pyVHR: a Python framework for remote photoplethysmography

Giuseppe Boccignone [1] , **Donatello Conte** [2] , **Vittorio Cuculo** [1] , **Alessandro D'Amelio** [Corresp., 1] , **Giuliano Grossi** [1] , **Raffaella Lanzarotti** [1] , **Edoardo Mortara** [1]

[1] PHuSe Lab - Dipartimento di Informatica, Università degli Studi di Milano, Milan, Italy

[2] Laboratoire d'Informatique Fondamentale et Appliquée de Tours, Université de Tours, Tours, France

Corresponding Author: Alessandro D'Amelio
Email address: alessandro.damelio@unimi.it

Remote photoplethysmography (rPPG) aspires to automatically estimate heart rate (HR) variability from videos in realistic environments. A number of effective methods relying on data-driven, model-based and statistical approaches have emerged in the past two decades. They exhibit increasing ability to estimate the blood volume pulse (BVP) signal upon which BPMs (Beats per Minute) can be estimated. Furthermore, learning-based rPPG methods have been recently proposed. The present pyVHR framework represents a multi-stage pipeline covering the whole process for extracting and analyzing HR fluctuations. It is designed for both theoretical studies and practical applications in contexts where wearable sensors are inconvenient to use. Namely, pyVHR supports either the development, assessment and statistical analysis of novel rPPG methods, either traditional or learning-based, or simply the sound comparison of well-established methods on multiple datasets. It is built up on accelerated Python libraries for video and signal processing as well as equipped with parallel/accelerated ad-hoc procedures paving the way to online processing on a GPU. The whole accelerated process can be safely run in real-time for 30 fps HD videos with an average speedup of around 5. This paper is shaped in the form of a gentle tutorial presentation of the framework.

1 pyVHR: A Python Framework for Remote
2 Photoplethysmography

3 **Giuseppe Boccignone**[1], **Donatello Conte**[2], **Vittorio Cuculo**[1], **Alessandro**
4 **D'Amelio**[1], **Giuliano Grossi**[1], **Raffaella Lanzarotti**[1], and **Edoardo Mortara**[1]

5 [1]**PHuSe Lab - Dipartimento di Informatica, Università degli Studi di Milano, 20133 Milan,**
6 **Italy**
7 [2]**Laboratoire d'Informatique Fondamentale et Appliquée de Tours (LIFAT-EA 6300),**
8 **Université de Tours, 37000 Tours, France**

9 Corresponding author:
10 Alessandro D'Amelio[1]

11 Email address: alessandro.damelio@unimi.it

12 ABSTRACT

13 Remote photoplethysmography (rPPG) aspires to automatically estimate heart rate (HR) variability from
14 videos in realistic environments. A number of effective methods relying on data-driven, model-based and
15 statistical approaches have emerged in the past two decades. They exhibit increasing ability to estimate
16 the blood volume pulse (BVP) signal upon which BPMs (Beats per Minute) can be estimated. Furthermore,
17 learning-based rPPG methods have been recently proposed. The present `pyVHR` framework represents
18 a multi-stage pipeline covering the whole process for extracting and analyzing HR fluctuations. It is
19 designed for both theoretical studies and practical applications in contexts where wearable sensors
20 are inconvenient to use. Namely, `pyVHR` supports either the development, assessment and statistical
21 analysis of novel rPPG methods, either traditional or learning-based, or simply the sound comparison of
22 well-established methods on multiple datasets. It is built up on accelerated Python libraries for video and
23 signal processing as well as equipped with parallel/accelerated ad-hoc procedures paving the way to
24 online processing on a GPU. The whole accelerated process can be safely run in real-time for 30 fps
25 HD videos with an average speedup of around 5. This paper is shaped in the form of a gentle tutorial
26 presentation of the framework.

27 1 INTRODUCTION

28 Heart rate variability can be monitored via photoplethysmography (PPG), an optoelectronic measurement
29 technology first introduced in Hertzman (1937), and then largely adopted due to its reliability and non-
30 invasiveness (Blazek and Schultz-Ehrenburg, 1996). Principally, this technique captures the amount of
31 reflected light skin variations due to the blood volume changes.

32 Successively, remote-PPG (rPPG) has been introduced. This is a contactless technique able to measure
33 reflected light skin variations by using an RGB-video camera as a virtual sensor (Wieringa et al., 2005;
34 Humphreys et al., 2007). Essentially, rPPG techniques leverage on the RGB color traces acquired over
35 time and processed to approximate the PPG signal. As a matter of fact, rPPG has sparked great interest by
36 fostering the opportunity for measuring PPG at distance (e.g. remote health assistance) or in all those
37 cases where contact has to be prevented (e.g. surveillance, fitness, health, emotion analysis) (Aarts et al.,
38 2013; McDuff et al., 2014; Ramírez et al., 2014; Boccignone et al., 2020b; Rouast et al., 2017). Indeed,
39 the rPPG research field has witnessed a growing number of techniques proposed for making this approach
40 more and more robust and thus viable in contexts facing challenging problems such as subject motion,
41 ambient light changes, low-cost cameras (Lewandowska et al., 2011; Verkruysse et al., 2008; Tarassenko
42 et al., 2014; Benezeth et al., 2018; Wang et al., 2016, 2015; Pilz et al., 2018; de Haan and van Leest,
43 2014). More recently, alongside the traditional methods listed above, rPPG approaches based on deep
44 learning (DL) have burst into this research field (Chen and McDuff, 2018; Niu et al., 2019; Yu et al., 2020;
45 Liu et al., 2020, 2021; Gideon and Stent, 2021; Yu et al., 2021).

46 The blossoming of the field and the variety of the proposed solutions raise the issue, for both

47 researchers and practitioners, of a fair comparison among proposed techniques while engaging in the
48 rapid prototyping and the systematic testing of novel methods. Under such circumstances, several reviews
49 and surveys concerning rPPG (McDuff et al., 2015; Rouast et al., 2018; Heusch et al., 2017a; Unakafov,
50 2018; Wang et al., 2016; McDuff and Blackford, 2019; Cheng et al., 2021; McDuff, 2021; Ni et al., 2021)
51 have conducted empirical comparisons, albeit suffering under several aspects, as discussed in Sec. 1.1.

52   To promote the development of new methods and their experimental analysis, in Boccignone et al.
53 (2020a) we proposed `pyVHR`, a preliminary version of a framework supporting the main steps of the
54 traditional rPPG pulse rate recovery, together with a sound statistical assessment of methods' performance.
55 Yet, that proposal exhibited some limits, both in terms of code organization, usability, and scalability, and
56 since it was suitable for traditional approaches only. Here we present a new version of `pyVHR` [1], with a
57 totally re-engineered code, which introduces several novelties.

58   First of all, we provide a dichotomous view of remote heart rate monitoring, leading to two distinct
59 classes of approaches: traditional methods (Sec. 3) and DL-based methods (Sec. 4). Moreover, concerning
60 the former, a further distinction is setup, concerning the Region Of Interest (ROI) taken into account, thus
61 providing both holistic and patch-based methods. The former takes into account the whole skin region,
62 extracted from the face captured in subsequent frames. Undoubtedly, it is the simplest approach, giving
63 satisfying results when applied on video acquired in controlled contexts. However, in more complex
64 settings the illumination conditions are frequently unstable, giving rise to either high variability of skin
65 tone or shading effects. In these cases the holistic approach is prone to biases altering subsequent analyses.
66 Differently, the patch-based approach employs and tracks an ensemble of patches sampling the whole
67 face. The rationale behind this choice is twofold. On the one hand, the face regions affected by either
68 shadows or bad lighting conditions can be discarded, thus avoiding uncorrelated measurements with the
69 HR ground-truth. On the other hand, the amount of observations available allows for making the final
70 HR estimate more robust, even through simple statistics (e.g., medians), while controlling the confidence
71 levels.

72   Second, the framework is agile, covers each stage of the pipeline that instantiates it, and it is easily
73 extensible. Indeed, one can freely embed new methods, datasets or tools for the intermediate steps (see
74 Sec. 6) such as for instance: face detection and extraction, pre- and post-filtering of RGBs traces or BVPs
75 signals, spectral analysis techniques, statistical methods.

76   `pyVHR` can be easily employed for many diverse applications such as anti-spoofing, aliveness
77 detection, affective computing, biometrics. For instance, in Sec. 7 a case study on the adoption of rPPG
78 technology for a Deepfake detection task is presented.

79   Finally, computations can be achieved in real-time thanks to the NVIDIA GPU (Graphics Processing
80 Units) accelerated code and the use of optimized Python primitives.

1.1 Related Works

82 In the last decade the rPPG domain has witnessed a flourish of investigations (McDuff et al., 2015; Rouast
83 et al., 2018; Heusch et al., 2017a; Unakafov, 2018; Wang et al., 2016; McDuff and Blackford, 2019;
84 Cheng et al., 2021; McDuff, 2021; Ni et al., 2021). Yet, the problem of a fair and reproducible evaluation
85 has been in general overlooked. It is undeniable that theoretical evaluations are almost infeasible, given
86 the complex operations or transformations each algorithm performs. Nevertheless, empirical comparisons
87 could be very informative if conducted in the light of some methodological criteria (Boccignone et al.,
88 2020a). In brief: pre/post processing standardization; reproducible evaluation; multiple dataset testing;
89 rigorous statistical assessment.

90   To the best of our knowledge, a framework respecting all these criteria was missing until the introduc-
91 tion of the early version of `pyVHR` Boccignone et al. (2020a).

92   In Heusch et al. (2017a) a Python collection of rPPG algorithms is presented, without claiming to be
93 complete in the method assessment.

94   Interestingly, in Unakafov (2018), the authors highlight the dependency of the pulse rate estimation
95 on five main steps: ROI-selection, pre-processing, rPPG method, post-processing, pulse rate estimation.
96 They present a theoretical framework to assess different pipelines in order to find out which combination
97 provides the most precise PPG estimation; results are reported on the DEAP dataset (Koelstra et al., 2011).
98 Unfortunately, no code has been made available.

---

[1]Freely available on GitHub: github.com/phuselab/pyVHR.

In Pilz (2019) a MATLAB toolbox is presented, implementing two newly proposed methods, namely Local Group Invariance (LGI) (Pilz et al., 2018) and Riemannian-PPGI (SPH) (Pilz, 2019), and comparing them to the GREEN channel expectation (Verkruysse et al., 2008) baseline, and two state-of-the-art methods, i.e. Spatial Subspace Rotation (SSR) (Wang et al., 2015), and Projection Orthogonal to Skin (POS) (Wang et al., 2016).

In McDuff and Blackford (2019) authors propose `iPhys`, a MATLAB toolbox implementing several methods, such as Green Channel, POS, CHROM (De Haan and Jeanne, 2013), ICA (Poh et al., 2010), and BCG (Balakrishnan et al., 2013). The toolbox is presented as a bare collection of method implementations, without aiming at setting up a rigorous comparison framework on one or more datasets. It is worth noticing that all these frameworks are suitable for traditional methods only. Table 1 summarizes at a glance the main differences between pyVHR and the already proposed frameworks.

	Lang.	Modular	Deep-Ready	Multi-Data	Stat. Assessment
pyVHR	Python	✓	✓	✓	✓
McDuff and Blackford (2019)	MATLAB	✗	✗	✗	✗
Heusch et al. (2017a)	Python	✗	✗	✗	✗
Pilz (2019)	MATLAB	✗	✗	✓	✗

Table 1. A comparison of the freely available rPPG frameworks. Check signs mark conditions fulfilled; crosses, those neglected.

2 INSTALLATION

The quickest way to get started with pyVHR is to install the miniconda distribution, a lightweight minimal installation of Anaconda Python.

