



**IM-TWIN: from Intrinsic Motivations
to Transitional Wearable INTelligent
companions for autism spectrum disorder**
a European funded project

***Processing of physiological signals, visual
info, and PlusMe interaction, first version***

Deliverable 2.1



This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 952095.

Project duration 36 months (November 2020, October 2023).
Consortium: Consiglio Nazionale delle Ricerche (ITA),
Universiteit Utrecht (NLD), Centre de Recherches
Interdisciplinaires (FRA), Università degli Studi di Roma
La Sapienza (ITA), Plux-Wireless Biosignals S.A. (PRT).

Deliverable data

Work Package:	2 Affective signal processing through the integration of multiple sources
Work Package leader:	UU
Deliverable beneficiary:	UU
Dissemination level:	Public
Due date:	31 th December (Month 26)
Type:	Report
Authors:	E.L. van den Broek, Lukas P.A. Arts, Francesco Montedori, Valerio Sperati

Acronyms of partners

CNR-ISTC	Consiglio Nazionale delle Ricerche, Istituto di Scienze e Tecnologie della Cognizione (Italy)
UU	Universiteit Utrecht (The Netherlands)
CRI	Centre de Recherches Interdisciplinaires (France)
LA SAPIENZA	Università degli Studi di Roma La Sapienza (Italy)
PLUX	Plux - Wireless Biosignals S.A. (Portugal)

Table of contents

1. Overview of the deliverable	4
2. Processing of physiological signals	4
2.1 Assessing dry electrode signal quality	4
2.2. Signal Quality Indicator (SQI)	13
2.3 Processing non-stationary signals and open-source repository	18
3. Processing of visual information	23
4. Processing of interaction between child, PlusMe and therapist	29
5. Conclusions and future developments	31

1. Overview of the deliverable

The core of this deliverable contains the following sections:

2. *Processing of physiological signals*, containing the subsections:
 - 2.1. *Assessing dry electrode signal quality*;
 - 2.2. *Signal Quality Indicator (SQI)*;
 - 2.3. *Processing non-stationary signals*, including open source code repository;
3. *Processing of visual information*, containing a description of the AI-based system to detect the eye contact between child and caregiver;
4. *Processing of interaction between child, PlusMe and therapist*, containing a description of the system to detect relevant social events during an experimental session;
5. *Conclusions and future developments*, containing the overall conclusions and the plan for the next system development.

2. Processing of physiological signals

2.1 Assessing dry electrode signal quality

2.1.1. Introduction

Extracting high-quality biosignals for reliable emotional state estimation of children is one of the main challenges of the ERC-funded Horizon 2020 IM-TWIN EIC path-finder project. Unfortunately, the situation is complicated even further as we cannot use the well-known, golden standard gelled (e.g., wet) electrodes. We want to avoid wet electrodes as they can only be used once and require a tedious attachment process. As children are generally not known for their patience and understanding during medical examinations, dry electrodes, which do not require gel, are preferred. Dry electrodes could be textile embedded, and the setup procedure would be as simple as putting on a tight-fitted t-shirt. This removes the need for medical assistance, and one can reuse the equipment many times.

However, dry electrodes have quite some disadvantages. As gel acts as a conductor between the skin and an electrode, measuring the tiny electrophysiological signals without gel is hard. Dry electrodes tend to have more difficulty measuring high-quality biosignals as the cushion of air between the skin and the electrode acts as a giant transistor, thereby masking the small biopotentials. Figure 1 from [2] shows the difference between wet and dry electrodes in more detail. Mathematically R_i is often larger than R_g .

Besides the absence of gel, dry electrodes are also not fixed to the skin. Therefore, biosignals recorded with dry electrodes are very susceptible to body movement as this causes the electrodes to shift across the skin. As sweat is the main conductor in dry electrodes, the shifting

causes the availability of sweat beneath the electrode to vary. Hence, conductance varies significantly, and signal artifacts arise.

A small pilot study was performed in which we analyzed an Electrocardiogram (ECG) measured with wet and dry electrodes on noise level and usability. In contrast to many other studies, we aimed to investigate the performance of an easy off-the-shelf implementation of dry textile-based electrodes in this study. We deliberately did not use advanced techniques to embed dry electrodes into the fabric to explore the challenges of textile-based electrodes and set a baseline performance on which improvements could be made.

2.1.2. Materials and Methods

This study aims to determine the primary challenges in using dry electrodes in textile-based applications. In this study, both dry and wet electrodes are used to compare performance. We focused mainly on exploration; we did not build custom-made hardware but used the commercially available Biosignalsplux kit from PLUX (Portugal) instead. The Biosignalsplux kit records up to 8 signals simultaneously using a small, portable recording unit. In turn, the recording unit sends the data via Bluetooth to a nearby computer. As the recording unit operates wirelessly, patients can move freely, and obtrusiveness is kept to a minimum.




Figure 1: Electrode-skin interaction of wet and dry electrodes from [2]. The corresponding electrical circuit representation of each situation is also shown.

2.1.2.1. Dry electrode setup

The dry electrode setup was created by attaching three commercially available non-gelled Ag/AgCl electrodes to a tight-fitted thermal shirt. The thermal shirt by Kipsta was selected such that the shirt sat tight around the entire upper body, thereby mimicking the textile decision of IM-TWIN. Three holes were punched in the thermal shirt for the positive (red), negative (black), and reference (white) electrode. Electrodes were placed in the V2 configuration of the 12-lead ECG scheme (see Figure 2a). Reference was placed on the sternum due to the relatively short

wiring (2-3 cm). Three electrode plates were attached to the electrodes such that the thermal shirt was clamped in between (see Figure 2b).



(a) Outside view

(b) Inside view

Figure 2: Dry electrode setup. ECG was measured using three non-gelled Ag/AgCl electrodes which were embedded into a tight fitted thermoshirt. Electrode plates were attached to the wires such that the thermoshirt was clamped in between.

The commercially available dry electrodes are generally delivered with a blue plastic stabilization ring attached to increase the fixation area. We noticed that this ring could sometimes hinder surface contact if the electrode was pressed into the skin at an angle. To remove this effect, an additional setup was created in which the blue plastic ring was removed (see Figure 3).



(a) Overview

(b) Closeup

Figure 3: To increase the chance of good contact between the skin and the electrode, an additional setup was created in which the blue stabilization ring was removed from the commercially available dry electrodes.

2.1.2.2. Wet electrode reference setup

In the wet electrode setup, the electrodes were pre-gelled and attached to the un-shaven skin in a similar V2 configuration as the dry electrode setup (see Figure 4). The skin was not shaven or cleaned before attachment to stay as close to the dry electrode setup as possible. This way, only the usage of gel and adhesive electrodes was added to the setup. Other factors were tried to be unchanged as well as possible.



Figure 4: Wet electrode setup. Gelled electrodes were attached to the skin using an adhesive membrane to function as a baseline for the dry electrode setup. Electrodes were configured in the same V2 configuration as the dry electrodes.

2.1.2.3. Biosignal choice

As this study focused mainly on exploration, we only investigated the performance of one biosignal; in this case, the ECG signal. ECG is the most used biosignal, and signal quality can be easily compared due to its low latency and continuous character. At the same time, ECG is one of the hardest signals to obtain, seen from a physiological standpoint. Where most electrophysiological signals are obtained using only one or two electrodes (Skin Temperature (SKT), Electrodermal Activity (EDA), Electromyography (EMG)), ECG uses three electrodes. The extra electrode acts as a reference or ground electrode.

Measuring ECG using the V2 configuration of the 12-lead ECG scheme allows one to measure the polarization and depolarization (i.e., contraction) of the atrial and ventricular muscles of the heart (see Figure 5). First, the atrial muscle is being depolarised, causing a small bump in the ECG called the P-Wave. Next, the large ventricle muscles are depolarised, causing the well-known QRS-complex to occur. Finally, repolarisation of the ventricle muscles

creates the T-Wave. Although ECG is believed to measure heart muscle activation, during the T-Wave, there is no muscle activation going on, only repolarisation. The heart is at rest.




Figure 5: Schematic representation of the waves and complexes found in an ECG signal. Each wave is generated by the polarization or depolarisation of the muscles surrounding the heart.

