How analog to digital conversion shapes heart rate information in PPG data

Francesco Lässig

Short Project
Institute of Neuroinformatics
May 2021

Supervisors: Dr. Hannah Bos, Prof. Dr. Giacomo Indiveri

Universität Zürich[UZH]

ETH zürich

Abstract

Asynchronous Delta Modulation (ADM) is a non-linear, threshold-based transformation used to create spike trains from analog signals. To analyze the effect of the ADM on PPG recordings from the DaLiA dataset, we explored the feasibility of inferring the heart rate from signals that were reconstructed from ADM spike trains and compared those results to the original PPG data. More specifically, we computed mutual information and R2 score (obtained by training a linear regressor predicting heart rate) based on discrete fourier transforms of the ADM reconstructions / raw PPG signals. These measurements were performed on spike trains of different frequencies by varying the ADM threshold parameter. Our results show that, for sufficiently high spiking rates, the ADM reconstruction resulted in overall higher mutual information and R2 score between frequency-domain feature vectors and ground truth heart rate than the raw PPG signal. We found evidence to support the hypothesis that this gain in mutual information and R2 score can mainly be attributed to an amplitude-normalizing effect of the ADM. We also demonstrate that amplitude normalization leads to similar effects as common machine learning preprocessing steps, when looking at feature vectors in the frequency domain.

Contents

1 Introduction

1.1 Goal

The goal of this project is to analyze spike trains generated by asynchronous delta modulation from PPG signals in terms of how much information (term used loosely) they contain about the heart rate. In particular, we want to investigate how the threshold parameter, and thus the mean spiking rate, affects the informational content. We also want to investigate mechanisms that might be responsible for changes in information.

1.2 Motivation

The more general aim of this project is to contribute to an ongoing effort of developing a neuromorphic device that detects abnormal heart rate variability in elderly people. Spiking neural networks running on neuromorphic chips are well-suited for this task due to multiple reasons. For one, SNNs running on neuromorphic chips have been shown to be powerful pattern-detection systems, in particular in the context of physiological data [Bur+21]. Moreover, neuromorphic synapses with exponentially decaying currents implement low-pass filters that are in theory able to linearly filter out frequencies containing noise. Moreover, because signal processing happens in real-time, thanks to time representing itself, and because of fast response times, thanks to parallelism and device mismatch [IS19], neuromorphic circuits are very fast at responding to relevant stimuli. Lastly, and perhaps most crucially, the power consumption of such a neuromorphic device would likely be vastly lower than that required for a functionally equivalent digital device. This fact can largely be attributed to the data-driven nature of neuromorphic processing (i.e. no energy is consumed if no data is coming in) in combination with encoding strategies that are sparse in space and in time [IS19].

1.3 PPG and ECG

PPG. Photoplethysmography is a technique used to measure changes in blood volume in small blood vessels [All07]. The technology is based on measuring the fraction of light from a light source that is reflected back from microvascular tissue it is shone on. The amount of light that is reflected varies as the blood volume in the microvascular tissue fluctuates. In most cases, this can be measured non-invasively. Generally speaking, the most salient component of the PPG signal is a pulsatile waveform that oscillates with a fundamental frequency that corresponds to the heart rate. However, the PPG signal contains many more components that are not causally linked to the heart beat. These 'noise' components can be due to factors such as respiration or body movements.

ECG. An electrocardiogram is a measurement of voltage differences across different parts of the human body induced by the electric activity in the heart

associated with the cardiac cycle [Wik21]. This measurement can be used to infer the heart rate of a person, which is exactly how the ground truth heart rate labels of the dataset we used were generated.

Figure 1: ECG next to clean PPG [All07]

When the PPG signal is mostly unobstructed by noise, the relationship to the ECG becomes visually evident, as seen in Figure 1.

1.4 ADM

Neurons in neuromorphic chips such as the DYNAP-SE neuromorphic chip are event-based and communicate with each other by using spikes [Mor+17]. PPG, however, is an analog signal. This means that, in order for a spiking network running on a neuromorphic device like the DYNAP-SE chip to process PPG data, we have to convert the analog signal to a spike train. This is where Asynchronous Delta Modulation (ADM) comes into play. A delta modulation A/D converter circuit implemented as described by Corradi and Indiveri [CI15] converts an analog signal to a spike train by generating events when the signal value changes up or down by specified amounts. For the purposes of this project, we created a software simulation of the ADM, which is further described in Section 2.1.

1.5 PPG-DaLiA Dataset

PPG-DaLiA, short for PPG dataset for motion compensation and heart rate estimation in Daily Life Activities, is a dataset introduced by Attila Reiss et al. [Rei+19]. Among other data, it consists of PPG measurements of 15 subjects performing the same 8 activities in sequence, as well as corresponding heart rate ground truth labels that were extracted from ECG measurements. The heart rate was extracted from the ECG data by using an 8 second wide sliding window that moves across the signal at 2 second intervals. The dataset was introduced to benchmark algorithms that infer heart rate from PPG signals.

2 Methods

2.1 ADM Implementation

Software Implementation. To generate a spike train from an analog signal, we used a software implementation of the Asynchronous Delta Modulation A/D Converter described by Corradi and Indiveri [CI15]. The exact implementation was based on the delta-modulation spike generation algorithm used by Nik Dennler [Den20]. The mechanism of this digital realization of the ADM is simple: As we iterate through the values of the sampled input signal, we check whether the value is bigger than a reference value plus a given threshold. If yes, we emit an 'up' spike (add a timestamp to the 'up' spike train) and reset the reference value to the current signal value. We do the same if the current value is lower than the reference value minus a threshold and emit these 'down' spikes in a different channel . At the beginning, the reference value is simply set to the first value of the signal. Optionally, a refractory period parameter can be set to prevent spike emissions and changes to the reference value for a given time window after each spike. However, regardless of the refractory period, spike times are restricted to the sampling times of the input signal. This means that, in practice, the refractory period of this software implementation has a minimum which is equal to the sampling period.

