A Geometric Deep Learning approach to blood pressure regression.

Ferdinando Zanchetta[*], **Angelica Simonetti**[†], **Giovanni Faglioni**[‡], **Andrea Malagoli**[§] and **Rita Fioresi**[*]

Abstract

A correct measurement of the Blood Pressure (BP) is crucial in monitoring the health status of a patient and in diagnosis of cardiac vascular diseases. In this proof of concept, we develop and test a new Geometric Deep Learning (GDL) approach to infer BP data from other biological parameters: the photoplethysmogram and the electrocardiogram (ECG). Our findings suggest that such a GDL approach shows great promise compared with many state of the art models that are used in the field.

1. Introduction

Human cardiovascular system is a complex dynamical system, with multiple feedbacks, that we can monitor via key parameters like blood pressure (BP), heart rate (HR), cardiac cycle evolution etc. New challenges for AI systems to interpret and predict the time evolution of such parameters call for new approaches (Zhang et al. (2013) and refs. therein).

In particular, systolic BP is crucial in diagnosis of cardiac vascular diseases (CVD) Amini et al. (2021); Lim et al. (2013), however traditional cuff-based BP measurement devices cannot give a continuous BP signal, which is key for diagnosis and correct predictions, and are not comfortable to use for long periods of time. Recently, new light portable devices, like the photoplethysmogram (PPG), enable to obtain some of the above mentioned parameters, but not the BP, which, however can be inferred by them with standard techniques together with more machine learning oriented algorithms (P. Su et al. (2018), Chen et al. (2000), and refs. therein).

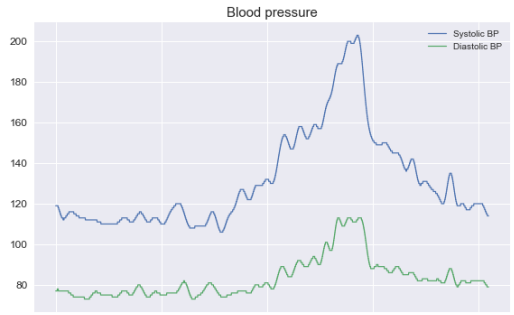

We plan to approach the question of inferring, using a graph neural network, the value of the BP, given the input signal of a PPG wearable device. Our paper represents a proof of concept contribution, we plan to furtherly improve our performances and results with a future study involving more data and variability in patient collection samples.

This is a summary of our contribution:

- This paper presents a novel Geometric Deep Learning approach to the question of determining the BP from PPG and ECG signals that is tested on real data;
- To the best of our knowledge, this is the first time that GDL methods, in particular Graph Neural Networks (Kipf and Welling (2017); Velickovic et al. (2018); Hamilton et al. (2017); Bronstein et al. (2021)), are employed for the solution of this question.