2.1.2.4. Design

This exploratory study examines the performance of several dry electrode implementations and one wet, reference electrode implementation. During all experiments, the participant was seated on a chair and tried to move as little as possible to prevent motion artifacts. Each implementation was tested for a brief period during which signal quality was visually assessed and noted down. As the study was meant for experimental purposes only, no strict study design was created beforehand. Especially during the dry electrode experiments, several variations were tested on the fly. All variations will be described in the results section in chronological order.

In the scope of the IM-TWIN project, we are most interested in the detection of heart rate and heart rate variability as these two are well-known estimators for emotional arousal. As the QRS-complex can be detected easiest, heart rate is generally estimated using QRS-complexes. By measuring the duration between two R-peaks (i.e., the RR-interval), one can easily obtain an estimation of heart rate and heart rate variability. Consequently, in this exploratory study, we will mainly focus on the possibility of extracting the QRS-complex. Successful extraction of the QRS-complex, independent of noise, makes a signal usable in a real-world application. We are only interested in the possibility of extracting the QRS-complex, not the actual extraction itself. Therefore, as was mentioned before, we will assess each signal only by eye instead of applying an actual processing pipeline.

2.1.3 Results

First, the reference ECG signal was recorded using the wet electrode setup (see Figure 4). The resulting signal was largely noise-free, as expected. The P, QRS, and T waves can be easily recognized as seen in Figure 6b. The well-known 50/60 Hz power-line interference probably causes the high-frequency noise imposed on the signal. Next, the dry electrode setup with the default blue plastic rings was tested.

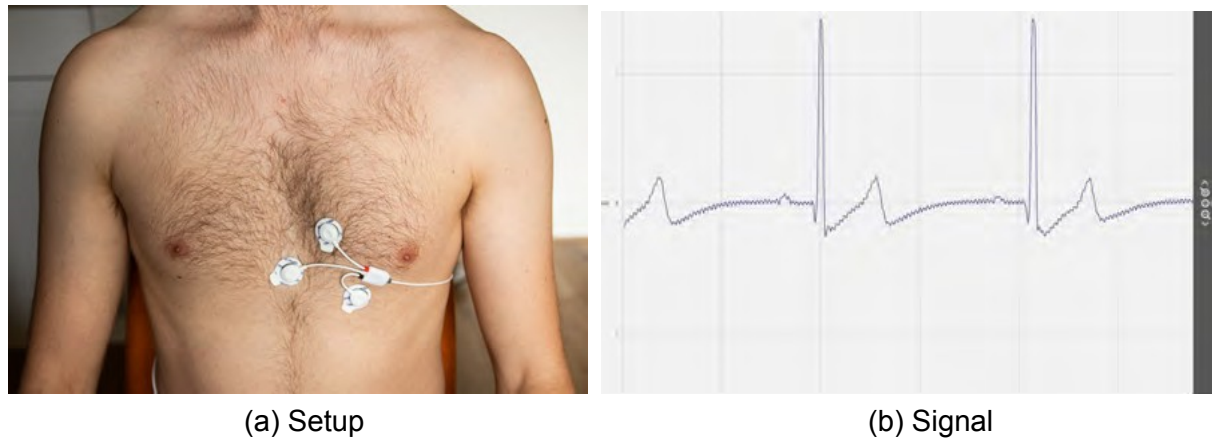


Figure 6: Setup and signal resulting from the wet, reference electrode implementation. As expected, the signal quality is on par and all ECG waves can be detected.

During this experiment, the electrodes were not pressed against the skin besides the pressure generated by the thermal shirt. One can easily see that the signal is unusable. None of the ECG waves are present.





Figure 7: Setup and signal resulting from the dry electrode implementation. The signal is unusable as no ECG wave is present in the signal.

After the disappointing results of the dry electrode implementation, we decided to see what would happen if the pressure on the electrodes was increased. Theoretically, the amount of surface contact of the electrodes on the skin would improve. Reducing the amount of air between the skin and the electrodes decreases R_i from Figure 1. Hence, signal quality would be expected to improve.

Indeed, pressing hard on the electrodes as depicted in Figure 8 improves ECG signal quality immensely beyond the point it is almost indistinguishable with the wet electrode ECG signal. The only difference being in the detection of the small P-wave. This wave can be detected in some heartbeats but is absent in others. While pressing, we noticed that the reference electrode (white) was floating above the skin. The cavity created by the chest muscles prevents the reference electrode from touching the sternum.

Next, we tried to improve skin-electrode contact by removing the blue plastic stabilization rings (see Figure 2b). Without these rings, the hypothesis is that surface contact is improved, resulting in better signal quality. Unfortunately, Figure 9 shows a different result. In contrast to the result with the stabilization ring, we got a signal that largely consists of a high-frequency, large-amplitude noise component. The standard ECG waves are unrecognizable. There are, however, hints of periodicity.


Based on the additional pressure experiment findings, we hypothesized that the signal could be heavily distorted because the reference electrode was floating above the sternum. Therefore, instead of pressing the electrode on the sternum, we tried another ground electrode placement by simply holding the reference electrode between the right hand's thumb and index finger.



(a) Setup

(b) Signal


Figure 8: Pressing on the dry electrodes improves signal quality immensely. QRS and T waves can be recognized easily. The P-wave is more difficult to see.



(a) Setup

(b) Signal

Figure 9: Removing the blue plastic stabilization ring around each electrode plate could potentially increase skin-electrode contact. The resulting signal indicates skin-electrode contact is improved as we now measure something in contrast to our first dry electrode experiment. However, the signal is nothing like a usable ECG signal.



(a) Setup

(b) Signal

Figure 10: Instead of pressing on all electrodes, we just held the reference electrode between the right hand's thumb and index finger to assure the reference electrode was attached to the body. The resulting signal is surprisingly workable.

Indeed, the signal quality can be improved significantly if the reference electrode can adequately do its job. Although signal quality is lower than the wet electrode or additional pressure signals, one could imagine obtaining a usable signal by filtering out the high-frequency powerline interference. In the resulting signal, R-peaks and even T-waves could be detected, albeit with a decent effort.

2.1.4 Conclusion

In this research report, we investigated the primary challenges of implementing dry electrodes in textile-based applications. As these dry electrodes remove the need for gel and sticky membranes, they can easily be implemented in t-shirts and used unobtrusively. However, as gel often ensures the much-needed skin-electrode conductivity, removing it has consequences. One of these consequences is that dry electrodes need a lot of external pressure to maintain continuous contact with the skin and prevent movement across the skin. Using only commercially available, unmodified hardware, this experiment explored the possibility of obtaining usable biosignals using a tight-fitted thermal shirt in combination with ungelled Ag/AgCl electrodes.

First, it should be noted that electrode placement did not affect signal quality as it was verified by the reference, wet electrode signal. This signal was of decent quality and, hence, confirmed the electrode placement used throughout the experiment. However, a factor that did influence signal quality was the amount and direction of pressure on the electrodes. As dry electrodes do not have gel to maintain conductivity, they rely on direct skin contact via pressure. A tight-fitted thermal shirt does provide a part of the needed external pressure. However, as was confirmed by removing the blue plastic stabilization rings, this pressure must be applied as perpendicularly as possible. Unfortunately, having wires hanging from the electrodes introduces additional shearing force, which changes the overall pressure direction. Hence, removing the non-conductive rings improved conductivity as the pressure did not have to be perpendicular anymore.

Furthermore, the experiment learned that significant additional pressure was needed only for electrodes that cannot be pressed against the skin by the thermal shirt alone. For example, the sternum electrode hovered a few centimeters above the skin because of the cavity that was created by the muscles in the chest. No matter how tight the shirt, these sorts of cavities are unavoidable. Therefore, avoiding these places at all is the only solution for maintaining good skin-electrode contact. Hence, another lead configuration should be chosen. Following the Einthoven principle [1], one should measure Lead I (i.e., horizontally across the body) to obtain a decent estimation of the QRS vector's magnitude. Further, as we are primarily interested in QRS complex detection, electrodes do not have to be close to the heart. Atrial depolarization causes an electrical field large enough to be measured across multiple places across the torso. Therefore, further research regarding alternative electrode positions could be interesting.