Our Settings. The sampling frequency of the PPG signal is 64 Hz. In the following, we fix the refractory period at zero (resulting in an effective refractory period of 1/64 seconds) and investigate the effect of the threshold parameter. Although 'up' and 'down' thresholds can be chosen independently, for the purposes of this project we always set them to the same value. Lower thresholds lead to higher mean spiking frequencies of the ADM, while the opposite holds for higher thresholds. The spiking thresholds were chosen in such a way as to create spike trains with similar mean spiking frequencies across subjects with different PPG value ranges.

2.2 Signal Reconstruction from Spike Train

Given the fact that we extract frequency-domain features from analog signals to measure how much information they contain about heart rate (as described in Section 2.3), we need to convert the ADM spike train back to an analog signal in order to be able to compare its informational content to its raw PPG counterpart from which it was created.

Gaussian 1D Filtering. To create a smooth analog signal, the spike train is convoluted with a 1D Gaussian filter with standard deviation sigma.

Gaussian Sigma. In general, we choose sigma as 10 times the sampling period. The effect of other sigma values is briefly explored in in Section 3.1.1.

2.3 PPG Spectrogram

The heart rate is easiest inferred from PPG by transforming the signal to the Fourier domain. So, to measure the informational content of the analog signals about the heart rate, we create a set of feature-label pairs, where the features correspond to discrete Fourier transforms of slices of the signal and the labels to the average heart rates in the corresponding slices. Using these preprocessing steps, Reiss et al. [Rei+19] successfully predicted the heart rate. As in the PPG DaLiA dataset, we choose a 8 second sliding window with a 6 second overlap.

2.4 Evaluation Metrics

We use two metrics to quantify the informational content of raw PPG data and ADM reconstructions: Mutual information and cross-validated R2 score of a linear regression. Both of them relate signal data in the frequency domain to scalar heart rate data that was obtained from ECG measurements. These measures are indicators of how well the heart rate can be inferred by a linear readout of the feature vectors.

2.4.1 Mutual Information

Mathematics. Mutual information is a measure for how much we can find out about a given random variable X by observing another random variable Y [LR09]. This quantity is minimized if X and Y are independent. This can be easily seen in the definition of mutual information as the KL Divergence between the joint distribution of X and Y, and the product of the marginal distributions of X and Y seen in Equation 1.

$$I(X;Y) = D_{KL}\Big(P_{XY}(x,y)||P_X(x)P_Y(y)\Big) \tag{1}$$

Approximation. The problem with computing mutual information is that we do not have direct access to the random variables. We have to approximate them in some form. In this project we use the *scikit-learn* implementation of an approximation method (sklearn.feature_selection.mutual_info_regression) based on a k-nearest-neighbor algorithm described by Kraskov et al. [KSG11]. The number of neighbors parameter was set to 10 for all mutual information estimations. This way of approximating mutual information only works between 2 scalar values. Since we are dealing with vectors in the frequency domain, we compute the mutual information for every frequency component individually. To get to a single value that lets us compare mutual information of different signals, we take the mean of the mutual information values of individual frequencies. We shall use the acronym 'MMI' when referring to mean mutual information. Please note that the different frequency components are not independent of each other and therefore simply adding up mutual information from all of the frequencies is not a measure for total mutual information [MLP12].

2.4.2 Linear Regression R2 Score

We perform closed-form linear regression on signal data in the frequency domain to predict the ground truth heart rate for every time window. To make sure that our model captures a generalized relationship between input and output we use randomized 5-fold cross-validation. To evaluate the predictions of the linear regression model we use the R2 score given by Equation 2.

$$R^2(y, \hat{y}) = 1 - \frac{\sum_i (y_i - \hat{y}_i)^2}{\sum_i (y_i - \bar{y})^2} \tag{2}$$

Here $\hat{y}$ corresponds to the predicted values, y is equal to the ground truth values, and $\bar{y}$ represents the mean over all ground truth values. The result is a measure of the fraction of the variance of the ground truth labels that is accounted for by the predictions. Scaling has no effect on this score, which makes it valuable for comparing heart rate predictions of different subjects, since different people can have different baselines. We shall abbreviate cross-validated linear regression R2 score to simply 'R2' or 'R2 score' in the following sections.

3 Results

3.1 Signal Reconstruction

3.1.1 Effect of Standard Deviation of Gaussian Filtering

In addition to parameters associated with the ADM procedure, the signal reconstruction includes another degree of freedom, namely the sigma of the Gaussian 1D filter used to transform the spike train back to an analog signal. In the following, we demonstrate that larger sigmas yield suppression of high frequencies.

Original Signal. To illustrate at least some of the effects that are associated with the value of the Gaussian sigma, let us take subject 3 as an example. Figure 2 shows the PPG spectrogram with the main frequency component corresponding to the heart rate (marked with a gray dashed line). We can faintly see harmonics in the upper part of the plot that correspond to frequency components with twice the frequency of the heart rate.

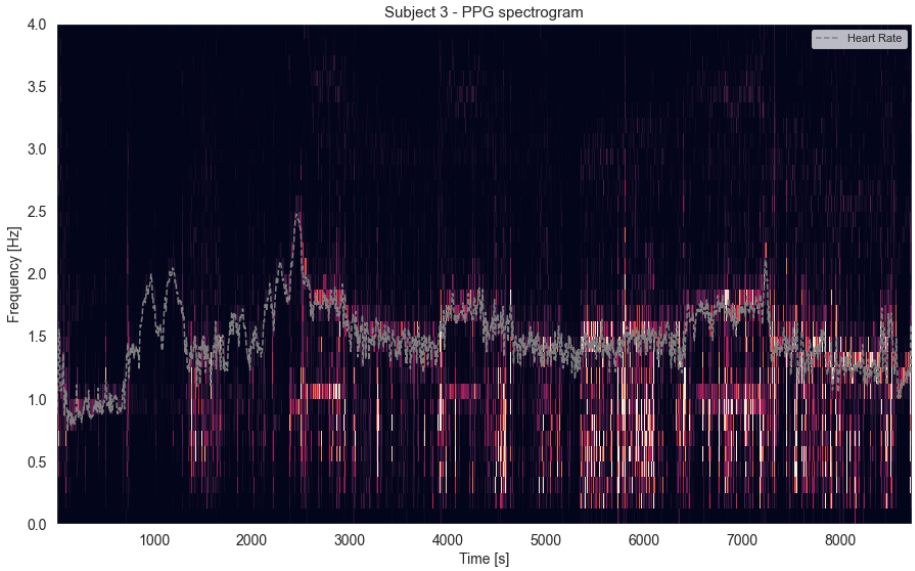

Figure 2: Subject 3 - Raw PPG spectrogram.