Once installed, create a new conda environment, automatically fetching all the dependencies based on the adopted architecture - with or without GPU -, by one of the following commands:

```
$ conda env create --file https://github.com/phuselab/pyVHR/blob/pyVHR_CPU/
    pyVHR_CPU_env.yml
```

for CPU-only architecture, or

```
$ conda env create --file https://github.com/phuselab/pyVHR/blob/main/pyVHR_env.yml
```

for a CPU architecture with GPU support. The latest stable release build of pyVHR can be installed, inside the newly created conda environment, with:

```
$ pip install pyvhr-cpu
```

for CPU-only, or

```
pip install pyvhr
```

for CPU with GPU support.

The source code for pyVHR can be found on GitHub at https://github.com/phuselab/pyVHR and it is distributed under the GPL-3.0 License. On GitHub, the community can report issues, questions as well as contribute with code to the project. The documentation of the pyVHR framework is available at https://phuselab.github.io/pyVHR/.

3 PYVHR PIPELINE FOR TRADITIONAL METHODS

In this section, we introduce the pyVHR modules to be referred by traditional rPPG methods. They are built on top of both APIs developed for the purpose, and open-source libraries. This pipeline follows a software design strategy that assemble sequential modules or stages, with the output of a stage serving as input to one or more subsequent stages. This responds to the need for the framework to be flexible and extensible in order to be more maintainable and improvable over time with innovative or alternative techniques.

3.1 The Pipeline Stages

The `Pipeline()` class implements the sequence of stages or steps that are usually required by the vast majority of rPPG methods proposed in the literature, in order to estimate the BPM of a subject, given a video displaying his/her face. Eventually, going through all these steps in `pyVHR` is as simple as writing a couple of lines of Python code:

```python
from pyVHR.analysis.pipeline import Pipeline

pipe = Pipeline()
time, BPM, uncertainty = pipe.run_on_video('/path/to/vid.avi')
```

Calling the `run_on_video()` method of the `Pipeline()` class starts the analysis of the video provided as argument and produces as output the `time` step of the estimated `BPM` and related `uncertainty` estimate. Figure 1 depicts the predicted BPM on `Subject1` of the UBFC[2] dataset (Bobbia et al., 2019) (blue trajectory). For comparison, the ground truth BPM trajectory (as recorded from a PPG sensor) is reported in red.

Figure 1. Prediction example. Predictions on the `Subject1` of the `UBFC` Dataset.

On the one hand the above-mentioned example witnesses the ease of use of the package by hiding the whole pipeline behind a single function call. On the other hand it may be considered too constraining as hinders the user from exploiting its full flexibility. Indeed, the `run_on_video()` method can be thought of as a black box delivering the desired result with the least amount of effort, relying on default parameter setting.

Nevertheless, some users may be interested in playing along with all the different modules composing the `pyVHR` pipeline and the related parameters. The following sections aim at describing in detail each of such elements. These are shown in Figure 2(b) and can be recapped as follows:

1. *Skin extraction:* The goal of this first step is to perform a face skin segmentation in order to extract PPG-related areas; the latter are subsequently collected in either a single patch (holistic approach) or a bunch of "sparse" patches covering the whole face (patch-wise approach).

2. *RGB signal processing:* The patches, either one or more, are coherently tracked and are used to compute the average colour intensities along overlapping windows, thus providing multiple time-varying RGB signals for each temporal window.

3. *Pre-filtering:* Optionally, the raw RGB traces are pre-processed via canonical filtering, normalization or de-trending; the outcome signals provide the inputs to any subsequent rPPG method.

4. *BVP extraction:* The rPPG method(s) at hand is applied to the time-windowed signals, thus producing a collection of heart rate pulse signals (BVP estimates), one for each patch.

---

[2]Available at: https://sites.google.com/view/ybenezeth/ubfcrppg.

168   5. *Post-filtering:* The pulse signals are optionally passed through a narrow-band filter in order to
169   remove unwanted out-of-band frequency components.

170   6. *BPM estimation:* A BPM estimate is eventually obtained through simple statistics relying on the
171   apical points of the BVP power spectral densities.
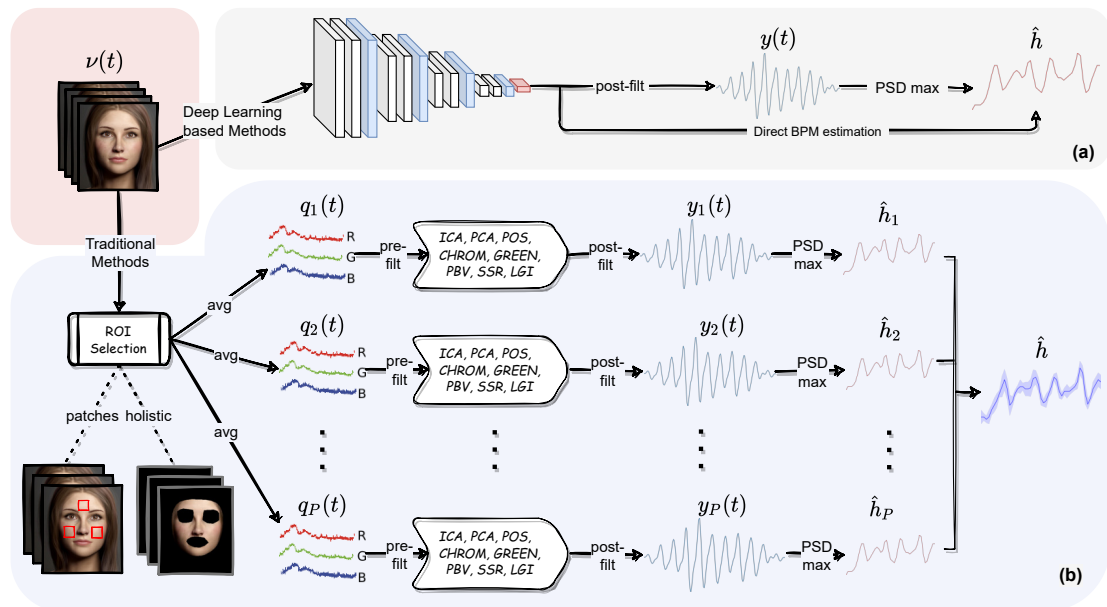
Figure 2. The pyVHR Pipeline at a glance. (a) The multi-stage pipeline of the `pyVHR` framework for
BPM estimate through PSD analysis exploiting end-to-end DL-based methods. **(b)** The multi-stage
pipeline for traditional approaches that goes through: windowing and patch collection, RGB trace
computation, pre-filtering, the application of an rPPG algorithm estimating a BVP signal, post-filtering
and BPM estimate through PSD analysis.

3.2 Skin Extraction

173 The skin extraction step implemented in `pyVHR` consists in the segmentation of the face region of the
174 subject. Typically, the regions corresponding to the eyes and mouth are discarded from the analysis. This
175 can be accomplished by `pyVHR` in two different ways, denoted as:

176   1. the *Convex-hull* extractor,

177   2. the *Face parsing* extractor.

178   The *Convex-hull* extractor considers the skin region as the convex-hull of a set of the 468 facial fiducial
179 points delivered by the MediaPipe face mesh (Google, 2021). The latter provides reliable face/landmark
180 detection and tracking in real-time. From the convex-hulls including the whole face, `pyVHR` subtracts
181 those computed from the landmarks associated to the eyes and mouth. The resulting mask is employed
182 to isolate the pixels that are generally associated to the skin. An example is shown in the left image of
183 Figure 3 on a subject of the LG-PPGI dataset (Pilz et al., 2018)[3].
184   Alternatively, the *Face parsing* extractor computes a semantic segmentation of the subject's face. It
185 produces pixel-wise label maps for different semantic components (e.g., hair, mouth, eyes, nose, etc...),
186 thus allowing to retain only those related to the skin regions. Face semantic segmentation is carried over
187 with BiSeNet (Yu et al., 2018), which supports real-time inference speed. One example is shown in the
188 right image of Figure 3.
189   Both extraction methods are handled in `pyVHR` by the `SignalProcessing()` class. The following
190 lines of code set-up the extractor with the desired skin extraction procedure:

---

[3]Available for download at https://github.com/partofthestars/LGI-PPGI-DB

(A) Convex Hull   (B) BiSeNet

Figure 3. Comparison of the two implemented skin extraction methods. Output of the Convex-hull approach (left) and face parsing by BiSeNet (right) on a subject of the LGI-PPGI dataset (Pilz et al., 2018).

```
191 1  from pyVHR.extraction.sig_processing import SignalProcessing
192 2
193 3  sig_processing = SignalProcessing()
194 4  if skin_method == 'convexhull':
195 5      sig_processing.set_skin_extractor(SkinExtractionConvexHull(target_device))
196 6  elif skin_method == 'faceparsing':
197 7      sig_processing.set_skin_extractor(SkinExtractionFaceParsing(target_device))
```

3.2.1 Holistic Approach

The skin extraction method paves the way to the RGB trace computation which is accomplished in a channel-wise fashion by averaging the facial skin colour intensities. This is referred to as the holistic approach, and within the `pyVHR` framework it can be instantiated as follows:

```
202 1  sig = sig_processing.extract_holistic(videoFileName)
```

3.2.2 Patch-based Approach

In contrast to the holistic approach, the patch-based one takes into account a bunch of localized regions of interest, thus extracting as many RGB traces as patches. Clearly, imaging photoplethysmography in unconstrained settings is sensitive to subjects changing pose, moving their head or talking. This calls for a mechanism for robust detection and tracking of such regions.

To such end, `pyVHR` again relies on the MediaPipe Face Mesh, which establishes a metric 3D space to infer the face landmark screen positions by a lightweight method to drive a robust and performant tracking. The analysis runs on CPU and has a minimal speed or memory footprint on top of the inference model.

The user can easily select up to 468 patches centered on a subset of landmarks and define them as the set of informative regions on which the subsequent steps of the pipeline are evaluated. An example of landmark extraction and tracking is shown in Figure 4. Note that eventually, a patch may disappear due to subject's movements, hence delivering only partial or none contribution.

Figure 4. Landmarks automatically tracked by MediaPipe and correspondent patch tracking on a subject of the LGI-PPGI dataset (Pilz et al., 2018).

215

216   It is worth noting how the user is allowed to arbitrarily compose its own set of patches by exploiting
217 `pyVHR` utility functions. In the example below, three patches have been selected corresponding to the
218 forehead, left and right cheek areas. Usually, several patches are chosen in order to better control the high
219 variability in the results and to achieve high level of confidence, while making smaller the margin of error.
220   As for the holistic approach, video loading and patch extraction are handled by few APIs available in
221 the `SignalProcessing()` class, as shown in the following script.

```
222 from pyVHR.extraction.utils import MagicLandmarks
223
224 ldmks_list = [MagicLandmarks.cheek_left_top[16], MagicLandmarks.cheek_right_top[14],
225     MagicLandmarks.forehead_center[1]]
226 sig_processing.set_landmarks(ldmks_list)
227 # set squares patches side dimension
228 sig_processing.set_square_patches_side(28.0)
229 #Extract square patches and compute the RGB trajectories as the channel-wise mean
230 sig = sig_processing.extract_patches(videoFileName, 'squares', 'mean')
```

3.2.3 RGB Signal Computation

231
232 In this step, the skin regions detected and tracked on the subject's face are split in successive overlapping
233 time windows. Next, the RGB traces associated to each region are computed by averaging their colour
234 intensities.