Moreover, in the scope of IM-TWIN, answering questions regarding optimal electrode placement concerning body type and size would also be interesting. Young children grow quickly in their first few years of development. Consequently, corresponding electrical properties of their bodies tend to change profoundly as well. Providing optimal electrode placement based on age, gender, and body type could, therefore, be very beneficial to ECG measurements in the IM-TWIN project. Furthermore, if we take this one step further, optimal electrode placement can also be calculated considering the dynamics of the textile movement across the body during posture changes. Both factors can then be combined by minimizing textile movement while maximizing electrical field potential and, hence, reaching optimal real-world ECG signal quality.

References

[1] John T Cacioppo, Louis G Tassinary, and Gary Berntson. *Handbook of psychophysiology* (4th ed.). Cambridge, UK: Cambridge University Press, 2017.

[2] Alper Cömert, Markku Honkala, and Jari Hyttinen. Effect of pressure and padding on motion artifacts of textile electrodes. *Biomedical engineering online*, 12(1):1–18, 2013.

2.2. Signal Quality Indicator (SQI)


The acquisition of quality physiological signals is essential for the effective functioning of the IM-TWIN system. Without proper data streams, the classification of affective states is not possible. However, physiological data streams acquired with the use of dry non-adhesive electrodes are vulnerable to numerous sources of noise, as discussed in the previous section. As some of these sources of noise are intrinsically linked to components of the IM-TWIN system (e.g. dry electrodes, comfortable T-Shirt), it is not possible to eliminate them completely. Consequently, the aim should not be the complete removal of noise, but instead to reduce it to an acceptable level, thus elevating the Signal-to-Noise Ratio (SNR) above a predefined quality threshold. To facilitate the development of algorithms for this purpose, we have designed the Signal Quality Indicator (SQI) to objectively assess the quality and usability of the signals.

Signal quality is a vague concept that can be evaluated in many ways, both qualitatively and quantitatively [1]. Qualitative analysis typically involves simply observing the signal, which is both quick and subjective. To this end, we sought to evaluate the quality quantitatively, if possible in an automated fashion, thereby allowing an unbiased, controlled analysis of the signal. Furthermore, improvements to the signal quality can be easily tracked and monitored over time. To achieve this goal, we have designed two signal quality indicators: i) a signal loss and ii) signal usability indicator.


2.2.1. Signal loss indicator

Body movements and other actions stress the skin-electrode contact by inducing shearing forces upon the electrode. If an electrode loses contact with the skin, no signal can be recorded anymore as the air between the electrode and the skin will act as a massive insulator preventing any electrophysiological signal from reaching the electrode. Consequently, signal loss is experienced as seen in the previous section. Unfortunately, it is not trivial to detect signal loss. Losing skin-electrode contact causes chaotic behavior which is difficult to distinguish reliably from noisy but usable ECG. We tackle this issue by identifying four types of signal loss:

1. **Signal capping** When an analog signal reaches the digital boundaries, it is automatically limited, losing the information above the limit.




2. **Flatlining** If a signal contains no information it flatlines. This type of signal loss is easily recognized.



3. **Powerful high-frequency noise** A very low signal-to-noise ratio. If a signal carries no information it often flatlines or amplifies its background noise.



4. **Low-amplitude** Sometimes, a signal does flatline but still hovers around due to tiny background noise fluctuations which are not strong enough to cause a type 3 signal loss.



Each type of signal loss has its own indicator algorithm. In the end, the masks of all four indicators are merged resulting in a final signal loss mask.

2.2.2. Signal usability indicator

Even if the signal loss indicator approves a signal, it could still be contaminated with motion artifacts. The signal loss indicator only checks if there is no information lost. It does not check whether information is buried under noise. Therefore, a second indicator has been developed: the signal usability indicator. It checks whether heartbeat detection can be performed reliably enough for heart rate analysis.

First, it performs heartbeat detection using the popular Pan-Tompkins algorithm on the remaining segments of data after signal loss removal. This algorithm developed in 1985 detects QRS peaks in ECG data and is still among the most used ECG algorithms today. Next, the interbeat intervals (IBIs) between two successive heartbeats are calculated. Missing a heartbeat causes an abnormally large IBI. Vice versa, detecting a false beat creates an unusually short IBI (see peaks in Figure 11).

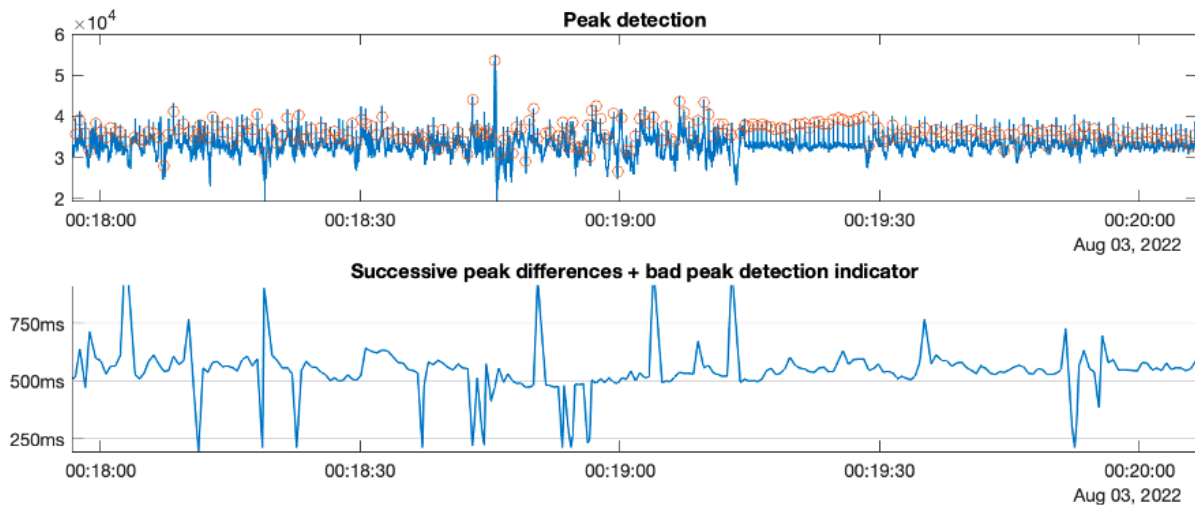


Figure 11: The two phases of usability assessment. First, ECG peak candidates are extracted from the incoming ECG signal using the Pan-Tompkins algorithm. The number of missing and falsely detected peaks is then assessed by calculating the successive peak differences. Missing a heartbeat causes an abnormally large difference. Vice versa, detecting a false beat creates an unusually short difference.

We use the number of misclassified beats as a quantitative metric for signal usability. Missing many beats or detecting many false beats makes it increasingly difficult to perform proper heart rate variability analysis. However, short peaks (i.e., misclassifying only one beat) are not a problem if scattered throughout the signal. Hence, we filter out short peaks by performing a median filter of width 5. By subtracting the median filtered signal from the original signal, we obtain a residue containing streaks of misclassified beats. Taking the residue's sum over 15s windows allows us to use a threshold of 1.5 seconds to classify 15s regions as usable or unusable. The result is a binary mask that's merged with the signal loss mask to obtain the final signal usability mask.

2.2.3. Results

We evaluated several pilot data samples with the SQI. The summary of one of these assessments is plotted in the figure below. The masks of all five quality indicators (the four signal loss indicators and the signal usability indicator) are merged into a binary 'good' and 'bad' signal mask represented as 'green' and 'red' regions, respectively. A 'good' signal provides good input for HR and HRV analysis. In contrast, HR and HRV measures obtained from 'Red' regions are unpredictable and unusable for affect classification.