The fact that higher frequencies in the frequency domain encode the heart rate in addition to the frequency components in the range of the actual heart rate is reflected by the two peaks in the mutual information graph of the raw PPG signal shown as the blue lines in Figure 5.

High Sigma. If we pass this signal to the ADM (with a fixed spiking threshold leading to a mean spiking rate of 37 Hz) and reconstruct an analogue signal with a Gaussian sigma of 10, we can tell by looking at the spectrogram shown in

Figure 3 that higher frequency components are greatly reduced. In other words, the ADM + Gaussian reconstruction exhibits a low-pass filtering behavior. This mirrors the 'Gaussian blur' effect associated with 2D Gaussian filters in image processing. This effect is also visible in the mutual information graph shown in Figure 5b as the second peak becomes less pronounced.

Figure 3: Subject 3 - Spectrogram of ADM reconstruction using sigma=10.

Low Sigma. On the other hand, if we set the Gaussian sigma to a lower value of 1, we see a completely different behavior, as seen in Figure 4. Here we do not just retain harmonics, but it looks like they are amplified as well. The mutual information graph from Figure 5a also reflects this. If we take the mean of the mutual information values across all frequencies, the second reconstructed signal (with the gaussian sigma set to 1) will win over the first one. This is largely due to the increase of the second peak in mutual information.

Effect on Mean Mutual Information. While using a lower sigma here might be the better option regardless of whether we factor in the mutual information gain of the harmonics or not, this example reveals a drawback of looking at mean mutual information across frequencies as a metric. This is because simply duplicating the same information across different frequency bands will yield a higher score. As a result, there is a danger of making this a competition about which signals contain more harmonics of the heart rate frequencies. To avoid this problem, in the following analyses we decided to constrain the frequencies considered for the computation of MMI / R2 to the range between 0.5 Hz and 2.5 Hz. This way the contribution of harmonics to mutual information can be at least somewhat restricted. However, while this likely reduces the effect of duplicate information across frequency bands, it does not prevent it from

Figure 4: Subject 3 - Spectrogram of ADM reconstruction using sigma=1.

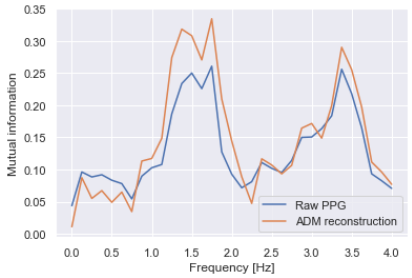

(a) ADM reconstruction with sigma=1.

(b) ADM reconstruction with sigma=10.

Figure 5: Mutual information across frequencies for both original PPG and ADM reconstruction using two different gaussian sigmas.

happening in the remaining range. As a result, mean mutual information needs to be treated as an indicator of overall mutual information with caution.

3.1.2 Amplitude Normalization

Visual Observation. Figure 6a shows a slice of subject 8's PPG measurement containing a combination of high-amplitude waves (at 3 locations) and low-amplitude waves in-between. The relative differences in amplitude between the two types of waves are quite high. Looking at the signal that was reconstructed

from the ADM spike train shown in Figure 6b, we can see that this relative difference in amplitude has been reduced. In other words, it looks like the amplitude of the signal varies less over time.

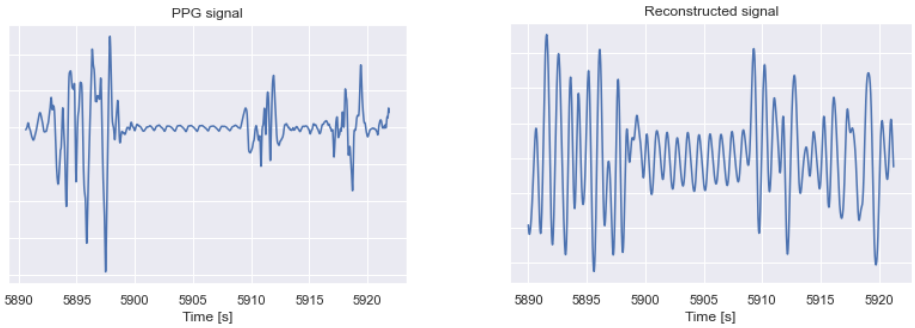

(a) Raw PPG Signal.

(b) ADM reconstruction.

Figure 6: 30 second slices of raw PPG and ADM reconstruction (mean spiking rate: 37 Hz) signals hinting towards amplitude normalization.

Measuring Amplitude Variability. To quantify this amplitude normalization effect more precisely, and to look at the whole PPG measurement of subject 8 and not just this particular time slice, we applied the following procedure to get a measure for amplitude variability: We divide the PPG signal into many consecutive, non-overlapping time windows and compute the variance for each. The result is a time series of variances over time (Figure 7a). The same can be done with the reconstructed signal, as shown in Figure 7b. If an oscillating signal has higher amplitudes in a particular time window A than it does in a different time window B, then the variance of the signal in A should be higher than the variance of the signal in B. This means that, the more the values in Figures 7a and 7b fluctuate, the more the amplitude varies over time. Taking the variance of these newly created time series for both signals should give us scalar measures for how much each of the signal's amplitudes changes over time. To ensure a fair comparison, each of the signals was first scaled to zero mean and unit variance before computing the variance time series. The time window size was chosen to be 2 seconds, which corresponds to the amount by which the sliding window moves at each step when creating the frequency-domain features used to compute MMI / R2.

Figure 8 shows amplitude variability (the variance of the time series shown in Figure 7) for all subjects for both raw PPG and ADM reconstruction. The ADM spiking threshold was chosen to be 10% of the absolute values of each signal, which corresponds to a mean spiking rate of roughly 37Hz. Every signal was first scaled to zero mean and unit variance before computing the variance across variances. Evidently, ADM reconstruction decreases the signal amplitude variability (measured as the variance of signal variances within 2 second time windows) for all subjects.