More formally, let us consider an RGB video $v \in \mathbb{R}^{w \times h \times 3 \times T}$ of T frames containing a face, split on P (possibly overlapped) patches. Once the i-th patch has been selected, an RGB signal $q_i(t)$ is computed. Denote $\{p_i^j(t)\}_{j=1}^{N_i}$ the set of N_i pixels belonging to the i-th patch at time t, where $p_i^j(t) \in [0, 255]^3$. Then, $q_i(t)$ is recovered by averaging on pixel colour intensities, i.e.,

$$q_i(t) = \frac{1}{N_i} \sum_{j=1}^{N_i} p_i^j(t), \qquad i = 1, \dots, P.$$

In the time-splitting process, fixed an integer $\tau > 0$, $q_i(t)$ is sliced into K overlapping windows of $M = W_s F_s$ frames, thus obtaining

$$q_i^k(t) = q_i(t)w(t - k\tau F_s), \qquad k = 0, \dots, K-1.$$

where F_s represents the video frame rate, W_s the window length in seconds, while w is the rectangular window defined as:

$$w(t) = \begin{cases} 1, & 0 \le t < M \\ 0, & \text{otherwise.} \end{cases} \tag{1}$$

235   In order for the signal segments to actually overlap, the overlap inequality $\tau < W_s$ must be verified.
236   Figure 5 shows how the above described patch-based split and tracking procedure is put in place.
237   In `pyVHR` , the extraction of the windowed RGB signals is computed by the following code snippet.

```
238 from pyVHR.extraction.utils import sig_windowing, get_fps
239
240 Ws = 6 #window lenght in seconds
241 overlap = 1 #window overlap in seconds
242 fps = get_fps(videoFileName)
243
244 windowed_sig, timesES = sig_windowing(sig, Ws, overlap, fps)
```

245   Notably, beside being able to switch between convex-hull and face parsing, the user can easily change
246 the main process parameters such as the window length and the amount of frame overlapping.

3.2.4 Methods for BVP estimation

247
248 Given that the framework can rely on holistic-wise and patch-wise processing, `pyVHR` estimates the BVP
249 signal either from a single trace or leveraging on multiple traces. In both cases it employs a wide range of
250 state of the art rPPG methods.
251   In particular, the windowed RGB traces $q_i^k(t)$ ($i = 1, \dots, P$, with $P = 1$ in the holistic case) of length K
252 are given in input to the rPPG method at hand, which outputs the signals $y_i^k(t)$ representing the estimated
253 BVP associated to the i-th patch in the k-th time window.

Figure 5. Patch tracking within a frame temporal window on a subject of the LGI-PPGI dataset (Pilz et al., 2018).

254   The many methods that have been proposed in the recent literature mostly differ in the way of
255 combining such RGB signals into a pulse-signal. The pool of methods provided by `pyVHR`, together
256 with a description of the main concepts grounding them, is provided in Table 2. A review of the
257 principles/assumptions behind each of the implemented algorithms is out of the scope of the present work.
258 The interested reader might refer to Wang et al. (2016); McDuff et al. (2015); Rouast et al. (2018).

259   Currently, the package implements the following methods for the estimation of the pulse signal from
260 the RGB traces: GREEN (Verkruysse et al., 2008), CHROM (De Haan and Jeanne, 2013), ICA (Poh et al.,
261 2010), LGI (Pilz et al., 2018), PBV (de Haan and van Leest, 2014), PCA (Lewandowska et al., 2011),
262 POS (Wang et al., 2016), SSR (Wang et al., 2015). However, the user may define any custom method for
263 estimating BVP by extending the `pyVHR.BVP.methods` module.

264   The BVP signal can be estimated in `pyVHR` as follows:

265
```
bvp = RGB_sig_to_BVP(windowed_sig, fps, method=cpu_POS)
```

266   Figure 6 depicts the BVP signals estimated by four different rPPG methods implemented in `pyVHR`
267 (POS, GREEN, CHROM, PCA), on the same time window using the holistic patch.

3.2.5 Pre and Post-Filtering

269 `pyVHR` offers simple APIs to apply filters on either the RGB traces $q_i(t)$ (pre-filtering) or the estimated
270 pulse signal $y_i(t)$ (post-filtering). A set of ready to use filters are implemented, namely:

271 • *Band Pass (BP) filter*: filters the input signal using a bandpass N-th order Butterworth filter with a
272   given passband frequency range.

273 • *Detrending*: subtracts offsets or linear trends from time-domain input data.

274 • *Zero-Mean*: Removes the DC component from a given signal.

275   However, the user can adopt any custom filter complying with the function signature defined in
276 `pyVHR.BVP.filters`. The following provides an example of how to detrend an RGB trace $q_i(t)$:

277
```
filtered_sig = apply_filter(sig, detrend)
```

278   Additionally, a Band-Pass filter can be applied on the estimated BVP signals $y_i(t)$ in order to rule out
279 the frequencies that leave outside the feasible range of typical heart rates (which is usually between 40Hz
280 and 200Hz):

281
```
filtered_bvp = apply_filter(bvp, BPfilter,
                params={'order':6,'minHz':0.65,'maxHz':4.0,'fps':fps})
```

Table 2. Traditional rPPG algorithms implemented in `pyVHR` .

Method	Description
GREEN (Verkruysse et al., 2008)	The Green (G) temporal trace is directly considered as an estimate of the BVP signal. Usually adopted as a *baseline* method.
ICA (Poh et al., 2010)	Independent Component Analysis (ICA) is employed to extract the pulse signal via Blind Source Separation of temporal RGB mixtures.
PCA (Lewandowska et al., 2011)	Principal Component Analysis (PCA) of temporal RGB traces is employed to estimate the BVP signal.
CHROM (De Haan and Jeanne, 2013)	A Chrominance-based method for the BVP signal estimation.
PBV (de Haan and van Leest, 2014)	Computes the signature of blood volume pulse changes to distinguish the pulse-induced color changes from motion noise in RGB temporal traces.
SSR (S2R) (Wang et al., 2015)	Spatial Subspace Rotation (SSR); estimates a spatial subspace of skin-pixels and measures its temporal rotation for extracting pulse signal.
POS (Wang et al., 2016)	Plane Orthogonal to the Skin (POS). Pulse signal extraction is performed via a projection plane orthogonal to the skin tone.
LGI (Pilz et al., 2018)	Local Group Invariance (LGI). Computes a feature representation which is invariant to action and motion based on differentiable local transformations.

(A) POS   (B) GREEN

(C) CHROM   (D) PCA

Figure 6. Predicted BVP signals. An example of estimated BVP signals on the same time window by four different methods.

3.2.6 From BVP to BPM

Given the estimated BVP signal, the beats per minute (BPM) associated to a given time window can be easily recovered via analysis of its frequency domain representation. In particular, `pyVHR` estimates the Power Spectral Density (PSD) of the windowed pulse signal $y_i^k(t)$ via discrete time Fourier transform (DFT) using the Welch's method. The latter employs both averaging and smoothing to analyze the underlying random process.

Given a sequence $y_i^k(t)$, call $S_i^k(\nu)$ its power spectra (periodogram) estimated via the Welch's method. The BPM is recovered by selecting the normalized frequency associated to the peak of the periodogram:

$$\hat{\nu}_i^k = \arg\max_{\nu \in \Omega} \left\{ S_i^k(\nu) \right\},$$

corresponding to the PSD maxima as computed by Welch's method on the range $\Omega = [39, 240]$ of feasible BPMs.

The instantaneous BPM associated to the k-th time window ($k \in 1, \ldots, K$) for the i-th patch ($i \in 1, \ldots, P$), is recovered by converting the normalized peak frequency $\hat{\nu}_i^k$ into an actual frequency,

$$\hat{h}_i^k = \hat{\nu}_i^k \frac{F_s}{L},$$

where F_s is the video frame rate and L is the DFT size. Figure 7 shows the Welch's estimates for the BVP signals of Figure 6. The peak in the spectrum represents the instantaneous Heart Rate ($\hat{h}_i^k$).
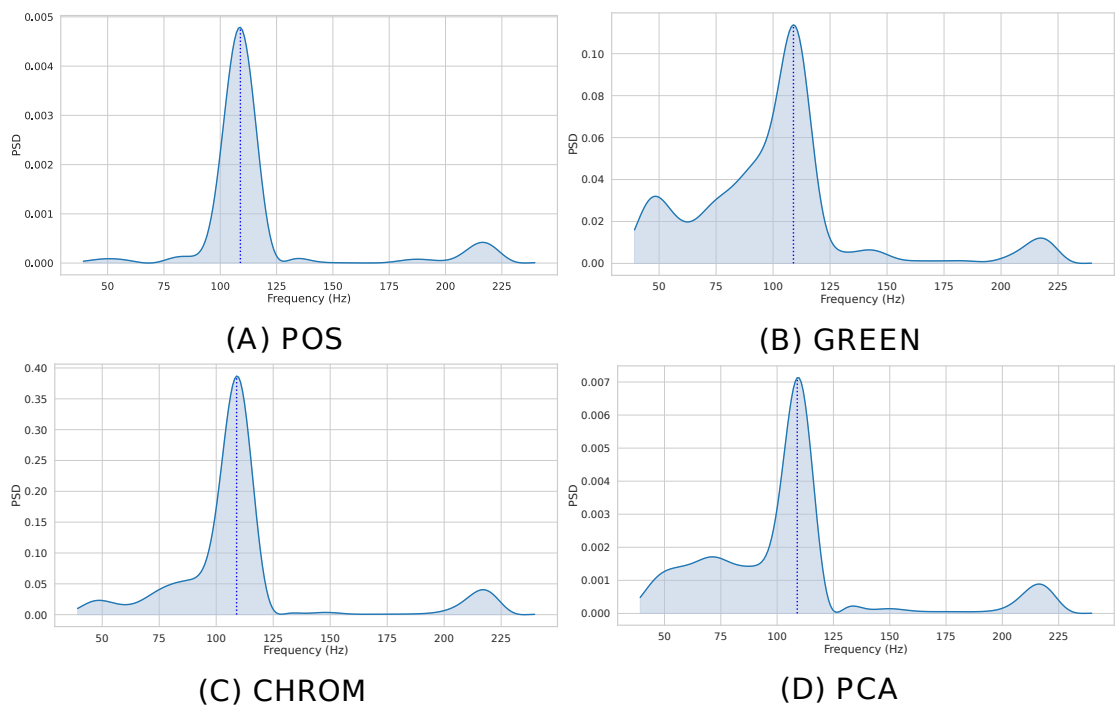
Figure 7. Estimated PSD. Estimated Power Spectral Densities (PSD) for the BVP signals plotted in Figure 6. The BPM estimate, given by the maxima of the PSD, is represented by the blue dashed line.

When multiple patches have been selected ($P > 1$), the predicted BPM for the k-th time window can be obtained resorting to simple statistical measures. Specifically, `pyVHR` computes the median BPM value of the predictions coming from the P patches.