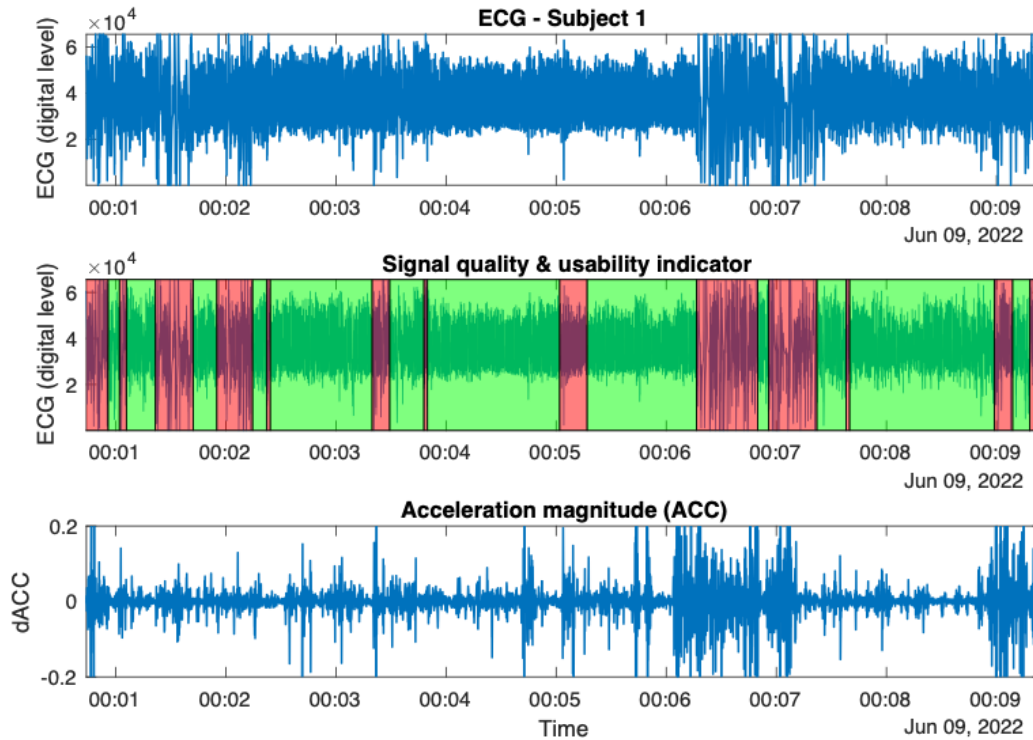


Figure 12: Signal quality assessment results on pilot data from one subject. The signal quality indicators are fused to one binary output signal that highlights ‘good’ and ‘bad’ signal segments as green and red shaded regions.

As seen in Figure 12, body movements have a detrimental effect on the quality of signals, and despite attempts to remove this effect, it is impossible to do so entirely. However, the effect of body movements can be reduced by improving the fit of the T-Shirt. An ideal fitting T-Shirt ensures a less harmful effect on the quality of the signals. Unfortunately, obtaining a good fit on very young children with Autism Spectrum Disorder (ASD) can be quite challenging. To address this issue, University of Utrecht (UU) has developed a real-time Signal Quality Indicator tool that can help practitioners get a perfect T-Shirt fit. The tool is intended to be used in combination with a startup protocol that tests the T-Shirt fit prior to the commencement of the measurement session.

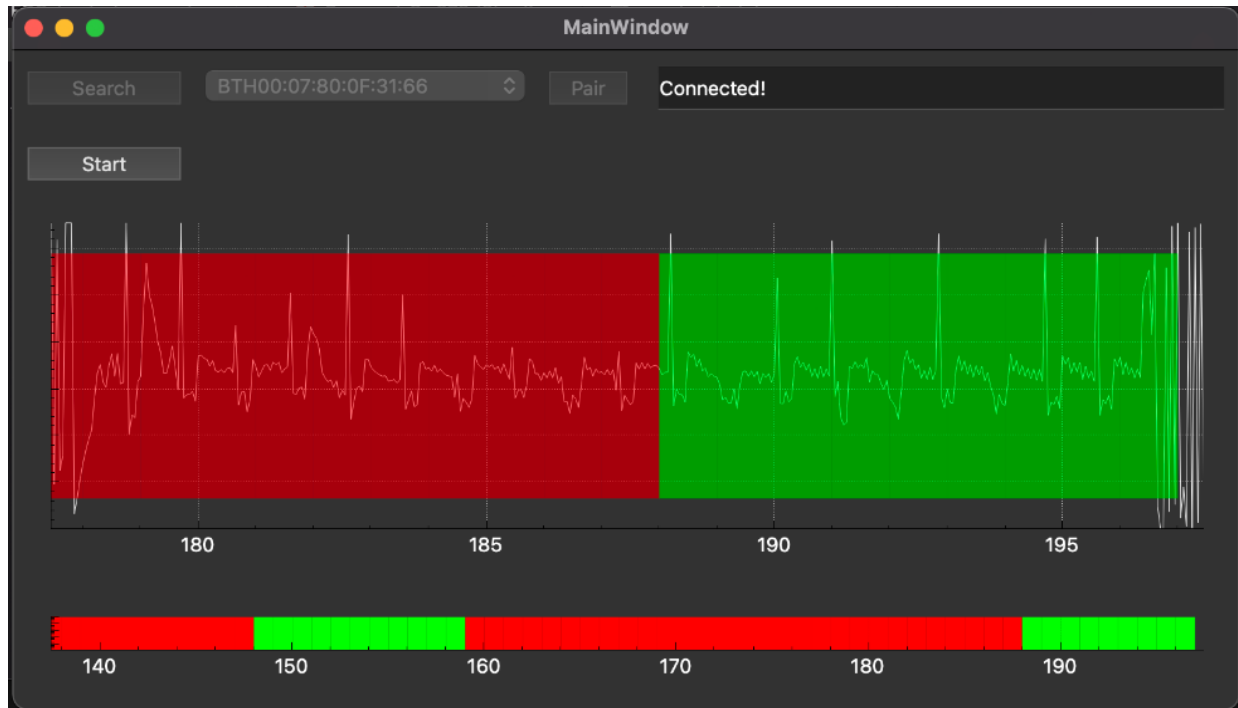


Figure 13: The real-time Signal Quality Indicator provides easy-to-interpret feedback about signal quality. When used in combination with properly defined startup protocols, signal quality can be assured in every measurement session.

2.2.4. Conclusion

In conclusion, the Signal Quality Indicator (SQI) provides a reliable and objective means of assessing the quality of physiological signals acquired with the IM-TWIN system. It identifies four types of signal loss and assesses signal usability by detecting missing and false heartbeats. By combining this information, the SQI creates a binary output signal that highlights 'good' and 'bad' signal segments. The real-time SQI tool is a valuable tool for practitioners attempting to obtain a perfect T-Shirt fit, as it provides real-time feedback about signal quality. As such, it can help ensure that the quality of the acquired signals is satisfactory for affective state classification. Furthermore, it aids development by allowing for easy tracking and monitoring of signal quality over time.

2.3 Processing non-stationary signals and open-source repository

The IM-TWIN system aims to classify the emotional state of a participant through analysis of various electrophysiological signals. These signals, originating from complex and dynamic systems such as the cardiovascular, respiratory, and integumentary systems, pose a challenge due to their inherent chaotic nature [1]. In raw form, these signals are often difficult to use as they are a result of multiple systems that are not always correlated to emotional state. Furthermore, they are easily distorted due to their low power levels. Therefore, it is crucial to separate meaningful patterns from background noise. In many physiological signals, meaningful patterns are defined as repeated patterns (see Figure 14). In the cardiovascular system, repeating patterns are correlated with heart rate, which is the repeated contraction of the heart muscle and serves as an important indicator of stress [1]. Repeating patterns in the respiratory system are created by the repeated in and exhalation of air by the lungs that are influenced by the emotional state of the body as well [1]. As such, the analysis of repeated patterns (e.g., frequencies) forms the basis of physiological signal processing, and hence, the basis of IM-TWIN's affective state classification.




Figure 14: In physiological signals, analysis of repeated patterns give an insight into the complex interplay between various systems that let the body function.

Frequency analysis of physiological signals is difficult, however. The non-stationary (e.g., dynamic signal characteristics) nature of these signals poses a challenge for traditional frequency analysis techniques such as Fourier analysis that assume stationary signals. In stationary signals, frequencies do not change over time (see Figure 15) [2]. However, in physiological signals this is not the case at all. In fact, the majority of valuable information is in the change of these frequencies over time. We are more interested in the change of heart rate than the actual heart rate itself. As such, using traditional frequency analysis on our non-stationary physiological signals is not possible [3,4].