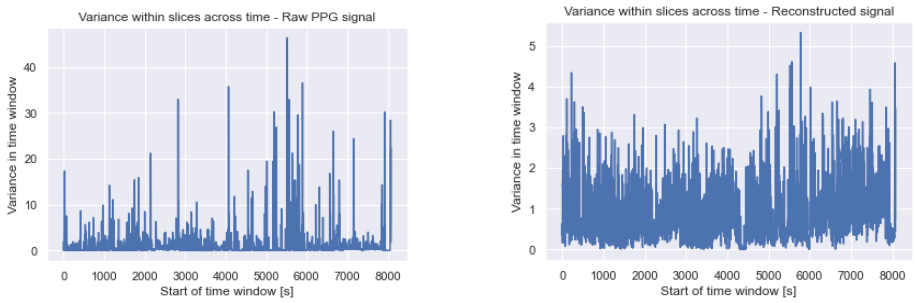

(a) Raw PPG Signal.

(b) ADM reconstruction.

Figure 7: Time series created by dividing the signals into consecutive, non-overlapping time windows of 2 seconds and computing their variance.

Amplitude Normalization Hypothesis. We hypothesize that amplitude normalization is at least partially responsible for the gains in MMI / R2 score observed in Section 3.3. To test this hypothesis, we preprocessed the data by scaling individual 2 second time windows of the PPG signals to unit-variance before creating the spectrogram, which is essentially a way to artificially minimize the above measure of amplitude variation. We will refer to a signal that was preprocessed like this as 'amplitude normalized'. In Section 3.3 we will look at the impact of this procedure on mutual information and R2 score and how it compares to ADM reconstruction. But before that, we show that amplitude normalization approaches behaviors of common machine learning preprocessing techniques.

Feature Vector Magnitude Variability. Figure 9 shows how a reduction in amplitude variability (both through ADM reconstruction and artificial amplitude normalization) translates to lower variability of the magnitude of the frequency-domain feature vectors, i.e. the variance of the L2 norms of the discrete Fourier transforms of the individual time windows. These are the feature vectors that are used for the computation of mutual information and R2 scores.

We can see that amplitude normalization greatly reduces the variability of feature vector L2 norms, and that ADM reconstruction approaches this behavior. Scaling every feature vector to unit-length would minimize this metric.

Please note that the window size used for the feature vector magnitude variability analysis corresponds to the window size which was used throughout this project to create feature vectors, namely 8 seconds. This is different from the window size of 2 seconds that we previously used for the computation of variances across variances. The amplitude normalization window size was chosen to be a smaller value because we expected it to be more effective at reducing amplitude variability due to the reasoning explained in Section 3.3.4.

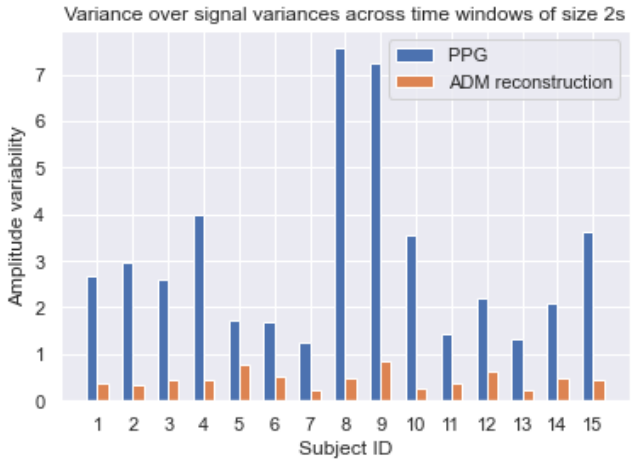

Figure 8: Amplitude variability measured as variance over variances of consecutive time windows.

Feature Variance Variability. Figure 10 shows how a reduction in amplitude variability translates to lower variability of the variances of individual features in our frequency-domain feature vectors (see Equation 4).

Artificial amplitude normalization greatly reduces the variability of feature variances, and ADM reconstruction approaches this behavior. Scaling every feature across datapoints to have unit-variance would minimize this metric.

Why look at feature vector magnitude variability and variability of variances of different features specifically? Because minimizing these measures is exactly what is commonly done with two common preprocessing steps from the machine learning world: Scaling all feature vectors to unit-length and scaling all features (across datapoints) to unit-variance.

3.2 Amplitude Normalization and Refractory Period

Despite the fact that in our software implementation of the ADM we set the refractory period parameter to a value of 0, there is still an implicit refractory period present. This is because our ADM implementation cannot generate spikes with a frequency higher than the sampling frequency. In other words, we are dealing with a refractory period of 1/64 seconds. This introduces an upper bound to any ADM reconstruction signal amplitude, because in the most extreme case there is a spike at every sampling point, but we can never go higher than that.

Let us consider the example of a sine wave: If the first half of a sine wave completely saturates the spike train (meaning that the ADM spikes at 64Hz), then it will not make a difference if we double the amplitude of the sine wave. The

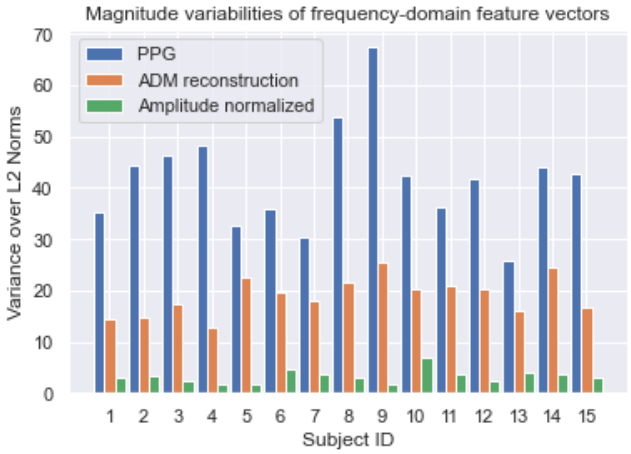

Figure 9: Feature vector magnitude variability. Y axis computed from feature vectors as seen in Equation 3, where x_{ij} corresponds to the j'th frequency component in the i'th feature vector and n is equal to the number of feature vectors.

$$y = Var[\{||x_{1,:}||_2, , ..., ||x_{n,:}||_2\}] \tag{3}$$

output amplitude of the reconstructed signal stays the same because the spike train does not change. Figure 11 demonstrates the amplitude normalization effect on a toy example.