Formally, call H^k the ordered list of P BPM predictions coming from each patch in the k-th time window; then:

$$\hat{h}^k = \text{median}(H^k) = \begin{cases} H^k\left[\frac{P-1}{2}\right] & \text{if P is odd} \\ \frac{\left(H^k\left[\frac{P}{2}-1\right]+H^k\left[\frac{P}{2}\right]\right)}{2} & \text{if P is even.} \end{cases} \tag{2}$$

298   Note that if the number of patches $P = 1$ (i.e. a single patch has been selected or the holistic approach
299   has been chosen), then:

$$\hat{h}^k = H^k[0]. \tag{3}$$

300   Moreover, when multiple patches have been selected, a measure of variability of the predictions can
301   be computed in order to quantify the uncertainty of the estimation. In particular, `pyVHR` computes the
302   Median Absolute Deviation (*MAD*) as a robust measure of statistical dispersion. The *MAD* is defined as:

$$MAD^k = \text{median}(|H^k - \hat{h}^k|). \tag{4}$$

303   Clearly, the *MAD* drops to 0 when $P = 1$.
304   Figure 8 depicts the distribution of predicted BPM in a given time window, when $P = 100$ patches are
305   employed. The results from different methods are shown for comparison. Note how the median is able to
306   deliver precise predictions, while the MAD represents a robust measure of uncertainty.
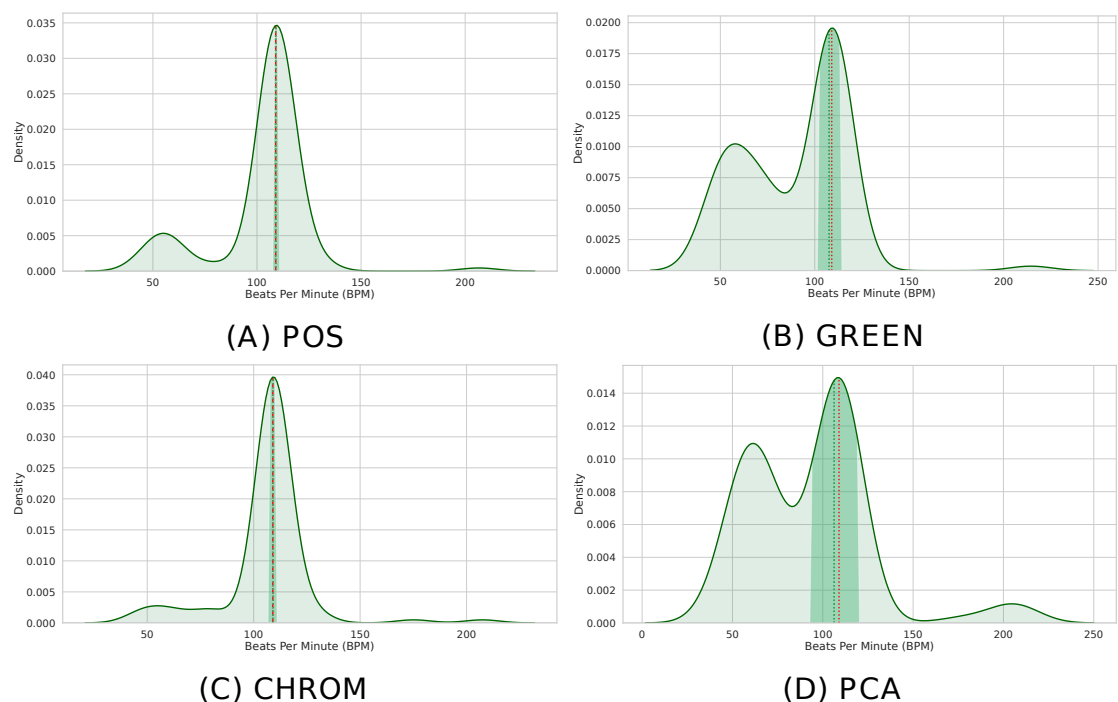
Figure 8. Distribution of BPM predictions by four methods on P patches. Kernel Density Estimates
(KDEs) of the predicted BPMs in a time window from $P = 100$ patches. The ultimate BPM prediction is
given by the median (green dashed line). The uncertainty estimate deliverd by the Median Absolute
Deviation (MAD) is shown in darker green. the red dashed line represents the actual BPM.

307   Computing the BPM from the BVP signal(s) can be easily accomplished in `pyVHR` as follows:

```
from pyVHR.BPM.BPM import BVP_to_BPM, multi_est_BPM_median

bpmES = BVP_to_BPM(bvp, fps)
# median BPM from multiple estimators BPM
bpm, uncertainty = multi_est_BPM_median(bpmES)
```

313   The result along with the ground-truth are shown in Figure 9.

3.3 Efficient computation and GPU acceleration

315   Most of the steps composing the pipeline described above are well suited for parallel computation. For
316   instance, the linear algebra operations involved in the pulse signal recovery from the RGB signal or, more

Figure 9. Comparison of predicted vs Ground truth BPMs using the patch-wise approach.
Predicted BPM (blue) for the `Subject1` of the `UBFC` Dataset. The uncertainty is plotted in shaded blue, while the ground-truth is represented by the red line.

generally, the signal processing steps (e.g. filtering, spectral estimation, etc..), not to mention the skin segmentation procedures from high resolution videos.

To such end, `pyVHR` exploits the massive parallelism of Graphical Processing Units (GPUs). It is worth mentioning that GPUs are not strictly required to run `pyVHR` code; nevertheless, in some cases, GPU accelerated code allows to run the pipeline in real-time.

Figure 10 shows the average per-frame time requirement for getting through the whole pipeline when using the `POS` method. It is worth noticing that, when using the Holistic approach (or equivalently one single patch), a video frame can be processed in less than 0.025 seconds, regardless of the adopted architecture (either CPU or GPU). This means that the whole pipeline can be safely run in real-time for videos at 30 frames per second (the 30 fps time limit is represented by the dashed green line).

Figure 10. Per-frame time requirements. Average time requirements to process one frame by the Holistic and Patches approaches when using CPU *vs.* GPU accelerated implementations. The green dashed line represents the real-time limit at 30 frames per second (fps).

Obviously, when multiple patches are employed (in the example of Figure 10, $P = 100$ patches are used), the average time required by CPUs to process a single frame rises up to about 0.12 seconds. Notably, the adoption of GPU accelerated code allows to run the whole pipeline in real-time, even when using a huge number of patches. Indeed, the ratio to CPU time and GPU time, i.e., the speedup defined

as $time_{\text{seq}}/time_{\text{parall}}$, is about 5. Remarkably, similar gain in performances are observed if adopting any other rPPG method.

The result shown in Figure 10 refers to the following hardware configuration: Intel Xeon Silver 4214R 2.40GHz (CPU), NVIDIA Tesla V100S PCIe 32GB (GPU). Similar results were obtained relying on a non-server configuration: Intel Core i7-8700K 4.70 GHz (CPU), NVIDIA GeForce GTX 960 2GB (GPU). The maximum RAM usage for 1 minute HD video analysis is 2.5 GB (average is 2 GB); the maximum GPU memory usage for 1 minute HD video analysis is 1.8 GB (average is 1.4 GB).

In the following it is shown how to enable CUDA GPU acceleration on different steps in the Pipeline:

- Skin extraction: Convex Hull and Face Parsing. The user can easily choose to run this step with CPU or GPU:

```
1    target_device = 'GPU' # or 'CPU'
2    sig_processing = SignalProcessing()
3    sig_processing.set_skin_extractor(
4            SkinExtractionConvexHull(target_device))
5    sig_processing.set_skin_extractor(
6            SkinExtractionFaceParsing(target_device))
7
```

- rPPG Methods: the package contains different version of the same method. For example the CHROM method is implemented for both CPU and GPU.

```
1    bvp = RGB_sig_to_BVP(sig, fps, device_type='cuda', method=cupy_CHROM)
2    bvp = RGB_sig_to_BVP(sig, fps, device_type='cpu', method=cpu_LGI)
3
```

- BPM Estimation:

```
1      bpmES = BVP_to_BPM_cuda(bvp, fps)    #GPU
2      bpmES = BVP_to_BPM(bvp, fps)         #CPU
3
```

3.4 GUI for Online Processing

Besides being used as a Python library, `pyVHR` makes available a Graphical User Interface (GUI). It provides access to most of the available functionalities, while showing the BPMs estimation process in real-time. It is straightforward to use and it allows for setting up the pipeline parameters and the operating mode, by choosing either a webcam or a video file.

To start the GUI, one can run the command:

```
$ Python pyVHR/realtime/GUI.py
```

Figure 11 shows a screenshot of the GUI during the online analysis of a video. On the top right are presented the video file name, the video FPS, resolution, and a radio button list to select the type of frame displayed. The original or segmented face can be visualized either selecting the *Original Video* or the *Skin* option, while the *Patches* radio button enables the visualization of the patches (in red). The *Stop* button ends the analysis, and results can be saved on disk by pushing the *Save BPMs* button.

4 `PYVHR` PIPELINE FOR DEEP-LEARNING METHODS

Recent literature in computer vision has given wide prominence to end-to-end deep neural models and their ability to outperform traditional methods requiring hand-crafted feature design. In this context, learning frameworks for recovering physiological signals were also born (Chen and McDuff, 2018; Niu et al., 2019; Yu et al., 2020, 2021; Gideon and Stent, 2021; Liu et al., 2021; Nowara et al., 2020).

The end-to-end nature of the DL based approaches is reflected by a much simpler pipeline; indeed, these methods typically require as input raw video frames that are processed by the DL architecture at hand and produce either a BVP signal or the estimated heart rate, directly. Figure 2(a) depicts at a glance the flow of stages involved in the estimation of heart rate using DL based approaches. Clearly, this gain in simplicity comes at the cost of having to train the model on huge amounts of data, not to mention the issues related to the assessment of the model's generalization abilities.
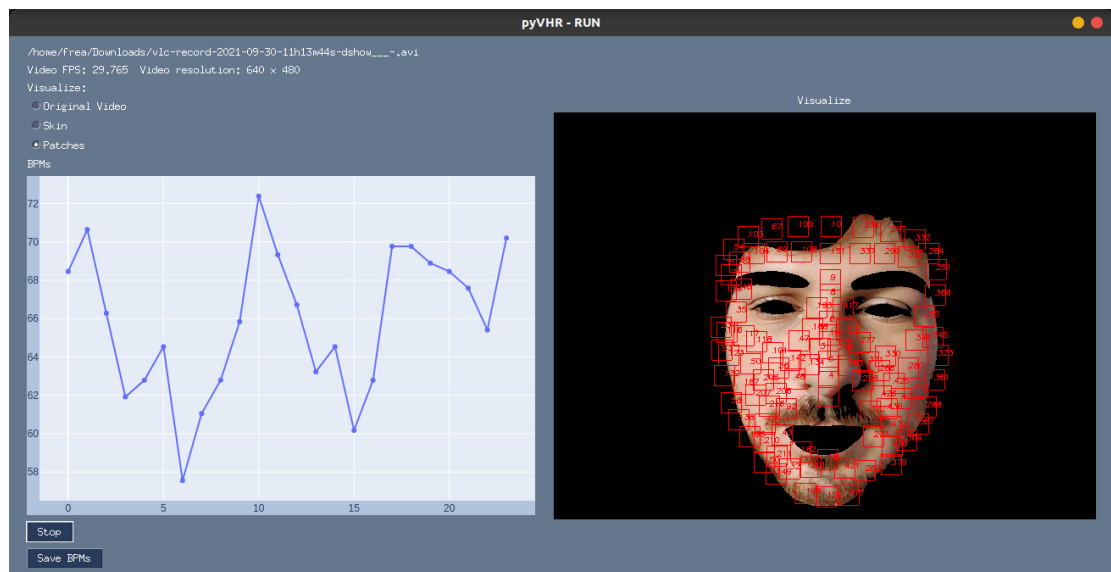
Figure 11. The Graphical User Interface. A screenshot of the Graphical User Interface (GUI) for online video analysis. The plot on the left shows the predicted BPMs, while on the right it is shown the processed video frames (captured with a webcam) with an example of the segmented skin and the tracked patches.

In the last few years the literature has witnessed a flourish of DL-based approaches (for two recent reviews see Cheng et al. (2021) and Ni et al. (2021)). Nonetheless, despite the claimed effectiveness and superior performances, few solutions have been made publicly available (both in terms of code and learned model weights). This raises issues related to proper reproducibility of the results and the method assessment. For instance, a recent efficient neural architecture called MTTS-CAN has been proposed in Liu et al. (2020) being a valuable contribution since the pre-trained model and code are released. It essentially leverages a tensor-shift module and 2D-convolutional operations to perform efficient spatial temporal modeling in order to enable real-time cardiovascular and respiratory measurements. MTTS-CAN can be framed as an end-to-end model since it does not need any pre-processing step before data is fed into the network, except performing trivial image normalizations. MTTS-CAN is included in the `pyVHR` framework, and below it is shown how practical is to extend the framework with similar DL-based approaches provided that the pre-trained model is available.