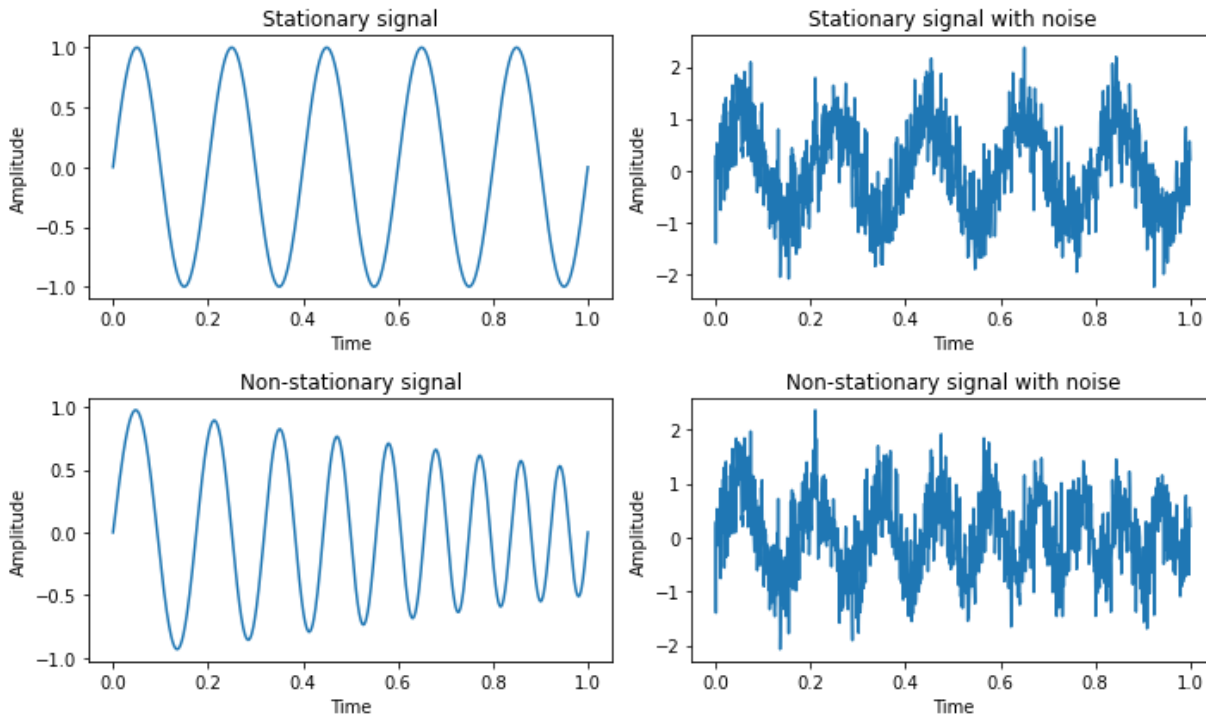


Figure 15: A comparison of stationary and non-stationary signals with and without added noise. Stationary signals exhibit characteristics that remain stationary in time, hence their name. Non-stationary signals, on the other hand, have dynamic characteristics which makes them harder to process. Especially in the presence of noise.

Fortunately, advanced algorithms such as the short-term Fourier transform (STFT) and the wavelet transform (WT) have been developed to overcome this problem. As the name implies, the STFT splits the signal into smaller, short segments and performs Fourier analysis on each segment. Although STFT adds the ability of tracking frequencies over time, this approach comes with the drawback of chopping up the signal [3]. This scrambles the original characteristics of the signal which makes frequency analysis not always accurate (see Figure 16). Hence, STFT is not an option for IM-TWIN. The WT, on the other hand, solves this problem by using base functions that can dilate and contract with frequency to represent the signal. As this lets the original signal unchanged, an unbiased high resolution across the entire frequency spectrum can be obtained [4].

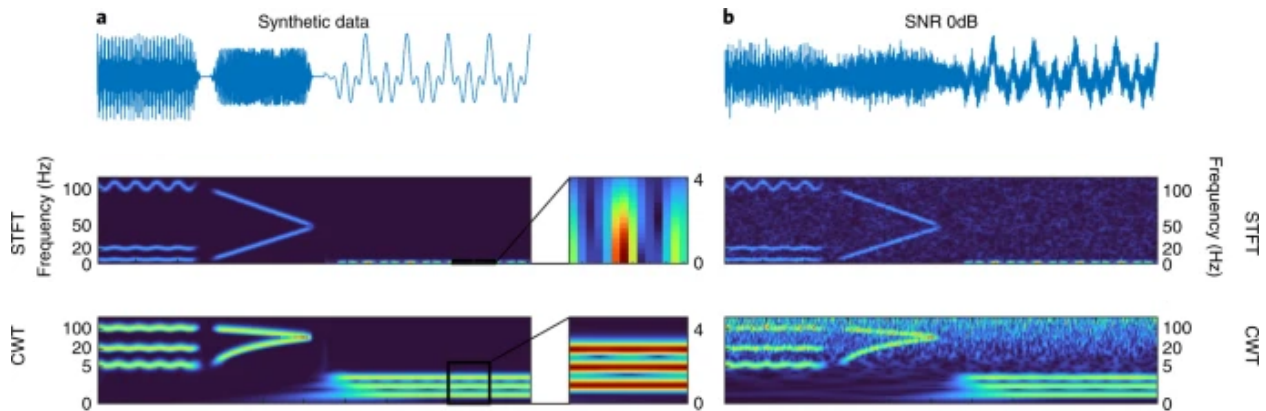


Figure 16: The Continuous Wavelet Transform (CWT) has much better resolution across a wider frequency range than the Short-term Fourier Transform (STFT). This is especially visible in the low frequency range where STFT cannot separate the three low frequency components. CWT can easily do this while still maintaining high resolution in the high frequency range.

The high, wide-band resolution of the WT, makes it well-suited for physiological signal analysis. However, high, wide-band resolution requires the use of many base functions, and hence, many calculations. Consequently, its high computational load limits its use in real-time applications [5]. To address this challenge, UU developed a fast implementation of CWT, which has proven to be 34 to 122 times faster than existing implementations. With the development of fast CWT (fCWT), a strong foundation has been created for the real-time IM-TWIN signal processing pipeline.

The fast Continuous Wavelet Transform (fCWT) developed by UU has gained significant traction in the field of physiological signal processing. The results of benchmarking fCWT against eight competitor implementations on three challenging datasets were published in a peer-reviewed article in Nature Computational Science and received positive feedback. fCWT has gained popularity among clinical, biomedical and computer scientists as it is maintained in an easy-to-access open-source repository (see Figure 17). Additionally, to make fCWT easily accessible, three interfaces to use fCWT in other languages (Python, Matlab and C++) have been created to reduce the need for specialized knowledge on wavelet theory or low-level programming. Besides being very usable for the next steps of IM-TWIN, the availability of these interfaces now allows researchers from various backgrounds to utilize the power of fCWT in their real-time applications, making the power of CWT accessible to a wider audience for the first time (see Figure 18).

fastlib / fCWT Public

Pin Unwatch 5 Fork 20 Star 78

<> Code Issues 4 Pull requests Actions Projects Wiki Security Insights Settings

main 1 branch 0 tags Go to file Add file Code

File	Description	Commit Date
fastlib Remove c0, is not needed here		f33a6de 2 weeks ago 88 commits
MATLAB	Matlab fix	2 weeks ago
img	Readme update	3 weeks ago
libs	Fixed small bugs in Python and Matlab package	2 weeks ago
src	Remove c0, is not needed here	2 weeks ago
tests	Added complex valued input	3 weeks ago
.DS_Store	Added example to Readme.md	last year
CITATION.cff	Update Readme	3 weeks ago
CMakeLists.txt	Matlab fix	2 weeks ago
LICENSE.txt	Initial commit	last year
README.md	Added benchmark notebook	2 weeks ago
benchmark.ipynb	Added benchmark notebook	2 weeks ago

About

The fast Continuous Wavelet Transform (fCWT) is a library for fast calculation of CWT.

c fast cpp optimization signal-processing matlab wavelet-transform cwt time-frequency-analysis

Readme Apache-2.0 license Cite this repository 78 stars 5 watching 20 forks

Releases

Figure 17: Open-source GitHub repository has gained significant popularity since its release with a total of 78 stars and 20 forks as of February 4th, 2023.

fCWT is published as an extended open access in **Nature Computational Science**: <https://www.nature.com/articles/s43588-021-00183-z>. The publication contains a link to supplementary materials, data, etc. Also a dedicated Code Ocean code repository is available via <https://doi.org/10.24433/CO.8389373.v1code>. The article was featured as the issue's cover article. Here, we refrain from providing an additional description of this work. Please consult the article for more information. It is provided with this deliverable.

