve mortality prediction performance consistently and significantly in various training scenarios, with F1-scores up to relative 100%.

Irregularity in Multimodalities

EHRs are multimodal, containing numerical time series, which are multivariate irregularly sampled (MISTS) (Shukla and Marlin 2020) with features collected at both regular and irregular time points, and clinical notes taken at irregular time intervals. The complex irregular temporal property and multimodal structure make modeling EHRs for medical predictions challenging. Prior works (Deznabi, Iyyer, and Fiterau 2021; Yang and Wu 2021) ignore the irregularity in each modality in multimodal fusion. To full integrate irregularity into multimodal representation learning, we propose to separately tackle irregularity in each single modality and fuse their representations across temporal steps to improve medical predictions. First of all, inspired by Mixture-of-Experts (MoE) (Shazeer et al. 2017; Jacobs

et al. 1991), which maintains a set of experts (neural networks) and learns a combination of the experts specific to each input via a gating mechanism, we leverage different time series embeddings as submodules and dynamically integrate and obtain gains from them. In addition, we cast clinical note representations with note-taking time as MISTS, and leverage a time attention mechanism (Shukla and Marlin 2021) to tackle irregularity in each dimension. Finally, we incorporate irregularity into multimodal representations across temporal steps by adopting a fusion method that interleaves self-attentions and cross-attentions (Vaswani et al. 2017). Extensive experiments and ablation studies show the effectiveness of our approaches in every single modality and multimodal fusion for medical predictions and the importance of modeling irregularity.

Future Direction

For my future direction, I am going to explore potential topics in bioNLP and multimodal learning. The large language models (LLM) have achieved impressive performance with in-context learning, which solves unseen tasks without any parameter update (Brown et al. 2020) by adding a few demonstrations as prompts. Recently, a few works have utilized LLM on several bioNLP tasks (Moradi et al. 2021; Agrawal et al. 2022), but how to leverage LLM in the biomedical domain is still under exploration. In the next half-year, I will focus on utilizing LLM to improve biomedical tasks including medical prediction in EHRs.

References

- Acuna, E.; and Rodriguez, C. 2004. The treatment of missing values and its effect on classifier accuracy. In *Classification, clustering, and data mining applications*, 639–647. Springer.
- Agrawal, M.; Hegselmann, S.; Lang, H.; Kim, Y.; and Sontag, D. 2022. Large Language Models are Zero-Shot Clinical Information Extractors. *arXiv preprint arXiv:2205.12689*.
- Ben-David, S.; Blitzer, J.; Crammer, K.; Kulesza, A.; Pereira, F.; and Vaughan, J. W. 2010. A theory of learning from different domains. *Machine learning*, 79(1): 151–175.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33: 1877–1901.
- Buuren, S. v.; and Groothuis-Oudshoorn, K. 2010. mice: Multivariate imputation by chained equations in R. *Journal of statistical software*, 1–68.
- Csurka, G. 2017. Domain adaptation for visual applications: A comprehensive survey. *arXiv preprint arXiv:1702.05374*.
- Desautels, T.; Calvert, J.; Hoffman, J.; Mao, Q.; Jay, M.; Fletcher, G.; Barton, C.; Chettipally, U.; Kerem, Y.; and Das, R. 2017. Using transfer learning for improved mortality prediction in a data-scarce hospital setting. *Biomedical informatics insights*, 9: 1178222617712994.
- Deznabi, I.; Iyyer, M.; and Fiterau, M. 2021. Predicting in-hospital mortality by combining clinical notes with time-series data. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, 4026–4031.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2020. Generative adversarial networks. *Communications of the ACM*, 63(11): 139–144.
- Jacobs, R. A.; Jordan, M. I.; Nowlan, S. J.; and Hinton, G. E. 1991. Adaptive mixtures of local experts. *Neural computation*, 3(1): 79–87.
- Khoshnevisan, F.; and Chi, M. 2020. An adversarial domain separation framework for septic shock early prediction across ehr systems. In *2020 IEEE International Conference on Big Data (Big Data)*, 64–73. IEEE.
- Moradi, M.; Blagec, K.; Haberl, F.; and Samwald, M. 2021. GPT-3 models are poor few-shot learners in the biomedical domain. *arXiv preprint arXiv:2109.02555*.
- Motiian, S.; Jones, Q.; Iranmanesh, S. M.; and Doretto, G. 2017a. Few-shot adversarial domain adaptation. *arXiv preprint arXiv:1711.02536*.
- Motiian, S.; Piccirilli, M.; Adjero, D. A.; and Doretto, G. 2017b. Unified deep supervised domain adaptation and generalization. In *Proceedings of the IEEE international conference on computer vision*, 5715–5725.
- Shazeer, N.; Mirhoseini, A.; Maziarz, K.; Davis, A.; Le, Q.; Hinton, G.; and Dean, J. 2017. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *arXiv preprint arXiv:1701.06538*.
- Shukla, S. N.; and Marlin, B. M. 2020. A survey on principles, models and methods for learning from irregularly sampled time series. *arXiv preprint arXiv:2012.00168*.
- Shukla, S. N.; and Marlin, B. M. 2021. Multi-time attention networks for irregularly sampled time series. *arXiv preprint arXiv:2101.10318*.
- Sun, Z.; Peng, S.; Yang, Y.; Wang, X.; and Li, F. 2019. A General Fine-tuned Transfer Learning Model for Predicting Clinical Task Acrossing Diverse EHRs Datasets. In *2019 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, 490–495. IEEE.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Yang, B.; and Wu, L. 2021. How to Leverage Multimodal EHR Data for Better Medical Predictions? *arXiv preprint arXiv:2110.15763*.
- Yoon, J.; Jordon, J.; and Van Der Schaar, M. 2018. Gain: Missing data imputation using generative adversarial nets. *arXiv preprint arXiv:1806.02920*.
- Ziser, Y.; and Reichart, R. 2016. Neural structural correspondence learning for domain adaptation. *arXiv preprint arXiv:1610.01588*.