As for the pipeline for traditional methods shown in previous section, `pyVHR` also defines a sequence of stages that allows to recover the time varying heart rate from a sequence of images displaying a face. These are detailed in the following.

4.1 The Stages for End-to-end Methods

Given a video displaying a subject face, the `DeepPipeline()` class performs the necessary steps for the rPPG estimate using a chosen end-to-end DL method. Specifically, the pipeline includes the handling of input videos, the estimation from the sequence of raw frames and, eventually, the pre/post-processing steps.

The following code snippet carries out the above procedure with few statements:

```
from pyVHR.analysis.pipeline import DeepPipeline

pipe = DeepPipeline()
time, BPM = pipe.run_on_video('/path/to/vid.avi', method='MTTS_CAN')
```

Figure 2(a) summarizes the steps involved in a `run_on_video()` call on a given input video. As in the pipeline using traditional methods (see Section 3), after a predetermined chain of analysis steps it produces as output the estimated `BPM` and related timestamps (`time`).

For instance, consider the `MTTS-CAN` model currently embedded into the `DeepPipeline()` class; it estimates the rPPG pulse signal from which the BPM computation can be carried out by following

410 the very same procedure outlined in Section 3.2.6, namely time windowing and spectral estimation.
411 Eventually, some optional pre/post filtering operations (Section 3.2.5) can be performed.
412   The following few lines of Python code allow to carry out the above steps explicitly:

```python
from pyVHR.extraction.sig_processing import SignalProcessing
from pyVHR.extraction.utils import get_fps
from pyVHR.BPM import BVP_to_BPM
from pyVHR.utils.errors import BVP_windowing

sp = SignalProcessing()
frames = sp.extract_raw('/path/to/videoFileName')
fps = get_fps('/path/to/videoFileName')

bvps_pred = MTTS_CAN_deep(frames, fps)

winsize = 6 #6 seconds long time window
bvp_win, timesES = BVP_windowing(bvp_pred, winsize, fps, stride=1)
bpm = BVP_to_BPM(bvp_win, fps)
```

427   In order to embed a new DL-method, the code above should be simply modified substituting the
428 function MTTS_CAN_deep with a new one implementing the method at hand, while respecting the same
429 signature (cfr. Sec. 6).

5 ASSESSMENT OF RPPG METHODS

431 Does a given rPPG algorithm outperforms the existing ones? To what extent? Is the difference in
432 performance significantly large? Does a particular post-filtering algorithm cause an increase/drop of
433 performance?
434   Answering all such questions, calls for a rigorous statistical assessment of rPPG methods. As
435 a matter of fact, although the field has recently experienced a substantial gain in interest from the
436 scientific community, it is still missing a sound and reproducible assessment methodology allowing to
437 gain meaningful insights and delivering best practices.
438   By and large, novel algorithms proposed in the literature are benchmarked on non-publicly available
439 datasets, thus hindering proper reproducibility of results. Moreover, in many cases, the reported results
440 are obtained with different pipelines; this makes it difficult to precisely identify the actual effect of the
441 proposed method on the final performance measurement.
442   Besides that, the performance assessment mostly relies on basic and common-sense techniques, such
443 as roughly rank new methods with respect to the state-of-the-art. These crude methodologies often make
444 the assessment unfair and statistically unsound. Conversely, a good research practice should not limit to
445 barely report performance numbers, but rather aiming at principled and carefully designed analyses. This
446 is in accordance with the growing quest for statistical procedures in performance assessment in many
447 different fields, including machine learning and computer vision (Demšar, 2006; Benavoli et al., 2017;
448 Torralba and Efros, 2011; Graczyk et al., 2010; Eisinga et al., 2017).
449   In the vein of its forerunner (Boccignone et al., 2020a), pyVHR deals with all such problems by means
450 of its statistical assessment module. The design principles can be recapped as follows:

- *Standardized pipeline:* When setting up an experiment to evaluate a new rPPG algorithm, the whole
  pipeline (except the algorithm) should be held fixed.

- *Reproducible evaluation:* The evaluation protocol should be reproducible. This entails adopting
  publicly available datasets and code.

- *Comparison over multiple datasets:* In order to avoid dataset bias, the analysis should be conducted
  on as many diverse datasets as possible.

- *Rigorous statistical assessment:* The reported results should be the outcome of proper statistical
  procedures, assessing their statistical significance.

459 The workflow of the Statistical Assessment Module is depicted in Figure 12.
460   In a nutshell, each video composing a particular rPPG dataset is processed by the pyVHR pipeline as
461 described above. Moreover, the package provides primitives for loading and processing real BVP signals
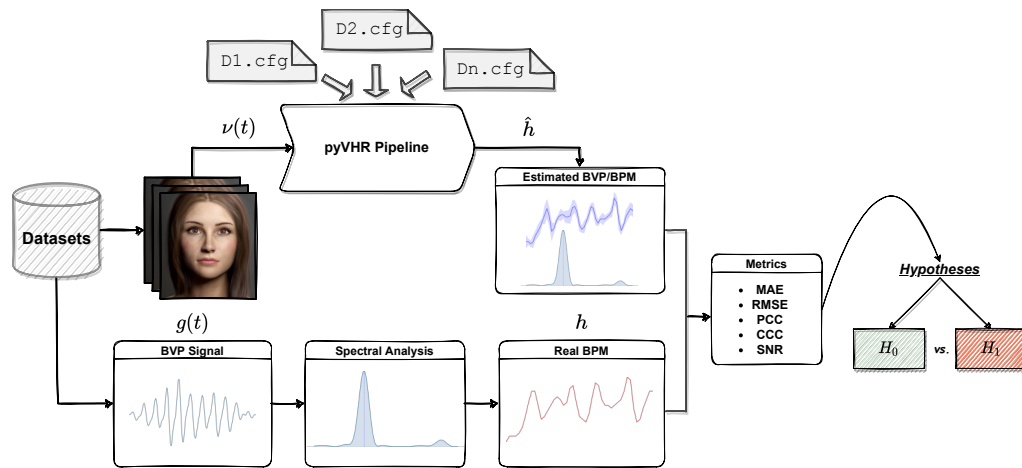
Figure 12. The Assessment Module at a glance. One or more datasets are loaded; videos are processed by the `pyVHR` pipeline while ground-truth BPM signals are retrieved. Predicted and real BPM are compared with standard metrics and the results are rigorously analyzed via hypothesis testing procedures.

as recorded from pulse-oximeters. Such signals undergo a treatment similar to the estimated BVP. In particular, the original BVP signal $g(t)$ is sliced into overlapping time windows; for each window the ground truth BPM h^k (the BPM associated to the k-th time window, with $k = 1, \ldots, K$) is recovered via maximization of the Power Spectral Density (PSD) estimate provided by the Welch's method.

Finally, the estimated ($\hat{h}^k$) and ground truth (h^k) BPM signals are compared with one another exploiting standard metrics (c.f.r Section 5.1). Eventually, statistically rigorous comparisons can be effortlessly performed (c.f.r Section 5.3).

Notably, the many parameters that make up each step of the pipeline (from the ROI selection method to the pre/post filtering operations, passing through the BVP estimation by one or multiple rPPG algorithms) can be easily specified in a configuration (`.cfg`) file. Setting up a `.cfg` file allows to design the experimental procedure in accordance with the principles summarized above. A brief description of the implemented comparison metrics and the `.cfg` file specifications are provided in the following Sections.

5.1 Metrics

`pyVHR` provides common metrics to evaluate the performance of one or more rPPG methods in estimating the correct heart rate (BPM) over time. These are briefly recalled here.

In order to measure the accuracy of the BPM estimate $\hat{h}$, this is compared to the reference BPM as recovered from contact BVP sensors h. To this end, the reference BVP signal $g(t)$ is splitted into overlapping windows, similarly to the procedure described in Section 3.2.4 for the estimated BVP, thus producing K windowed signals g^k ($k \in 1, \ldots, K$). The reference BPM is found via spectral analysis of each window, as described in Section 3.2.6. This yields the K reference BPM h^k to be compared to the estimated one $\hat{h}^k$ by adopting any of the following metrics:

Mean Absolute Error (MAE) The Mean Absolute Error measures the average absolute difference between the estimated $\hat{h}$ and reference BPM h. It is computed as:

$$\text{MAE} = \frac{1}{K} \sum_k |\hat{h}^k - h^k|.$$

Root Mean Squared Error (RMSE). The Root-Mean-Square Error measures the difference between quantities in terms of the square root of the average of squared differences, i.e.

$$\text{RMSE} = \frac{1}{K} \sqrt{\sum_k (\hat{h}^k - h^k)^2}.$$

Pearson Correlation Coefficient (PCC). Pearson Correlation Coefficient measures the linear correlation between the estimate $\hat{h}$ and the ground truth h. It is defined as:

$$PCC = \frac{\sum_k (\hat{h}^k - \hat{\mu})(h^k - \mu)}{\sigma_1 \sigma_2},$$

here $\hat{\mu}$ and μ denote the means of the respective signals, while σ_1 and σ_2 are their standard deviations.

Concordance Correlation Coefficient (CCC). The Concordance Correlation Coefficient (Lawrence and Lin, 1989) is a measure of the agreement between two quantities. Like Pearson's correlation, CCC ranges from -1 to 1, with perfect agreement at 1. It is defined as:

$$CCC = \frac{2\sigma_{12}}{(\hat{\mu} - \mu)^2 + \sigma_1^2 + \sigma_2^2}$$

where $\hat{\mu}$ and μ denote the means of the prediceted and reference BPM traces, respectively. Likewise, σ_1 and σ_2 are their standard deviations, while σ_{12} is their covariance.

Signal to Noise Ratio (SNR). The SNR (De Haan and Jeanne, 2013) measures the ratio of the power around the reference HR frequency plus the first harmonic of the estimated pulse-signal and the remaining power contained in the spectrum of the estimated BVP. Formally it is defined as:

$$SNR = \frac{1}{K} \sum_K 10 \log_{10} \left(\frac{\sum_v \left(U^k(v) S^k(v) \right)^2}{\sum_v \left(1 - U^k(v) \right) S^k(v))^2} \right)$$

where $S^k(v)$ is the power spectral density of the estimated BVP in the k-th time window and $U^k(v)$ is a binary mask that selects the power contained within ± 12 BPM around the reference Heart Rate and its first harmonic.