The fCWT article also received considerable attention. In the next issue of **Nature Computational Science** an article titled Revisiting signal analysis in the big data era was dedicated to our work: <https://www.nature.com/articles/s43588-022-00210-7>. This article underlines the relevance of our work for big data processing and machine learning with all possible data and in all possible contexts. As such it follows the arguments of the fCWT article.

A full page in one of the Netherlands' leading news papers NRC was devoted to fCWT, see: <https://www.nrc.nl/nieuws/2022/02/03/uit-een-hersensignaal-is-veel-meer-detail-te-halen-a4084223>. Moreover, the Utrecht University launched a press release, see: <https://www.uu.nl/nieuws/honderd-keer-snellere-signaalverwerking-met-nieuwe-rekentechniek>. The article has also been made available via open repositories such as ResearchGate (see: <https://www.researchgate.net/publication/358153539> The fast continuous wavelet transformation fCWT for real-time high-quality noise-resistant time-frequency analysis), Academia.edu (see: <https://www.academia.edu/75520883/The-fast-continuous-wavelet-transformation-fCWT-for-real-time-high-quality-noise-resistant-time-frequency-analysis>), and Semantic Scholar (see: <https://www.semanticscholar.org/paper/The-fast-continuous-wavelet-transformation-%28fCWT%29-Arts-Broek/1183d7641948819e5f3370ddc1f7bbcc7f37a3dc>), amongst several others.

At a local conference fCWT was presented via a poster, which also provides an abstract. The poster is provided with this deliverable.



Figure 18: Researchers showed enthusiasm about the release of fCWT to the public. Not only IM-TWIN benefits from a faster, more accurate signal processing pipeline. Projects around the world in various fields can be resurrected and brought back to life with the power of fCWT.

References

1. John T Cacioppo, Louis G Tassinary, and Gary Berntson. *Handbook of psychophysiology* (4th ed.). Cambridge, UK: Cambridge University Press, 2017.
2. Boukouvala, E., Miridakis, N. & Veloni, A. *Digital and Statistical Signal Processing* (CRC Press, 2019).

3. Boashash, B. *Time–Frequency Signal Analysis and Processing: A Comprehensive Reference* 2nd edn (Academic, 2016).
4. Addison, P. S. Introduction to redundancy rules: the Continuous Wavelet Transform comes of age. *Philos. Trans. R. Soc. A Math. Phys. Eng. Sci.* 376, 20170258 (2018).
5. Smeets, H., Ceriotti, M. & Marrón, P. J. Adapting recursive sinusoidal software oscillators for low-power fixed-point processors. *ACM Trans. Embedded Comput. Syst.* 19, 1–26 (2020).

3. Processing of visual information

As shown in a previous deliverable¹, CNR-ISTC developed a system for the automatic detection of facial expressions – *Facial Expression Recognition*-FER algorithm – to be used in conjunction with *PlusMe* interactive toy, during real-time activities involving children. However, exploratory tests showed how the system may not be performing enough in a real scenario, as FER requires very marked facial expressions to work properly; further analysis (including attempts to improve the algorithm and the *machine learning* training) confirmed this limitation.

For this reason CNR-ISTC adapted the system for a new purpose – more addressable from a technical point of view – still relevant for ASD and for neurodevelopmental disorders in general: the automatic detection of eye contact between child and caregiver. This behavioural index is extremely important, both in therapy and during the assessment of diagnosis; the lack of eye contact in children indicates generally an atypical social behaviour, and it is one of the most evident symptoms of ASD.

The gaze detection software allows us to measure reliably and objectively, the amount of eye contact events between child and therapist (namely, *how many times* and *for how long* the child looks at the caregiver’s eyes). It is important to note that this social behaviour is currently evaluated by the researchers who analyse the data, through the manual scoring of the videos of the experimental session, a very time-consuming operation, which often undergoes subjective evaluations.

To make the software usable in a real-therapy scenario, the CNR-ISTC embedded a micro camera² in customised, 3D printed glasses (Figure 19), for the therapist. The camera is then connected to a *Raspberry Pi*, a small portable, self-powered computer, placed in the therapist’s pocket. The *Raspberry* can then, according to the scenario, save the video of the therapy session for subsequent data analysis (in this scenario the system ensures 30 fps³), or broadcast the data to a laptop, via Wi-Fi connection, for real-time data analysis (in this scenario the system ensures between 20 and 25 fps, due to performance reasons).

¹ Deliverable 3.3 “*Plusme AI-augmente behaviour and IM-TWIN 1*”, <https://im-twin.eu/deliverables/>.

² The model used is the CMT-8MP-IMX179-W510 USB endoscope camera module, by <http://camera-module.com/>

³ *Frame Per Seconds, FPS*




Figure 19: *left: the micro USB camera, selected by CNR to be used for eye contact detection; right: the 3D printed glasses embedded in the micro camera; the glasses are intended to be worn by the therapist during the activity with the child.*

The software used for eye contact detection is based on the recent work of Chong (2020)⁴. In this paper, the authors developed a deep convolutional neural network model to automatically detect eye contact in videos. As a dataset, they used 4,339,879 annotated images composed of 103 subjects with different demographic backgrounds of whom 57 subjects have a diagnosis of ASD. The network achieves an overall accuracy of 0.936 and recall of 0.943 on 18 validation subjects, and its performance is equal to that of 10 trained human coders, with an average accuracy of 0.918 and recall of 0.946.

As shown in Figure 20, the performance of the system, tested in the lab, is very reliable, as the gaze detection seems to be very accurate, also with “noisy” images (i.e., images where the human face is neither in the centre of the frame nor in the most of the frame). This result was confirmed by an Inter-Rater Reliability (IRR) test, used to compare the software performance with the human rater scoring. The IRR was based on the response of 3 human raters who analysed, frame by frame, the same video processed by the software; the IRR provided a $k = 80$, $k = 0.83$ and $k = 0.81$ respectively for the first, second and third rater, which means a very good agreement between humans rating and the software performance.

⁴ Chong, E., Clark-Whitney, E., Southerland, A. *et al.* Detection of eye contact with deep neural networks is as accurate as human experts. *Nat Commun* **11**, 6386 (2020). <https://doi.org/10.1038/s41467-020-19712-x>



Figure 20: a test run in the lab, to assess the accuracy of the eye contact detection.

CNR and SAPIENZA run then some pilots (November and December 2022) using the system in an actual therapy scenario; as shown in Figures 21 and 22, the system proved to be both reliable and easy to use for the therapist; moreover, the use of a micro camera, hidden in standard glasses seems to be not noticed by the child⁵, an important aspect to run experiments in an ecological environment.

Finally, it is important to note that, as previously stated, the system can work also in real-time, for a potential use combined with the panda *PlusMe* toy (Figure 23), i.e., to make the toy automatically trigger rewarding patterns, in response to child's positive social behaviour, just like the eye contact toward the caregiver⁶.

⁵ The selected child for the pilot was diagnosed with Language Disorder.

⁶ This potential use could be evaluated in a future pilot, to reinforce the eye contact behaviour.



Figure 21: 3 images showing how the gaze detection system properly detects the eye contact between the child and therapist, who wears the sensorised glasses (on left the environmental camera images, on right the glasses camera images).




Figure 22: 3 images showing how the gaze detection system properly detects the missing eye contact between the child and therapist, who wears the sensorised glasses (on left the environmental camera images, on right the glasses camera images).