[*] University of Bologna [†] University of Lancaster [‡] Nabla2 [§] VST

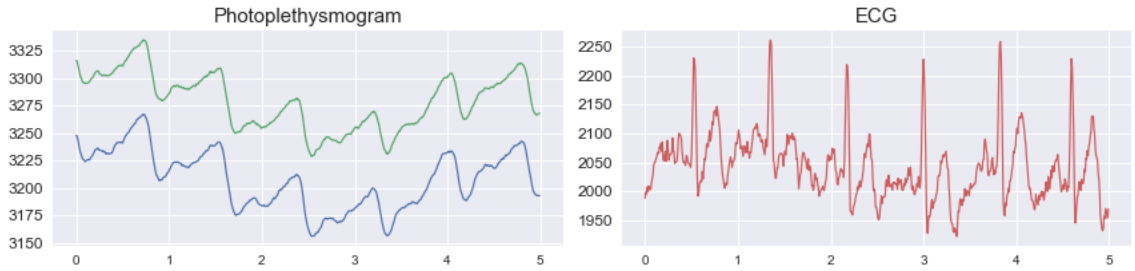

2. Dataset, Data Representation and Preparation

Our paper is a "proof of concept", hence our dataset consists of a sample taken from one healthy male patient only, monitored by professional health care workers. The dataset includes three time series: ECG, PPG and BP. The first two were obtained by a PPG device and the third one was obtained from a Finapres medical device. The PPG, ECG and BP of the subject were recorded at sampling frequency of 512 Hz (the BP was collected with a different frequency and then interpolated to have a measurement at 512 Hz) for 26 minutes of activity comprising both rest and non rest status. Corrupted data has been manually removed to ensure the quality of the signal (about 30% of the data has been dropped). Finally a convolution-based smoother has been applied to remove some noise.
To train and test our algorithm, we created a train/valid/test split with ratios 0.7/0.1/0.2. As customary when dealing with time series, this split was performed without shuffle. The final model to be put into production should predict, for a given 5s long series of data, the average Systolic Blood Pressure (SBP) and Diastolic Blood Pressure (DBP) of the patient, to be updated every second with the new average predicted from the last 5s series of data. As a consequence, as a final step, our train/valid/test sets have been split accordingly. Finally, we have kept only one data point every 4 (so that we can think at our signal as if it was recorded at a frequency of 128 Hz) to ease the computational burden.

3. The model

We write the (cleaned and prepared) time series signals as:

$$X_t = (x_1, \dots, x_T), \quad Y_t = (y_1, \dots, y_T), \quad Z_t = (z_1, \dots, y_T) \quad P_t = (p_1, \dots, p_T)$$

where T is the total number of time steps in our database, and the data is being given with a frequency of f Hz. The first two series X, Y are obtained via PPG, Z is the ECG signal and the third one is the SBP (or the DBP) data (see Sec. 2).
X, Y, Z give rise to a $3 \times T$ dataset that we split in a sequence $\{S_i\}_{i=1}^{N}$ of 3-dimensional time series having length of 5 seconds and with a step of one second (see Sec. 2). Accordingly, we can see each S_i as a matrix of size $3 \times I$, where $I = 5 \cdot f$, whose elements will be denoted as $s_{i,j,k}$ for $j = 0, 1, 2$ and $k = 0, ..., I-1$ (i.e. $s_{i,0,k}$ are obtained from the series X_t, etc.). Each S_i is then transformed into a graph G_i, having set of nodes $\{p_{i,j,k}\}$ in bijection with $\{s_{i,j,k}\}$ and whose construction can be summarised by the following picture.

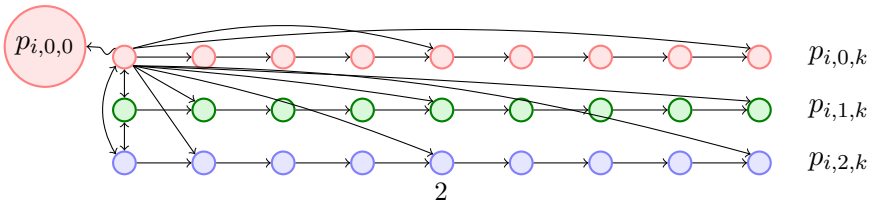

To each node $p_{i,j,k}$ is given the feature $s_{i,j,k}$ (more features can be added) and for any fixed k we create an edge $p_{i,a,k} \to p_{i,b,k+h}$ where $h = 0, 1, d, 2d, 3d, ..., \min(\lfloor \frac{I-k}{d} \rfloor d, f)$, $a, b = 0, 1, 2$ and $a \neq b$ if $h = 0$ for a given, customizable, integer divisor d of f. Note that in the previous diagram only adjacent connections and all the edges of the node $p_{i,0,0}$ are depicted and $d = 4$. We then obtain a series of topologically identical graphs G_i. We obtain a sequence $\{L_i\}_{i=1}^{N}$ of 1-dimensional labels from P_t by taking the mean of the SBP signal over the time windows relative to each S_i. Each graph is then fed to a Graph Neural Network (GNN) that is given by the concatenation of an encoder and a decoder. As possible decoders, we tested an average pooling layer and a flatten+linear layer. The encoder consists of a concatenation of graph convolutions that is described in the following diagram where BN denotes a Batch Normalisation, AF denotes an Activation Function and n is the number of layers:

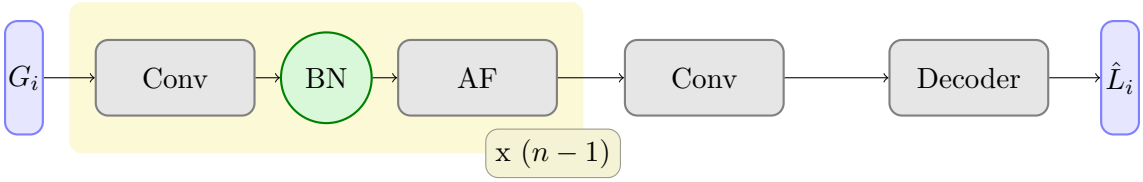

In our validation, we tested as convolutional layers the GCN, Sage and GAT convolutions (Kipf and Welling (2017), Hamilton et al. (2017), Velickovic et al. (2018)) and, when we used GAT, we did not use a BN layer. We call the resulting architectures TGCN, TSAGE and TGAT. Compared with other models traditionally used to model time series data (e.g. Recurrent Neural Networks, see Goodfellow et al. (2016)), the structure of our model has many similarities with recent Temporal Convolutional Networks (TCNs, see Bai et al. (2018)).