5.2 The configuration (.cfg) file

The `.cfg` file allows to set up the experimental procedure for the evaluation of models. It is structured into 6 main blocks that are briefly described here:

Dataset. This block contains the information relative to a particular rPPG dataset, namely its name, and its path.

```
1      [DATASET]
2      dataset     = DatasetName
3      path = None
4      videodataDIR= /path/to/vids/
5      BVPdataDIR  = /path/to/gt/
6      ...
7
```

Filters. It defines the filtering methods to be eventually used in the pre/post filtering phase. In the following example a band-pass butterworth filter of 6-th order is defined, with a passing band between 40Hz and 240Hz.

```
1      [BPFILTER]
2      path = None
3      name = BPfilter
4      params = {'minHz':0.65, 'maxHz':4.0, 'fps':'adaptive', 'order':6}
5
```

RGB Signal. Defines all the parameters for the extraction of the RGB signal (e.g. ROI selection method, temporal windowing size, number and type of patches to be used, etc.).

```
1      [SIG]
2      ...
3      winSize = 6
4      skin_extractor = convexhull
5      approach = patches
6      ...
7
```

BVP. Sets up the rPPG method to be adopted for the estimation of the BVP signal. Multiple methods can be provided in order to compare them. In this example two methods will be analyzed, namely `POS` and `GREEN` (adopting their CPU implementations).

```
1    [BVP]
2    methods = ['POS', 'GREEN']
3    ...
4
```

Methods. It allows to configure each rPPG method to be analyzed (e.g. eventual parametrs and pre/post filters). The two methods chosen above are configured here. In particular, `POS` will not employ any pre/post filtering, while for the `GREEN` method, the above-defined band pass filter will be applied for both *pre* and *post* filtering.

```
1    [POS]
2    ...
3    name = cpu_POS
4    device_type = cpu
5    pre_filtering = []
6    post_filtering = []
7
8    [GREEN]
9    ...
10   name = cpu_GREEN
11   device_type = cpu
12   pre_filtering = ['BPFILTER']
13   post_filtering = ['BPFILTER']
14
```

The experiment on the dataset defined in the `.cfg` file can be simply launched as:

```
1  from pyVHR.analysis.pipeline import Pipeline
2
3  pipe = Pipeline()
4  results = pipe.run_on_dataset('/path/to/config.cfg')
5  results.saveResults("/path/to/results.h5")
```

In the above code, the `run_on_dataset` method from the `Pipeline` class, parses the `.cfg` file and initiates a pipeline for each rPPG method defined in it. The pipelines are used to process each video in the dataset. Concurrently, ground truth BPM data is loaded and comparison metrics are computed w.r.t. the predictions (cfr. Figure 12). The results are delivered as a table containing for each method the value of the comparison metrics computed between ground truth and predicted BPM signals, on each video belonging to the dataset, which are then saved to disk. The same considerations hold for the definition of `.cfg` files associated to DL-based methods. Clearly, in this case the information related to the `RGB Signal` block are unnecessary.

5.3 Significance Testing

Once the comparison metrics have been computed for all the considered methods, the significance of the differences between their performance can be evaluated. In other words, we want to ensure that such difference is not drawn by chance, but it represents an actual improvement of one method over another.

To this end, `pyVHR` resorts to standard statistical hypothesis testing procedures. Clearly, the results eventually obtained represent a typical repeated measure design, in which two or more pipelines are compared on paired samples (videos). A number of statistical tests are available in order to deal with such state of affairs.

In the two populations case, typically, the paired t-test is employed; alternatively some non-parametric versions of the paired t-test are at hand, namely the Sign Test or the Wilcoxon signed ranks Test; in general the latter is preferred over the former due to its higher power. For the same reason it is recommended to adopt the parametric paired t-test instead of the non-parametric Wilcoxon test. However, the use of the paired t-test is subject to the constraint of normality of the populations. If such condition is not met, a non-parametric test should be chosen.

Similarly, with more than two pipelines, repeated measure ANOVA is the parametric test that is usually adopted. Resorting to ANOVA, requires Normality and Heteroskedasticity (equality of variances) conditions to be met. Alternatively, when these cannot be ensured, the Friedman Test is chosen.

In `pyVHR` the Normality and Heteroskedasticity conditions are automatically checked via the Shapiro-Wilk Normality test and, depending on the Normality with Levene's test or Bartlett's tests for homogeneity of the data.

In the case of multiple comparisons (ANOVA/Friedman), a proper post-hoc analysis is required in order to establish the pairwise differences among the pipelines. Specifically, the Tukey post-hoc Test is adopted downstream to the rejection of the null hypothesis of ANOVA (the means of the populations are equal), while the Nemenyi post-hoc Test is used after the rejection of the Friedman's null hypothesis of equality of the medians of the samples.

Besides the significance of the differences, it is convenient to report their magnitude, too. The *effect size* can be computed via the Cohen's d in case of Normal of populations; the Akinshin's γ is used otherwise.

5.3.1 The two populations case

`pyVHR` automatically handles the above significance testing procedure within the `StatAnalysis()` class, by relying on the *Autorank* Python package (Herbold, 2020). `StatAnalysis()` ingests the results produced at the previous step and runs the appropriate statistical test on a chosen comparison metric:

```
1 from pyVHR.analysis.stats import StatAnalysis
2
3 st = StatAnalysis("/path/to/results.h5")
4 # -- box plot statistics (medians)
5 st.displayBoxPlot(metric='CCC')
6
7 #testing
8 st.run_stats(metric='CCC')
```

The output of the statistical testing procedure is reported as follows:

> The Shapiro-Wilk Test rejected the null hypothesis of normality for the populations POS ($p < 0.01$) and GREEN ($p < 0.01$). (...) the Wilcoxon's signed rank test has been chosen to determine the differences in the central tendency; median (*MD*) and median absolute deviation (*MAD*) are reported for each population. The test rejected the null hypothesis ($p < 0.01$) that population POS ($MD = 1.344 \pm 1.256$, $MAD = 0.688$) is not greater than population GREEN ($MD = 2.297 \pm 3.217$, $MAD = 1.429$). Hence, we assume that the median of POS is significantly larger than the median of GREEN with a large effect size ($\gamma = -0.850$).

As it can be observed, the appropriate statistical test for two non-normal populations has been properly selected. The Concordance Correlation Coefficient (CCC) for the method POS turned out to be significantly larger than the CCC of the method GREEN. Besides being significant, such difference is substantial, as witnessed by the *large* effect size.

5.3.2 The more-than-two populations case

Suppose now to structure the above `.cfg` in order to run three methods instead of two. This would be as simple as extending the **"BVP"** and **"Methods"** blocks as follows:

```
1    ###BVP###
2    [BVP]
3    methods = ['POS', 'GREEN', 'CHROM']
4    ...
5    ###METHODS###
6    [POS]
7    ...
8
9    [GREEN]
10   ...
11
12   [CHROM]
13   ...
14   name = CHROM
15   pre_filtering = ['BPFILTER']
16   post_filtering = ['BPFILTER']
```

636 Re-running the statistical analysis would yield the following output:

637 The Shapiro-Wilk Test rejected the null hypothesis of normality for the populations CHROM
638 ($p < 0.01$), POS ($p < 0.01$), and GREEN ($p < 0.01$). Given that more than two populations
639 are present, and normality hypothesis has been rejected, the non-parametric Friedman test is
640 chosen to inspect the eventual significant differences between the medians of the populations.
641 The post-hoc Nemenyi test is then used to determine which differences are significant. The
642 Friedman test rejected the null hypothesis ($p < 0.01$) of equality of the medians of the
643 populations CHROM ($MD = 1.263 \pm 1.688$, $MAD = 0.515$, $MR = 1.385$), POS ($MD =$
644 1.344 ± 1.513, $MAD = 0.688$, $MR = 1.769$), and GREEN ($MD = 2.297 \pm 4.569$, $MAD =$
645 1.429, $MR = 2.846$). (...) the post-hoc Nemenyi test revealed no significant differences
646 within the following groups: CHROM and POS, while other differences are significant.

647 Notably, the presence of more than two non-normal populations leads to the choice of the non-parametric
648 Friedman Test as omnibus test to determine if there are any significant differences between the median
649 values of the populations.
650 The box-plots showing the distributions of CCC values for all methods on the UBFC dataset is
651 provided in Figure 13, while the output of the post-hoc Nemenyi test can be visualized through the Critical
652 Difference (CD) diagram (Demšar, 2006) shown in Figure 14; CD Diagrams show the average rank of
653 each method (higher ranks meaning higher average scores); models whose difference in ranks does not
654 exceed the CD_α ($\alpha = 0.05$) are joined by thick lines and cannot be considered significantly different.

Figure 13. Box Plots showing the CCC values distribution for the POS, CHROM and GREEN methods on the UBFC2 dataset.

5.3.3 Comparing Deep and Traditional Pipelines

656 How does a given DL-based rPPG method compares to the above mentioned traditional approaches? The
657 following code snippet allows to run both the traditional and deep pipelines. The results are saved to
658 the same folder, which is then fed as input to the StatAnalysis class; the join_data=True flag
659 allows to merge the results yielded by the two pipelines, thus enabling the statistical comparison between
660 the chosen methods.

```
661  from pyVHR.analysis.stats import StatAnalysis
662  from pyVHR.analysis.pipeline import Pipeline, DeepPipeline
663
664  #Pipeline for Traditional Methods
```

Figure 14. Results of the statistical assessment procedure. CD diagram displaying the results of the Nemenyi post-hoc test on the three populations (`POS`, `CHROM` and `GREEN`) of CCC values on the `UBFC2` dataset.

```
5   traditional_pipe = Pipeline()
6   traditional_results = traditional_pipe.run_on_dataset('/path/to/trad_config.cfg')
7   traditional_results.saveResults("/path/to/results_folder/traditional_results.h5")
8
9   #Pipeline for Deep Methods
10  deep_pipe = DeepPipeline()
11  deep_results = deep_pipe.run_on_dataset('/path/to/deep_config.cfg')
12  deep_results.saveResults("/path/to/results_folder/deep_results.h5")
13
14  #Statistical Analysis
15  st = StatAnalysis("/path/to/results_folder/", join_data=True)
16  # -- box plot statistics
17  st.displayBoxPlot(metric='SNR')
18  #Significance testing on the SNR metric
19  st.run_stats(metric='SNR')
```

In this case, the Signal-to-Noise Ratio (`SNR`) has been chosen as comparison metric; Figure 15 qualitatively displays the results of the comparison of the above mentioned traditional methods with the `MTTS-CAN` DL-based approach (Liu et al., 2020). The outcome of the statistical assessment is shown in the CD diagram of Figure 16.

6 EXTENDING THE FRAMEWORK

Besides assessing built-in methods on public datasets included in the framework, the platform is conceived to allow the addition of new methods or datasets. This way, it is possible to assess a new proposal, comparing it against built-in methods, and testing it on either already included datasets or on new ones, this exploiting all the pre- and post-processing modules made available in `pyVHR` . The framework extension can be achieved following simple steps as described in the subsequent subsections.

6.1 Adding a new method

In this section we show how to add to the `pyVHR` framework either a new traditional or learning-based method called MY_NEW_METHOD.

In the first case, to exploit the `pyVHR` built-in modules the new function should receive as input a `signal` in the shape produced by the built-in pre-processing modules, together with some other parameters required by the method itself. Specifically, this results in a signature of the form:

```
1   MY_NEW_METHOD(signal, **kargs)
```

where `signal` is a Numpy array in the form (`P, 3, K`); `P` is the number of considered patches (it can be 1 if the holistic approach is used), 3 is the number of RGB Channels and `K` is the number of frames. `**kargs` refers to a dictionary that contains all the parameters required by the method at hand. A proper function implementing an rPPG method must return a BVP signal as a Numpy array of shape (`P,K`).

Figure 15. **Box Plots showing the SNR values distribution for the `POS`, `CHROM`, `MTTS−CAN` and `GREEN` methods on the `UBFC1` dataset.**

701   In case of DL-based method, the new function should receive as input the raw `frames` as a Numpy
702 array in the form `(H,W,3,K)`, where `H,W` denote the frame dimensions. The output of the new method
703 could be either a BVP signal or the HR directly.
704   Accordingly, the signature becomes:

```
705 1 MY_NEW_METHOD(frames, fps)
```

706   Both for traditional and DL-based method, the function call `MY_NEW_METHOD` can now be embedded
707 into the proper `Pipeline`, and assessed as described earlier. In order to do so, the `.cfg` file should be
708 tweaked as follows:

```
709 1 [MY_NEW_METHOD]
710 2 path = 'path/to/module.py'
711 3 name = 'MY_NEW_METHOD'
712 4 ...
```

713   Moreover, the `methods` block of the `.cfg` file is supposed to contain a specific listing describing
714 `MY_NEW_METHOD`, providing the path to the Python module encoding the method and its function name.