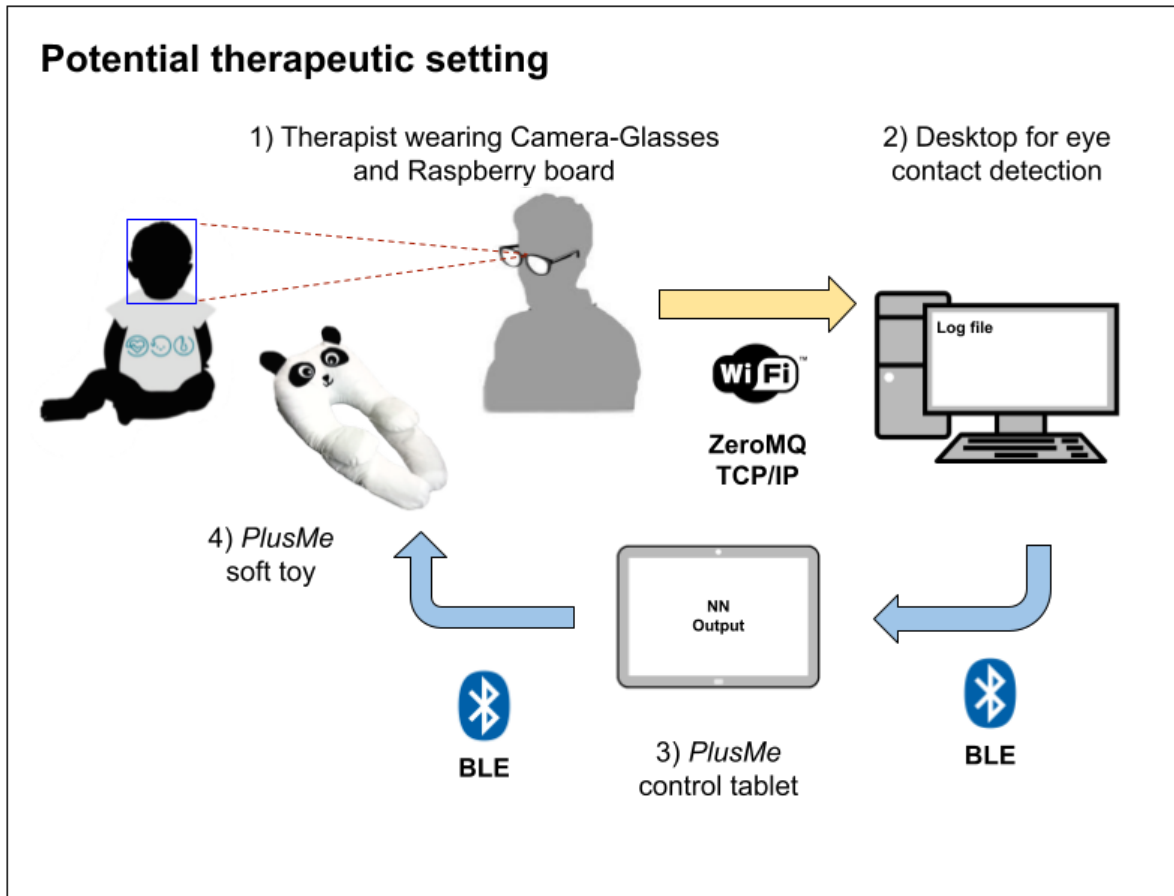


Figure 23: A possible version of the IM-TWIN system, where PlusMe toy emits rewarding responses when the eye contact between child and therapist is detected.

To conclude the section, it must be noted that the automatic detection of a child’s facial expression is a functionality which is postponed for the unsatisfactory FER performance to an actual therapeutic scenario. This is probably caused by, at least in part, the use of a non-appropriate training set (the *Real World Affective State Database* - RAF-DB⁷). However, the new system based on the sensorised glasses, lets to envisage the forthcoming collection of a new dataset, during the next experiments in SAPIENZA and CRI; the new data could then be used for refining the original training set, and then improve the FER performance to acceptable results.

⁷ See <http://www.whdeng.cn/raf/model1.html> for further information.

4. Processing of interaction between child, *PlusMe* and therapist

The physical interaction between child and *PlusMe* is an interesting source of information; if combined with the visual data collected by the sensorised glasses (described in the previous section), it can provide valuable information about the social engagement between child and therapist (see Figure 24).

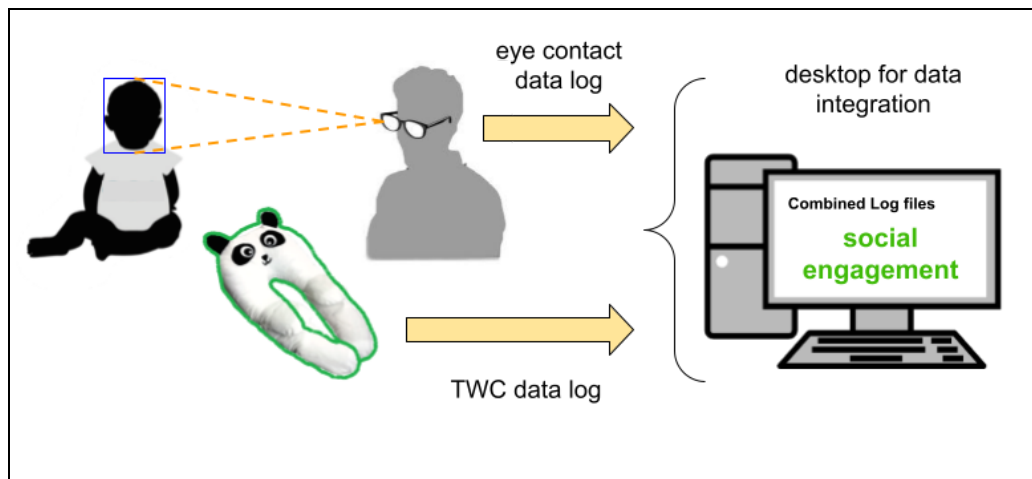


Figure 24: Schema for data integration from the sensorised glasses and a TWC toy.

In this regard, CNR-ISTC started to develop a software framework for integrated data collection – partially based on the main IM-TWIN architecture⁸ – using the log files provided by the glasses for eye contact detection, and by a generic *Transitional Wearable Companion* (TWC) toy.

The first test was run with the octopus *X-8*, a new TWC prototype specifically designed for *turn-taking* games⁹. The toy presents some improvements in respect to the panda *PlusMe*; in particular *X-8*, thanks to additional magnetic sensors, can detect the identity of the user, i.e. it can distinguish between the child’s and the caregiver’s touches. This information is particularly relevant as it allows a better characterisation of the physical interactions detected by the toy embedded electronics.

As shown in Figure 25, the system allows the combination of the logs data (standard XML files) by *X-8* and by the sensorised glasses, through the production of a new “augmented” log file, where all information is synchronised.

⁸ The IM-TWIN system is partially described in the deliverable 3.3 “PlusMe AI-augmented behaviour and IM-TWIN 1”, available at <https://im-twin.eu/deliverables/>

⁹ See videos at page <https://im-twin.eu/video/> and publication “*X-8: An Experimental Interactive Toy to Support Turn-Taking Games in Children with Autism Spectrum Disorders*”.

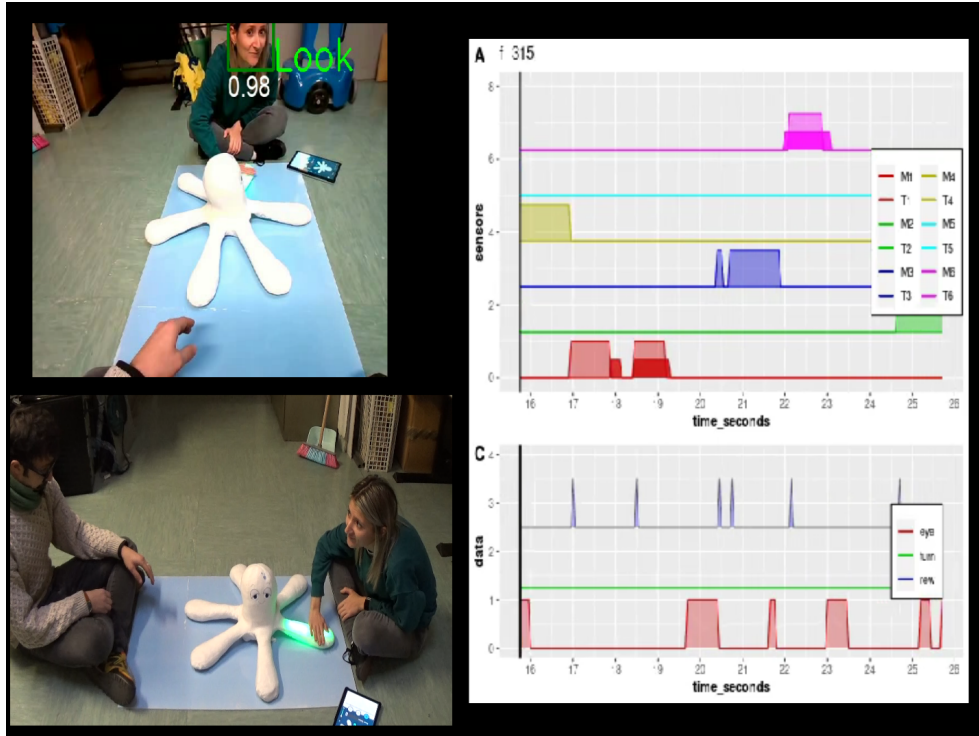


Figure 25 First feasibility test for data integration; the two researchers play the roles of child and caregiver. **Top left:** glasses camera view, for eye contact detection; **Bottom left:** environmental camera view; **Top right:** log of data collected by X-8; the graph 'A' shows the 6 tentacles activation values (one colour per tentacle), where higher activation is due to child's touch, and lower activation to caregiver's touch (the plot shows 10 seconds). **Bottom right:** log of data collected by the glasses: the graph 'C' shows the child's eye contact detection (red line), synchronised with the X-8 log, shown in graph 'A'.