3.3 Mutual Information / R2 Score

3.3.1 Overall

ADM Spiking Rate and MMI / R2. As a first step, we measured the relationship between ADM mean spiking rate and MMI / R2 over the whole dataset. This means that every time window and its associated heart rate label were taken as a data point, regardless of subject or activity. The effect of the average spiking rate of the ADM-generated spike train on the MMI / R2 of the reconstructed signal is shown in Figure 12 by the blue curve.

For high threshold values, which translates to low mean firing rates of the ADM, the MMI / R2 between the reconstructed signal and the ground truth heart rate are quite low. This makes sense intuitively, since higher ADM thresholds yield coarser granularity. The fact that MMI / R2 increase with higher spiking rates is also not surprising given the same reasoning. However, considering that the ADM cannot add any new information to the signal, we expected MMI / R2 of the reconstructed signal to plateau roughly at the height of the MMI / R2 of the

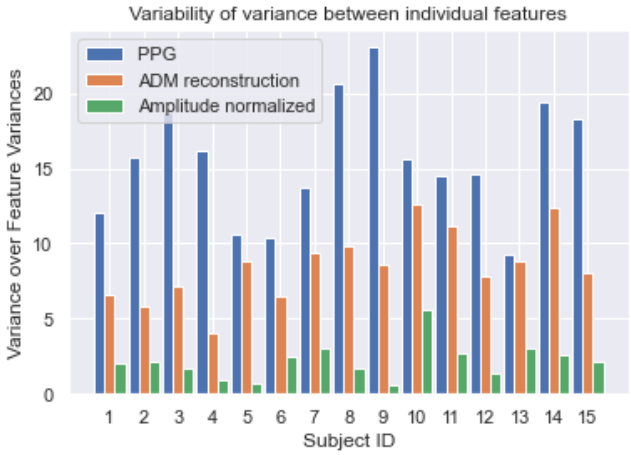

Figure 10: Feature variance variability. Y axis computed from feature vectors as seen in Equation 4, where x_{ij} corresponds to the j'th frequency component in the i'th feature vector and k is equal to the number of frequency components.

$$y = Var[\{Var[x_{:,1}], ..., Var[x_{:,k}]\}] \tag{4}$$

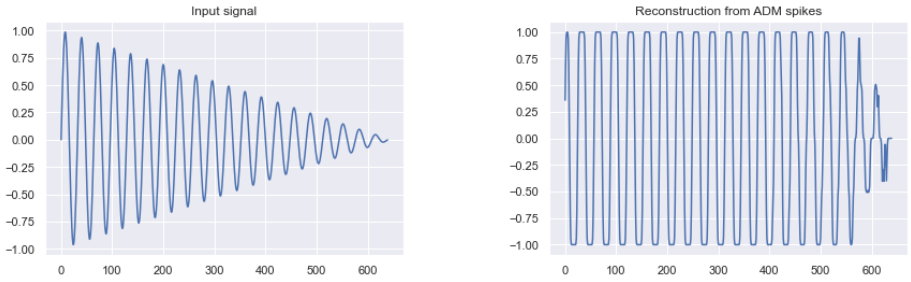

(a) Input signal.

(b) ADM reconstruction.

Figure 11: A toy example showing how ADM reconstruction can lead to amplitude normalization. The ADM spiking threshold for both up and down channels was set to 0.02.

raw PPG signal. Evidently, the data did not match our expectations. Please note that, while the ADM cannot possibly add new information to the signal, it can improve the signal-to-noise ratio and thus increase mutual information. The MMI / R2 scores start to surpass the ones from the raw signal at a mean spiking frequency of around 10 Hz. This minimum mean spiking rate required of the ADM for the reconstructed signal to surpass the raw signal in terms of MMI / R2 will be referred to as 'faithful reconstruction threshold' FRT, which

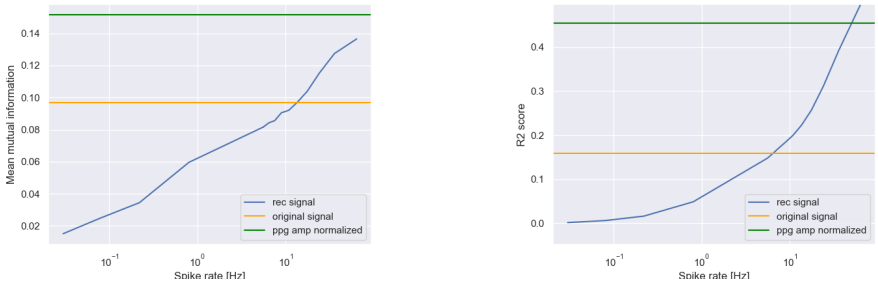

(a) Mean mutual information.

(b) R2 score.

Figure 12: Effect of mean ADM spiking rate on MMI / R2 of reconstructed signal. Raw PPG and amplitude normalized PPG (normalized using 2 second windows) included for comparison.

is measured in Hz.

Mutual Information Across Frequencies. Looking at the blue and orange curves of Figure 13, we can see how our ADM signal reconstruction affects mutual information within frequency bands. Our analysis of the mutual information suggests that the improvement in mutual information of the reconstructed signal compared to the raw PPG signal cannot simply be attributed to a linear filtering operation that removes frequencies that are likely to contain motion artifacts or other types of noise. This is because, as is visible in Figure 13, every individual frequency in the relevant range (between 1 and close to 3 Hz) shows mutual information gains in the reconstructed signal. In Section 3.4.2 we further validate the claim that linear filtering in the frequency domain is not responsible for the gains in MMI / R2. This observation gives further credibility to the amplitude normalization hypothesis, since that would be a different operation than linear frequency filtering.

ADM reconstruction vs. Artificial Amplitude Normalization. The green curves in Figure 12 show the MMI / R2 scores obtained by the PPG signal after artificial amplitude normalization as described in Section 3.1.2. For mean mutual information, it seems like the ADM reconstruction score approaches the amplitude normalization score as the spiking rate increases, giving further credibility to the amplitude normalization hypothesis. We also see that the normalization replicates much of the regression score improvements.

The green curve in Figure 13 also shows mutual information within frequency bands for the amplitude normalized PPG signal. We can see that the amplitude normalization MI correlates with the MI of the ADM reconstruction and constitutes an upper bound at every point.