6.2 Adding a new dataset

716 Currently `pyVHR` provides APIs for handling five datasets commonly adopted for the evaluation of rPPG
717 methods, namely LGI-PPGI (Pilz et al., 2018), UBFC (Bobbia et al., 2019), PURE (Stricker et al., 2014),
718 MAHNOB-HCI (Soleymani et al., 2011), and COHFACE (Heusch et al., 2017a). However, the platform
719 allows to add new datasets favoring the method assessment on new data. A comprehensive list of the
720 datasets that are typically employed for rPPG estimation and evaluation is reported in Table 3.
721   The framework conceives datasets as a hierarchy of classes (see Fig.17) that allows to describe a new
722 dataset by inheriting from the `Dataset` base class and implementing few methods for loading videos
723 and ground truth PPG data.
724   Specifically, the following two functions should be supplied:

725   • a `loadFilenames()` function to load video files in a Python list; this function has no inputs and
726   defines two class variables, namely `videoFilenames` and `BVPFilenames`. These are both
727   Python lists containing, respectively, video and ground-truth BVP filenames from the dataset);

728   • a `readSigfile(filename)` function loading and returning the ground-truth BVP signal given
729   a video filename.

Figure 16. Results of the statistical assessment procedure. CD diagram displaying the results of the Nemenyi post-hoc test on the four populations (`POS`, `CHROM`, `MTTS-CAN` and `GREEN`) of SNR values on the `UBFC1` dataset.

Dataset	Subjects	Task/Condition	`pyVHR`
UBFC1 (Bobbia et al., 2019)	8	Stationary	✓
UBFC2 (Bobbia et al., 2019)	42	Interaction	✓
PURE (Stricker et al., 2014)	10	Multiple	✓
LGI-PPGI (Pilz et al., 2018)	25 (6 available)	Multiple	✓
MAHNOB-HCI (Soleymani et al., 2011)	27	Emotion Elicitation	✓
COHFACE (Heusch et al., 2017b)	40	Stationary	✓
UBFC-Phys (Sabour et al., 2021)	56	Stress Test	✗
AFRL (Estepp et al., 2014)	25	Multiple	✗
MMSE-HR (Zhang et al., 2016)	140	Simulating Facial Expressions	✗
OBF (Li et al., 2018)	106	Multiple	✗
VIPL-HR (Niu et al., 2018)	107	Multiple	✗
ECG-Fitness (Špetlík et al., 2018)	17	Physical Activities	✗

Table 3. A list of datasets commonly used for rPPG. The left-most column collects the dataset names and introducing papers; second column, the number of subjects involved; third column, the task or condition under which data have been collected (Stationary: subject are asked to sit still; Interaction: emulation of a human-computer interaction scenario via a time sensitive mathematical game; Multiple: more than one condition has been considered while recording subjects, such as Steady, Talking, Head Motion etc; Physical Activities: subjects are recorded while performing activities such as speaking, rowing, exercising on a stationary bike etc; Stress Test: participants are subject to tasks with different levels of difficulty inspired by the Trier Social Stress Test; Emotion Elicitation: participants were shown fragments of movies and pictures apt at eliciting emotional reactions). In the last column, datasets whose handling APIs are currently available in pyVHR have been checked.

730   The new dataset can then be included in the testing via the `.cfg` file as described in the paragraph
731 *Dataset* of Section 5.2. As for the addition of new method, also in case of adding a new dataset the `.cfg`
732 file should be completed by specifying the path pointing to the new dataset class:

```
1    [DATASET]
2    dataset     = DatasetName
3    path = /path/to/datasetClass/
4    videodataDIR= /path/to/videos/
5    BVPdataDIR  = /path/to/gtfiles/
6    ...
7
```

Figure 17. **Class diagram of dataset hierarchy of classes.**

7 CASE STUDY : DEEPFAKE DETECTION WITH PYVHR

DeepFakes are a set of DL based techniques allowing to create fake videos by swapping the face of a person by that of another. This technology has many diverse applications such as expression re-enactment (Bansal et al., 2018) or video de-identification (Bursic et al., 2021). However, in recent years the quality of deepfakes has reached tremendous levels of realism, thus posing a series of treats related to the possibility of arbitrary manipulation of identity, such as political propaganda, blackmailing, and fake news (Mirsky and Lee, 2021).

As a consequence, efforts have been devoted to the study and the development of methods allowing to discriminate between real and forged videos (Tolosana et al., 2020; Mirsky and Lee, 2021). Interestingly enough, one effective approach is represented by the exploitation of physiological information (Hernandez-Ortega et al., 2020; Ciftci et al., 2020; Qi et al., 2020) . Indeed, signals originating from biological action such as heart beat, blood flow, or breathing are expected to be (in large part) disrupted after face-swapping. Therefore, methods such as remote PPG can be adopted in order to evaluate their presence.

In the following, it is shown how `pyVHR` can be effectively employed to easily perform a DeepFake detection task. To this end, we rely on the FaceForensics++[4] dataset (Rössler et al., 2019) consisting of 1000 original video sequences (mostly frontal face without occlusions) that have been manipulated with four automated face manipulation methods.

Each video, either original or swapped is fed as input to the `pyVHR` pipeline; then, the estimated BVPs and the predicted BPMs can be analyzed in order to detect DeepFakes. It is reasonable to imagine that the BVP signals estimated on original videos would have much lower complexity if compared with the swapped ones, due to the stronger presence of PPG related information that would be possibly ruled out during swapping procedures. As a consequence, BVP signals from DeepFakes would perhaps exhibit higher levels of noise and hence more complex behaviour.

There exist many ways of measuring the complexity of a signal; here we choose to compute the Fractal Dimension (FD) of BVPs; in particular the Katz's method (Katz, 1988) is employed.

The FD of the BVP estimated from the i-th patch on the k-th time window (D_i^k) can be computed as (Katz, 1988):

$$D_i^k = \frac{\log_{10}(L/a)}{\log_{10}(d/a)},$$

where L is the sum of distances between successive points, a is their average, and d is the maximum distance between the first point and any other point of the estimated BVP signal.

The FD associated to a given video can then be obtained via averaging:

---

[4]Available at: https://github.com/ondyari/FaceForensics.

$$\hat{FD}_{vid} = \frac{1}{PK} \sum_{i=0}^{P} \sum_{k=0}^{K} D_i^k.$$

Similarly, one could consider adopting the average Median Absolute Deviation (MAD) of the BPM predictions on a video as a predictor of the presence of DeepFakes:

$$\hat{MAD}_{vid} = \frac{1}{K} \sum_{k=0}^{K} MAD^k.$$

Figure 18 shows how the FaceForensics++ videos lie in the 2-dimensional space defined by the average Fractal Dimension ($\hat{FD}$) of predicted BVPs using the POS method and the average MADs of BPM predictions ($\hat{MAD}$), when considering the original and swapped videos with the *FaceShifter* method.

Figure 18. Deepfake detection results. The 1000 FaceForensics++ original videos (blue) and their swapped versions (yellow) represented in the 2-D space of BVP Fractal Dimension *vs.* BPMs average MAD. The green and red half-spaces are simply learned via a Linear SVM.

It is easy to see how adopting these simple statistics on pyVHR's predictions allows to discriminate original videos from DeepFakes. In particular, learning a baseline Linear SVM for the classification of Real *vs.* Fake videos generated by the *FaceShifter* method, yields an average 10-fold Cross-Validation Accuracy of $91.41\% \pm 2.05$. This result is comparable with state of the art approaches usually adopting much more complex solutions.

CONCLUSIONS

In recent years, the rPPG-based pulse rate recovery has attracted much attention due to its promise to reduce invasiveness, while granting higher and higher precision in heart rate estimation. In particular, we have witnessed the proliferation of rPPG algorithms and models that accelerate the successful deployment in areas that traditionally exploited wearable sensors or ambulatory monitoring. These two trends, combined together, have fostered a new perspective in which advanced video-based computing techniques play a fundamental role in replacing the domain of physical sensing.

In this paper, in order to allow the rapid development and the assessment of new techniques, we presented an open and very general framework, namely pyVHR . It allows for a careful study of every step, and no less important, for a sound comparison of methods on multiple datasets.

pyVHR is a re-engineered version of the framework presented in Boccignone et al. (2020a) but exhibiting substantial novelties:

792 • Ease of installation and use.

793 • Two distinct pipelines for either traditional or DL-based methods.

794 • Holistic or patch processing for traditional approaches.

795 • Acceleration by GPU architectures.

796 • Ease of extension (adding new methods or new datasets).

797 The adoption of GPU support allows the whole process to be safely run in real-time for 30 fps HD
798 videos and an average speedup ($time_{seq}/time_{parall}$) of around 5.

799 Besides addressing the challenges of remote Heart Rate monitoring, we also expect that this framework
800 will be useful to researchers and practitioners from various disciplines when dealing with new problems
801 and building new applications leveraging rPPG technology.

802 REFERENCES

803 Aarts, L. A., Jeanne, V., Cleary, J. P., Lieber, C., Nelson, J. S., Oetomo], S. B., and Verkruysse, W. (2013).
804   Non-contact heart rate monitoring utilizing camera photoplethysmography in the neonatal intensive
805   care unit — a pilot study. *Early Human Development*, 89(12):943 – 948.
806 Balakrishnan, G., Durand, F., and Guttag, J. (2013). Detecting pulse from head motions in video. In
807   *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3430–3437.
808 Bansal, A., Ma, S., Ramanan, D., and Sheikh, Y. (2018). Recycle-gan: Unsupervised video retargeting.
809   In *Proceedings of the European conference on computer vision (ECCV)*, pages 119–135.
810 Benavoli, A., Corani, G., Demšar, J., and Zaffalon, M. (2017). Time for a change: a tutorial for
811   comparing multiple classifiers through bayesian analysis. *The Journal of Machine Learning Research*,
812   18(1):2653–2688.
813 Benezeth, Y., Li, P., Macwan, R., Nakamura, K., Gomez, R., and Yang, F. (2018). Remote heart rate
814   variability for emotional state monitoring. In *2018 IEEE EMBS International Conference on Biomedical
815   & Health Informatics (BHI)*, pages 153–156. IEEE.
816 Blazek, V. and Schultz-Ehrenburg, U. (1996). *Quantitative Photoplethysmography: Basic Facts and
817   Examination Tests for Evaluating Peripheral Vascular Funktions [ie Functions]*. VDI-Verlag.
818 Bobbia, S., Macwan, R., Benezeth, Y., Mansouri, A., and Dubois, J. (2019). Unsupervised skin tissue
819   segmentation for remote photoplethysmography. *Pattern Recognition Letters*, 124:82–90.
820 Boccignone, G., Conte, D., Cuculo, V., D'Amelio, A., Grossi, G., and Lanzarotti, R. (2020a). An open
821   framework for remote-ppg methods and their assessment. *IEEE Access*, 8:216083–216103.
822 Boccignone, G., de'Sperati, C., Granato, M., Grossi, G., Lanzarotti, R., Noceti, N., and Odone, F. (2020b).
823   Stairway to elders: Bridging space, time and emotions in their social environment for wellbeing. In
824   *ICPRAM*, pages 548–554.
825 Bursic, S., D'Amelio, A., Granato, M., Grossi, G., and Lanzarotti, R. (2021). A quantitative evaluation
826   framework of video de-identification methods. In *2020 25th International Conference on Pattern
827   Recognition (ICPR)*, pages 6089–6095. IEEE.
828 Chen, W. and McDuff, D. (2018). Deepphys: Video-based physiological measurement using convolutional
829   attention networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages
830   349–365.
831 Cheng, C.-H., Wong, K.-L., Chin, J.-W., Chan, T.-T., and So, R. H. (2021). Deep learning methods for
832   remote heart rate measurement: A review and future research agenda. *Sensors*, 21(18):6296.
833 Ciftci, U. A., Demir, I., and Yin, L. (2020). How do the hearts of deep fakes beat? deep fake source
834   detection via interpreting residuals with biological signals. In *2020 IEEE International Joint Conference
835   on Biometrics (IJCB)*, pages 1–10. IEEE.
836 De Haan, G. and Jeanne, V. (2013). Robust pulse rate from chrominance-based rppg. *IEEE Transactions
837   on Biomedical Engineering*, 60(10):2878–2886.
838 de Haan, G. and van Leest, A. (2014). Improved motion robustness of remote-PPG by using the blood
839   volume pulse signature. *Physiological Measurement*, 35(9):1913–1926.
840 Demšar, J. (2006). Statistical comparisons of classifiers over multiple data sets. *J. Mach. Learn. Res.*,
841   7:1–30.