4. Experimental Results and Discussion

Model	MAE ± std (SBP)	MAE ± std (DBP)
TGCN	15.71 ± 10.95	6.66 ± 2.22
TSAGE	18.7 ± 9.5	3.47 ± 1.8
TGAT	**6.34 ± 3.71**	**1.87 ± 1.49**
DL benchmark (different dataset)	12.51 ± 12.61	8.3 ± 9.84
XGBoost (our dataset)	6.62 ± 4.11	1.98 ± 1.54

We summarise in this table some results we obtained running our model on our dataset. As a metric, we report the Mean Absolute Error (MAE) and the Standard Deviation (std) in mmHg. We then choose $d = 8$, we use a flatten+linear decoder, we set the number of layers to be equal to d and we add the time series of the derivatives of the given signals as features. As it can be expected, the TGAT architecture is the one displaying best results among our models.

The scores displayed are obtained on our test set by averaging the results of 10 different trainings of the chosen model on our train set.

As a benchmark, we record for reference the best result obtained in Schrumpf et al. (2021) which tests popular Deep Learning models on raw PPG data for BP prediction (with both a dataset and a problem different from ours) and we test the popular XGBoost on our dataset (which is a reasonable benchmark, see Che et al. (2019)). We include the results obtained from our models on the Diastolic Blood Pressure (DBP) as well.

References

M. Amini, F. Zayeri, and M. Salehi. Trend analysis of cardiovascular disease mortality, incidence, and mortality-to-incidence ratio: Results from global burden of disease study 2017. *BMC Public Health*, 21, 2021.

Shaojie Bai, J. Zico Kolter, and Vladlen Koltun. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *arXiv*, 2018.

Michael M. Bronstein, Joan Bruna, Taco Cohen, and Petar Velickovic. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv*, 2021.

Xinze Che, Min Li, Wenbo Kang, Fujun Lai, and Jiaxin Wang. Continuous blood pressure estimation from two-channel PPG parameters by xgboost. *IEEE International Conference on Robotics and Biomimetics, ROBIO*, 2019.

W. Chen, T. Kobayashi, S. Ichikawa, Y. Takeuchi, and T. Togawa. Continuous estimation of systolic blood pressure using the pulse arrival time and intermittent calibration. *Medical and Biological Engineering and Computing*, 28:569–574, 2000.

Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016.

William L. Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems*, pages 1024–1034, 2017.

Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *5th International Conference on Learning Representations, ICLR*, 2017.

S. S. Lim, T. Vos, A. D. Flaxman, G. Danaei, K. Shibuya, M. A. AlMazroa H. AdairRohani, M. Amann, H. R. Anderson, K. G. Andrews, and et al. A comparative risk assessment of burden of disease and injury attributable to 67 risk factors and risk factor clusters in 21 regions, 1990–2010: a systematic analysis for the global burden of disease study 2010. *The lancet*, 380(9859):2224–2260, 2013.

X. R. P. Su, Y. T. Ding, J. Zhang, F. Liu, Miao, and N. Zhao. Long-term blood pressure prediction with deep recurrent neural networks. *IEEE EMBS International Conference on Biomedical & Health Informatics (BHI)*, pages 323–328, 2018. doi: doi:10.1109/BHI.2018.8333434.

Fabian Schrumpf, Patrick Frenzel, Christoph Aust, Georg Osterhoff, and Mirco Fuchs. Assessment of deep learning based blood pressure prediction from PPG and rppg signals. *IEEE Conference on Computer Vision and Pattern Recognition Workshops, CVPR*, 2021.

Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks. *6th International Conference on Learning Representations, ICLR*, 2018.

Y. T. Zhang, Y. L. Zheng, H. Y. Lin, W. H.and Zhang, and X. L. Zhou. Challenges and opportunities in cardiovascular health informatics. *IEEE Trans. Biomed. Eng.*, 60: 633–642, 2013.