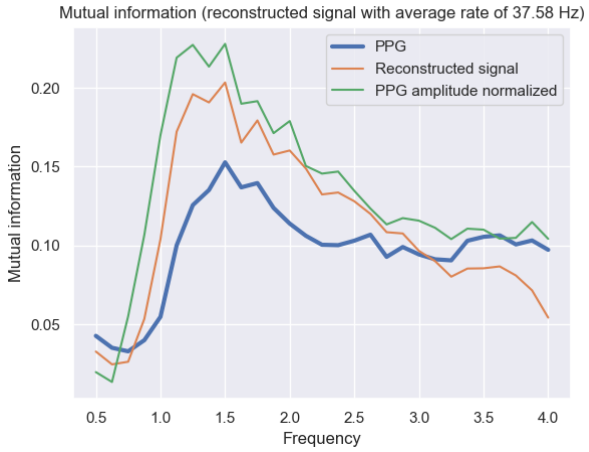

Figure 13: Mutual information of individual frequency components.

3.3.2 Across Subjects

ADM Reconstruction MMI / R2 Gains. The blue bars in Figure 14 show, for each subject, the absolute improvements made by the best ADM reconstruction (among different threshold parameter selections) in terms of MMI / R2 compared to the scores computed from the raw PPG signal. We can see that, while ADM reconstruction leads to consistent R2 score gains, MMI does not increase in every subject. However, overall we see mostly improvements which are expressed to varying degrees depending on the subject.

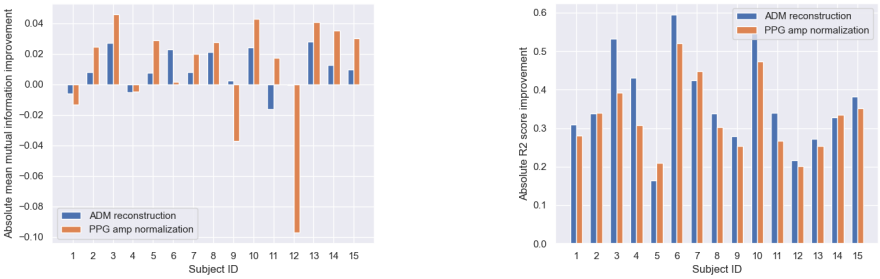

(a) Mean mutual information.

(b) R2 score.

Figure 14: Difference between best ADM reconstruction (blue) and artificially amplitude normalized PPG (orange) scores, and raw PPG scores for all subjects.

FRT and Mean Heart Rate. Not only do we see variance in terms of the best ADM reconstruction improvement across subjects, but we also observe differences in FRTs. As a result, we hypothesized that higher heart rates might on average need a higher ADM spiking frequency, implying a higher FRT. To

Metric	Spearman correlation coefficient	P-value
MMI	0.54	0.089
R2 score	0.78	0.00063

Table 1: Spearman correlations between FRT and mean heart rate across subjects.

explore this possibility, we computed the FRT for every subject and compared them with the respective subjects' mean heart rates, shown in Figure 15. The Spearman correlation coefficients are listed in Table 1. Although the P-value for the mutual information case might not be completely satisfactory, these results point toward the tendency hypothesized above. The reason why there are fewer points on the left plot is because the ADM reconstruction did not reach the PPG score for all subjects.

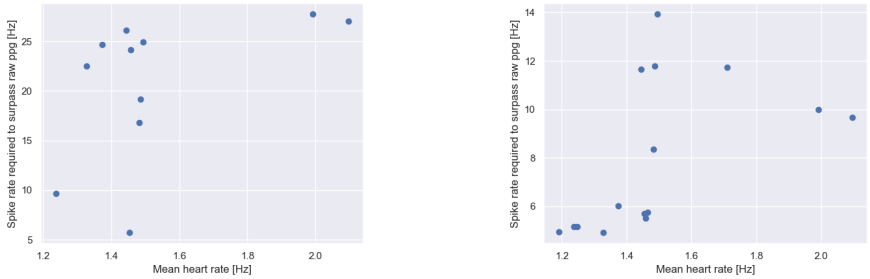

(a) Mean mutual information.

(b) R2 score.

Figure 15: The faithful reconstruction thresholds (FRT) in terms of mean ADM spiking rates per subject.

ADM reconstruction vs. Artificial Amplitude Normalization. The orange bars in Figure 14 visualize the improvements in MMI / R2 made by artificial amplitude normalization. In the MMI case, we see that for 9 out of the 11 subjects the PPG normalization can (at least) replicate the gains of the ADM reconstruction. In the regression score case, amplitude normalization replicates most of the gains in all subjects individually.

3.3.3 Across Activities

Big Picture. Figure 16 shows that ADM reconstruction improves R2 scores of all activities and MMI values of all activities except number 2, where it makes no difference. If we look at the scores of artificial amplitude normalization, all MMI gains of ADM reconstruction are (at least) matched, except for activity 3. A similar pattern emerges from the regression scores: Amplitude normalization mostly replicates the gains in all activities except number 3.

Zooming In. If we just look at individual activities of a single subject, we end

up with much smaller data sets used to compute MMI / R2. However, as seen in Figure 17 (which was created using data only from subject 7), MMI and R2 still show similar improvements from ADM reconstruction and artificial amplitude normalization.

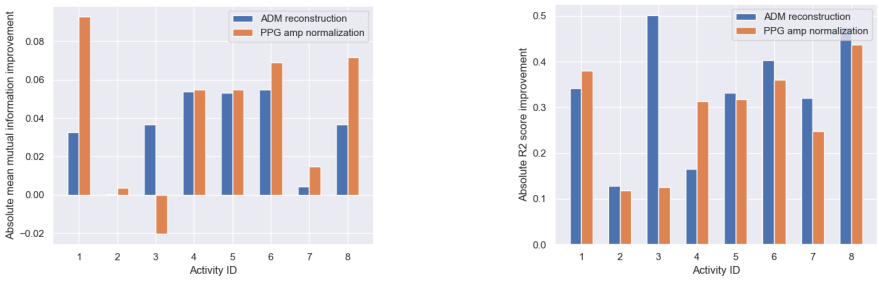

(a) Mean mutual information. (b) R2 score.