Eisinga, R., Heskes, T., Pelzer, B., and Te Grotenhuis, M. (2017). Exact p-values for pairwise comparison of friedman rank sums, with application to comparing classifiers. *BMC Bioinformatics*, 18(1):68.

Estepp, J. R., Blackford, E. B., and Meier, C. M. (2014). Recovering pulse rate during motion artifact with a multi-imager array for non-contact imaging photoplethysmography. In *2014 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pages 1462–1469. IEEE.

Gideon, J. and Stent, S. (2021). The way to my heart is through contrastive learning: Remote photoplethysmography from unlabelled video. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3995–4004.

Google (2021). Mediapipe face mesh.

Graczyk, M., Lasota, T., Telec, Z., and Trawiński, B. (2010). Nonparametric statistical analysis of machine learning algorithms for regression problems. In Setchi, R., Jordanov, I., Howlett, R. J., and Jain, L. C., editors, *Knowledge-Based and Intelligent Information and Engineering Systems*, pages 111–120. Springer Berlin Heidelberg.

Herbold, S. (2020). Autorank: A python package for automated ranking of classifiers. *Journal of Open Source Software*, 5(48):2173.

Hernandez-Ortega, J., Tolosana, R., Fierrez, J., and Morales, A. (2020). Deepfakeson-phys: Deepfakes detection based on heart rate estimation. *arXiv preprint arXiv:2010.00400*.

Hertzman, A. B. (1937). Photoelectric plethysmography of the fingers and toes in man. *Proceedings of the Society for Experimental Biology and Medicine*, 37(3):529–534.

Heusch, G., Anjos, A., and Marcel, S. (2017a). A reproducible study on remote heart rate measurement. *CoRR*, abs/1709.00962.

Heusch, G., Anjos, A., and Marcel, S. (2017b). A reproducible study on remote heart rate measurement. *arXiv preprint arXiv:1709.00962*.

Humphreys, K., Ward, T., and Markham, C. (2007). Noncontact simultaneous dual wavelength photoplethysmography: A further step toward noncontact pulse oximetry. *Review of Scientific Instruments*, 78(4):044304.

Katz, M. J. (1988). Fractals and the analysis of waveforms. *Computers in biology and medicine*, 18(3):145–156.

Koelstra, S., Muhl, C., Soleymani, M., Lee, J.-S., Yazdani, A., Ebrahimi, T., Pun, T., Nijholt, A., and Patras, I. (2011). Deap: A database for emotion analysis; using physiological signals. *IEEE transactions on affective computing*, 3(1):18–31.

Lawrence, I. and Lin, K. (1989). A concordance correlation coefficient to evaluate reproducibility. *Biometrics*, pages 255–268.

Lewandowska, M., Rumiński, J., Kocejko, T., and Nowak, J. (2011). Measuring pulse rate with a webcam - a non-contact method for evaluating cardiac activity. In *2011 Federated Conference on Computer Science and Information Systems (FedCSIS)*, pages 405–410.

Li, X., Alikhani, I., Shi, J., Seppanen, T., Junttila, J., Majamaa-Voltti, K., Tulppo, M., and Zhao, G. (2018). The obf database: A large face video database for remote physiological signal measurement and atrial fibrillation detection. In *2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018)*, pages 242–249. IEEE.

Liu, X., Fromm, J., Patel, S., and McDuff, D. (2020). Multi-task temporal shift attention networks for on-device contactless vitals measurement. *Advances in Neural Information Processing Systems*, 33:19400–19411.

Liu, X., Jiang, Z., Fromm, J., Xu, X., Patel, S., and McDuff, D. (2021). Metaphys: few-shot adaptation for non-contact physiological measurement. In *Proceedings of the Conference on Health, Inference, and Learning*, pages 154–163.

McDuff, D. (2021). Camera measurement of physiological vital signs. *arXiv preprint arXiv:2111.11547*.

McDuff, D. and Blackford, E. (2019). iPhys: An open non-contact imaging-based physiological measurement toolbox. In *2019 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 6521–6524. IEEE.

McDuff, D., Gontarek, S., and Picard, R. (2014). Remote measurement of cognitive stress via heart rate variability. In *2014 36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 2957–2960.

McDuff, D. J., Estepp, J. R., Piasecki, A. M., and Blackford, E. B. (2015). A survey of remote optical photoplethysmographic imaging methods. In *2015 37th Annual International Conference of the IEEE*

897   *Engineering in Medicine and Biology Society (EMBC)*, pages 6398–6404.

898 Mirsky, Y. and Lee, W. (2021). The creation and detection of deepfakes: A survey. *ACM Computing*
899   *Surveys (CSUR)*, 54(1):1–41.

900 Ni, A., Azarang, A., and Kehtarnavaz, N. (2021). A review of deep learning-based contactless heart rate
901   measurement methods. *Sensors*, 21(11):3719.

902 Niu, X., Han, H., Shan, S., and Chen, X. (2018). Vipl-hr: A multi-modal database for pulse estimation
903   from less-constrained face video. In *Asian Conference on Computer Vision*, pages 562–576. Springer.

904 Niu, X., Shan, S., Han, H., and Chen, X. (2019). Rhythmnet: End-to-end heart rate estimation from face
905   via spatial-temporal representation. *IEEE Transactions on Image Processing*, 29:2409–2423.

906 Nowara, E., McDuff, D., and Veeraraghavan, A. (2020). The benefit of distraction: Denoising remote
907   vitals measurements using inverse attention. *arXiv preprint arXiv:2010.07770*.

908 Pilz, C. (2019). On the vector space in photoplethysmography imaging. In *Proceedings of the IEEE/CVF*
909   *International Conference on Computer Vision Workshops*, pages 0–0.

910 Pilz, C. S., Zaunseder, S., Krajewski, J., and Blazek, V. (2018). Local group invariance for heart rate
911   estimation from face videos in the wild. In *Proceedings of the IEEE Conference on Computer Vision*
912   *and Pattern Recognition Workshops*, pages 1254–1262.

913 Poh, M.-Z., McDuff, D. J., and Picard, R. W. (2010). Non-contact, automated cardiac pulse measurements
914   using video imaging and blind source separation. *Optics express*, 18(10):10762–10774.

915 Qi, H., Guo, Q., Juefei-Xu, F., Xie, X., Ma, L., Feng, W., Liu, Y., and Zhao, J. (2020). Deeprhythm:
916   Exposing deepfakes with attentional visual heartbeat rhythms. In *Proceedings of the 28th ACM*
917   *International Conference on Multimedia*, pages 4318–4327.

918 Ramírez, G. A., Fuentes, O., Crites, S. L., Jimenez, M., and Ordonez, J. (2014). Color analysis of facial
919   skin: Detection of emotional state. *2014 IEEE Conference on Computer Vision and Pattern Recognition*
920   *Workshops*, pages 474–479.

921 Rössler, A., Cozzolino, D., Verdoliva, L., Riess, C., Thies, J., and Nießner, M. (2019). FaceForensics++:
922   Learning to detect manipulated facial images. In *International Conference on Computer Vision (ICCV)*.

923 Rouast, P. V., Adam, M. T., Cornforth, D. J., Lux, E., and Weinhardt, C. (2017). Using contactless
924   heart rate measurements for real-time assessment of affective states. In *Information Systems and*
925   *Neuroscience*, pages 157–163. Springer.

926 Rouast, P. V., Adam, M. T. P., Chiong, R., Cornforth, D., and Lux, E. (2018). Remote heart rate
927   measurement using low-cost rgb face video: a technical literature review. *Frontiers of Computer*
928   *Science*, 12(5):858–872.

929 Sabour, R. M., Benezeth, Y., De Oliveira, P., Chappe, J., and Yang, F. (2021). Ubfc-phys: A multimodal
930   database for psychophysiological studies of social stress. *IEEE Transactions on Affective Computing*.

931 Soleymani, M., Lichtenauer, J., Pun, T., and Pantic, M. (2011). A multimodal database for affect
932   recognition and implicit tagging. *IEEE transactions on affective computing*, 3(1):42–55.

933 Špetlík, R., Franc, V., and Matas, J. (2018). Visual heart rate estimation with convolutional neural network.
934   In *Proceedings of the british machine vision conference, Newcastle, UK*, pages 3–6.

935 Stricker, R., Müller, S., and Gross, H.-M. (2014). Non-contact video-based pulse rate measurement on a
936   mobile service robot. In *The 23rd IEEE International Symposium on Robot and Human Interactive*
937   *Communication*, pages 1056–1062. IEEE.

938 Tarassenko, L., Villarroel, M., Guazzi, A., Jorge, J., Clifton, D. A., and Pugh, C. (2014). Non-contact
939   video-based vital sign monitoring using ambient light and auto-regressive models. *Physiological*
940   *Measurement*, 35(5):807–831.

941 Tolosana, R., Vera-Rodriguez, R., Fierrez, J., Morales, A., and Ortega-Garcia, J. (2020). Deepfakes and
942   beyond: A survey of face manipulation and fake detection. *Information Fusion*, 64:131–148.

943 Torralba, A. and Efros, A. A. (2011). Unbiased look at dataset bias. In *CVPR 2011*, pages 1521–1528.
944   IEEE.

945 Unakafov, A. M. (2018). Pulse rate estimation using imaging photoplethysmography: generic framework
946   and comparison of methods on a publicly available dataset. *Biomedical Physics & Engineering Express*,
947   4(4):045001.

948 Verkruysse, W., Svaasand, L. O., and Nelson, J. S. (2008). Remote plethysmographic imaging using
949   ambient light. *Opt. Express*, 16(26):21434–21445.

950 Wang, W., den Brinker, A. C., Stuijk, S., and de Haan, G. (2016). Algorithmic principles of remote ppg.
951   *IEEE Transactions on Biomedical Engineering*, 64(7):1479–1491.

Wang, W., Stuijk, S., and De Haan, G. (2015). A novel algorithm for remote photoplethysmography: Spatial subspace rotation. *IEEE transactions on biomedical engineering*, 63(9):1974–1984.

Wieringa, F. P., Mastik, F., and Steen, A. F. W. v. d. (2005). Contactless multiple wavelength photoplethysmographic imaging: A first step toward "spo2 camera"technology. *Annals of Biomedical Engineering*, 33(8):1034–1041.

Yu, C., Wang, J., Peng, C., Gao, C., Yu, G., and Sang, N. (2018). Bisenet: Bilateral segmentation network for real-time semantic segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 325–341.

Yu, Z., Li, X., Niu, X., Shi, J., and Zhao, G. (2020). Autohr: A strong end-to-end baseline for remote heart rate measurement with neural searching. *IEEE Signal Processing Letters*, 27:1245–1249.

Yu, Z., Shen, Y., Shi, J., Zhao, H., Torr, P., and Zhao, G. (2021). Physformer: Facial video-based physiological measurement with temporal difference transformer. *arXiv preprint arXiv:2111.12082*.

Zhang, Z., Girard, J. M., Wu, Y., Zhang, X., Liu, P., Ciftci, U., Canavan, S., Reale, M., Horowitz, A., Yang, H., et al. (2016). Multimodal spontaneous emotion corpus for human behavior analysis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3438–3446.