Figure 16: Difference between best ADM reconstruction (blue) and artificially amplitude normalized PPG (orange) scores, and raw PPG scores for all activities.

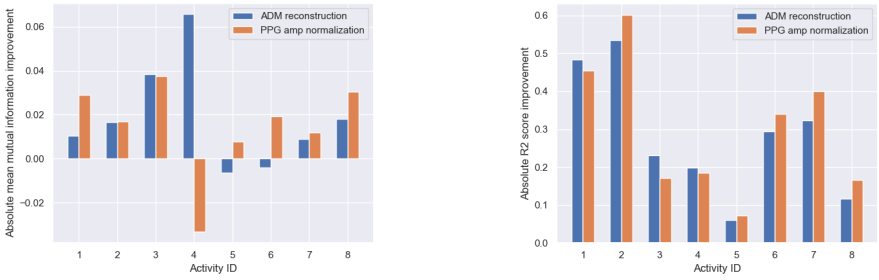

(a) Mean mutual information. (b) R2 score.

Figure 17: Difference between best ADM reconstruction (blue) and artificially amplitude normalized PPG (orange) scores, and raw PPG scores for all activities using only data from Subject 7.

3.3.4 Effect of Amplitude Normalization Window Size

In all of the above results, amplitude normalization was carried out by using a window size of 2 seconds. Figure 18 shows how varying this window size affects the MMI / R2 scores that are obtained from the artificially amplitude normalized signal, and how this compares to the best scores obtained from the ADM reconstruction and the score obtained from the raw PPG.

Although for mutual information a window size of 1 second seems to be a sweet spot, the general trend is that larger window sizes lead to lower MMI / R2 scores. This can likely be attributed to the following effect: The larger the normaliza-

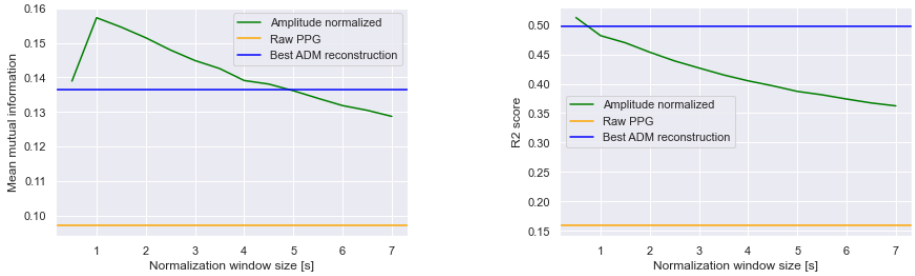

(a) Mean mutual information.

(b) R2 score.

Figure 18: MMI / R2 for artificially amplitude normalized PPG signal using different window sizes. Best ADM reconstruction and raw PPG included for comparison.

tion time window, the higher the chance that the contribution of particularly high or small amplitudes get cancelled out by other waves within the same time window. Let us consider the example of having large-amplitude waves immediately followed by small-amplitude waves. If the large and small amplitude waves fall within separate time windows, their amplitudes will be scaled down and up respectively. However, if we increase the time window size and they fall into the same window the two waves can only be scaled by the same factor, so either both get scaled up or down. In the extreme case, the time window is as large as the whole signal duration, in which case no amplitude normalization can take place.

3.4 Controls

In this section we present some control experiments with the purpose of testing whether the gain in MMI / R2 is really due to the nonlinear threshold operation of the ADM reconstruction process (its amplitude normalization effect in particular), and cannot be explained by a linear filtering mechanism.

3.4.1 Gaussian 1D Filter Alone

To test whether the ADM reconstruction improvements can also be achieved without the ADM, we performed Gaussian 1D filtering with different sigmas directly on the PPG data (instead of the spike trains). Figure 19 shows that by applying this procedure, no significant improvements in MMI / R2 are achieved.

3.4.2 Low-pass Filter Alone

As seen in Section 3.1.1, and also due to properties of the ADM regarding filtering of certain frequencies [Jar95], the whole ADM reconstruction process is expected to exhibit low-pass filtering behavior. To test whether the gains in

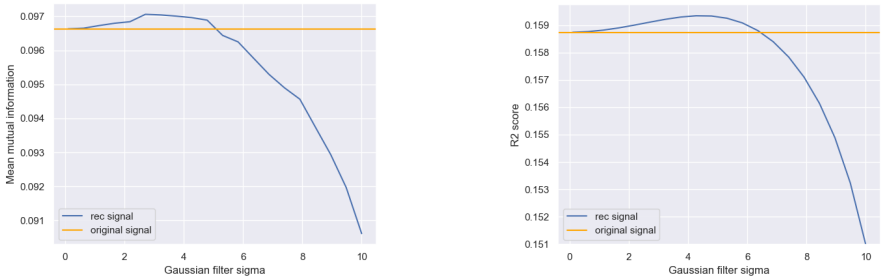

(a) Mean mutual information. (b) R2 score.

Figure 19: MMI / R2 for all PPG data preprocessed using Gaussian filtering like in the ADM reconstruction pipeline.

MMI / R2 scores could be attributed to filtering out higher frequencies with a classical low-pass filter, we filtered the PPG signal using 4th order low-pass butterworth filters with different cutoff frequencies. We then computed MMI / R2 scores of the filtered signal and compared it to the scores of the raw PPG signal. Figure 20 shows that the MMI of the filtered signal approaches the one from the raw PPG as the cutoff frequency increases. In other words, as fewer frequencies get filtered out, the filtered signal's MMI increases, but never reaches and definitely does not surpass the one from the raw PPG signal. For R2, there does seem to be a marginal benefit of filtering the PPG signal with a cutoff frequency of around 2 Hz. However, given the insignificant magnitude of the gain, there is a high chance this can be attributed to random fluctuations of the regression algorithm.

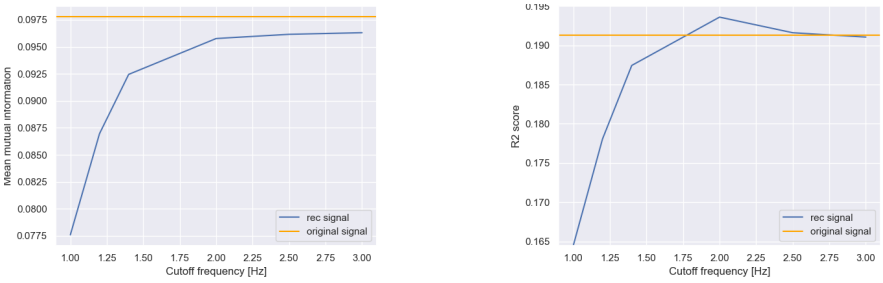

(a) Mean mutual information. (b) R2 score.

Figure 20: MMI / R2 for all PPG data preprocessed using 4th-order butterworth low-pass filters.

4 Discussion

ADM Reconstruction and Feature Normalization. The results show that the ADM reconstruction, at high enough spiking frequencies, has a higher mutual information with the ground truth heart rate and results in better R2 scores when applying linear regression to predict the heart rate. The results also suggest that these improvements in mutual information and R2 score can be mainly attributed to the amplitude-normalizing effect of the ADM reconstruction for the following reasons: Firstly, ADM clearly performs a type of amplitude normalization across all subjects, as shown in Section 3.1.2. Secondly, if we artificially force this amplitude variability metric (that is decreased by the ADM reconstruction) to its minimum value, we can, for the combined data set, achieve the same improvements in mutual information and almost the same improvements in R2 score. Drawing this conclusion would be consistent with the success of two common preprocessing techniques from the machine learning world: scaling feature vectors to unit-length and scaling feature values across data points to unit variance. This is because, as we showed in Section 3.1.2, lower amplitude variability is accompanied by lower variability of the L2 norms of feature vectors created by applying FFT to the signal, and also by lower variability of the variance of these features across data points. Both of these effects approach common standardization / normalization effects of the above mentioned techniques used in machine learning.

Broader Context. Given that the ADM transformation in the end constitutes a preprocessing step to create a signal that the DYNAP-SE neuromorphic chip can process, it only represents a small part of the heart rate variability anomaly detection pipeline that is the overarching aim which this project contributes towards. Still, amplitude normalization can be an important ingredient in a device which captures signals with varying strengths.

Limitations and Future Work. One indicator that amplitude normalization might not be the entire story emerged in Section 3.3.3. While most of the gains could be explained by amplitude normalization in 7 out of 8 activities, activity 3 did not follow that pattern. It would be interesting to find out why activity 3 showed different results and possibly discover a different mechanism by which ADM reconstruction leads to mutual information and R2 score improvements. One way to approach this could be to differentiate between signal and noise components of the PPG and analyze how they are affected by the ADM individually. Such an analysis could also help to quantify the extent to which the MMI / R2 gains are dependent on the specific dataset we used.

Considering the fact that the implicit refractory period played a key role in the amplitude normalization performed by the ADM reconstruction, it seems promising to further investigate the effects of this parameter to either get even better normalization, or to get similar normalization at lower energies. This project only explored a single setting. It might also be fruitful to investigate the relationship between this (and other) ADM parameter(s) and the artificial

amplitude normalization window size that best replicates the MMI / R2 gains of the ADM reconstruction.

Aggregating mutual information across different frequencies by simply taking their mean is problematic because information might be duplicated across frequency bands. This metric could be improved by finding ways of accounting for duplicated information, or by only looking at single dominant frequencies instead of aggregating them.

Although the ADM reconstruction showed improvements in mutual information and R2 scores, these metrics might not be indicative of how well the downstream neuromorphic chip can process this data. In the end, different ADM parameters have to be tested in the context of the end-to-end process of the spiking network to evaluate their effects in a real-world setting.

References

[Jar95] David Jarman. *A Brief Introduction to Sigma Delta Conversion*. [Online; accessed 16-May-2021]. 1995. URL: https://ecee.colorado.edu/~mcclurel/intersilan9504.pdf.

[All07] John Allen. "Photoplethysmography and its application in clinical physiological measurement". In: *Physiological measurement* 28.3 (2007), R1.

[LR09] P. E. Latham and Y. Roudi. "Mutual information". In: *Scholarpedia* 4.1 (2009). revision #186917, p. 1658. DOI: 10.4249/scholarpedia.1658.

[KSG11] Alexander Kraskov, Harald Stögbauer, and Peter Grassberger. "Erratum: estimating mutual information [Phys. Rev. E 69, 066138 (2004)]". In: *Physical Review E* 83.1 (2011), p. 019903.

[MLP12] Alberto Mazzoni, Nikos K. Logothetis, and Stefano Panzeri. *The information content of Local Field Potentials: experiments and models*. 2012. arXiv: 1206.0560 [q-bio.NC].

[CI15] Federico Corradi and Giacomo Indiveri. "A neuromorphic event-based neural recording system for smart brain-machine-interfaces". In: *IEEE transactions on biomedical circuits and systems* 9.5 (2015), pp. 699–709.

[Mor+17] Saber Moradi et al. "A scalable multicore architecture with heterogeneous memory structures for dynamic neuromorphic asynchronous processors (DYNAPs)". In: *IEEE transactions on biomedical circuits and systems* 12.1 (2017), pp. 106–122.

[IS19] Giacomo Indiveri and Yulia Sandamirskaya. "The importance of space and time for signal processing in neuromorphic agents: the challenge of developing low-power, autonomous agents that interact with the environment". In: *IEEE Signal Processing Magazine* 36.6 (2019), pp. 16–28.

[Rei+19] Attila Reiss et al. "Deep ppg: Large-scale heart rate estimation with convolutional neural networks". In: *Sensors* 19.14 (2019), p. 3079.

[Den20] Nik Dennler. "Unsupervised and Online Vibration Anomaly Detection Using Spiking Neural Networks". MA thesis. Institute of Neuroinformatics, 2020.

[Bur+21] Karla Burelo et al. "A spiking neural network (SNN) for detecting high frequency oscillations (HFOs) in the intraoperative ECoG". In: *Scientific Reports* 11.1 (2021), pp. 1–10.

[Wik21] Wikipedia contributors. *Electrocardiography — Wikipedia, The Free Encyclopedia*. [Online; accessed 16-May-2021]. 2021. URL: https://en.wikipedia.org/w/index.php?title=Electrocardiography&oldid=1019872377.