The new log file contains information about the course of the experimental session, sampled at 20 Hz (X-8 saves data each 50 ms), for example:

- status of the X-8 tentacles activations (touched / not touched), including the identity of the user who touches the toy (child / caregiver);
- triggering of toy sensory reward (lights and sounds), including the type of reward;
- type of game, currently selected by the caregiver (in X-8, 3 different games with increasing complexity are available¹⁰);
- if the rule of the selected game requires a turn, the specification of the game round (e.g.: 0: game with no turn; 1: child's turn; 2: caregiver's turn);
- eye contact detection, between child and caregiver.

During the first lab test, the 2 researchers simulated a 8 minute-long session, where the 3 games available in X-8 games were tested. Then, the main log file was inspected, matching it to the videos collected by the environmental and glasses cameras; the analysis showed that the information is synchronised with the events in the experimental session, confirming that the data

¹⁰ In *PlusMe*, the available games are currently 6.

is reliable and usable for further statistical analysis.

Importantly, the statistical processing of the log could reveal also relevant social events, e.g. looking at the temporal sequences of occurrences: for example, observing the sequence of child's touch, production of reward by the toy, and shifting of the child's attentional focus from the toy to the therapist, it could be possible to quantitatively estimate the joint attention.

This new software component is an important improvement for the IM-TWIN system: the automated analysis of the log can in fact release the human rater from an heavy, time-demanding task, namely the visual inspection of the videos of the experimental sessions, necessary to collect the child behavioural data. In other words, the researcher could have the results available immediately, after the test with the child. Moreover the results, being objective, would no longer require additional human raters to assert the accuracy of observed data (for example through the *Inter Rater Reliability* – IRR test).

As shown in the previous Figures 21 and 22, ISTC-CNR and SAPIENZA researchers are now running some pilot tests with children, to refine the system and to start collecting data in real scenarios. Once the system is verified, the software improvements developed in X-8 will be tested also with the panda *PlusMe* toy. Moreover, when the physiological data collected by the sensorised t-shirts will be available, it will be included in the main log file, to produce a complete evaluation of the social and emotional engagement of the child, during the experimental session.

5. Conclusions and future developments

The presented work in this report has contributed to the development of the IM-TWIN affective classification system as well as to the research community as a whole. Initial experiments with dry-electrodes and the PLUX hardware system turned out to be very fruitful as it emphasised the difficulty of obtaining high-quality physiological signals in noisy real-world environments once again. In response to this challenge, UU developed a two-faceted approach consisting of the guidance of practitioners in obtaining good quality data streams during the experiments using a real-time easy-to-understand Signal Quality Indicator and by creating fCWT; a fast, general, and noise-resilient algorithm that can handle all sorts of challenging data streams. fCWT was validated on highly challenging physiological datasets. By tackling the challenge of noisy data at both the input and processing stages of IM-TWIN, a solid foundation for further development has been created.

Next steps involve the development of processing pipelines specific to the physiological signals used in IM-TWIN. More specifically, pipelines for the ElectroCardioGram and ElectroDermal Activity data streams. Only decent quality ECG and EDA processed with high-resolution pipelines will lead to accurate and robust affective state classification in the everyday therapeutic situations that IM-TWIN will face.

Concerning the integration of information from different sources, in this report we described how the data collected by the sensorised glasses for eye contact detection, and the data about

the physical interaction detected by a generic TWC toy, can be combined to produce a reliable, accurate description of the social interaction between child and therapist. This aspect, currently tested in a lab proof of concept trial, will be tested in a real scenario, during a therapeutic session with children with neurodevelopmental disorders.

An update of the whole system will be provided in the next version of the current document, the deliverable 2.2 "*Processing of physiological signals, visual information and PlusMe interaction, second version*", due in July 2023.

Background

- Signals (i.e., data streams) are everywhere
- All **signals** require **processing**
- Frequency-analysis reveals underlying dynamics
- The **Fourier transformation (FT)** is often used
- **FT** is fast but **ignores time**

Introduction


- The **Continuous Wavelet Transform (CWT)** translates signals to time-frequency domain
- **CWT** is accurate but **too slow**
- We present the **fast CWT (fCWT)**, both fast and accurate!

How does a CWT work?


Wavelets stretch and shrink to measure different frequencies and move to measure the position of these frequencies using convolution.



Time-Amplitude




Time-Frequency spectrum



Results

Algorithm

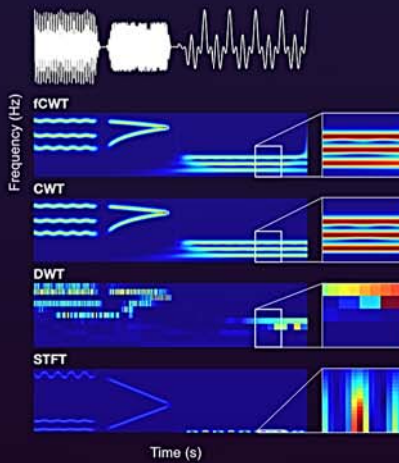


Benchmark against other CWT implementations



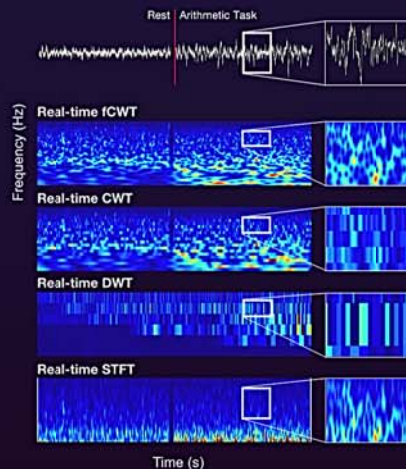
Synthetic data

Superb and equal accuracy when compared to CWT and other methods.



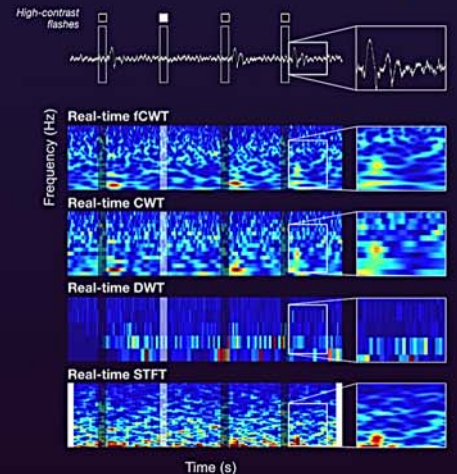
Electroencephalography (EEG)

Real-time, high-resolution analysis of noisy and non-stationary EEG



High-density in vivo analysis

Real-time, simultaneous analysis of electrodes in the visual cortex



Discussion

- **Exact calculation** instead of approximations
- Performance is **wavelet-type independent**
- CWT underlies **many time-frequency methods**
- **Inherits both CWT's advantages and limitations**

Conclusion

- **fCWT allows acceleration** in science & engineering, industry and health
- Bring offline research to **real-time practice**
- Source-code is published **open-source**

Output

- Arts, L.P.A., van den Broek, E.L. The fast continuous wavelet transformation (fCWT) for real-time, high-quality, noise-resistant time-frequency analysis. *Nature Comput Sci* (2022). DOI: 10.1038/s43568-021-00183-z
- Arts, L. P. A. & van den Broek, E. L. The fast continuous wavelet transformation (fCWT): real-time, high quality, noise-resistant, time-frequency analysis (Code Ocean, 2021); DOI: 10.24433/CO.6389373.v1
- Arts, L. P. A. & van den Broek, E. L. (2022) The fast continuous wavelet transformation (fCWT): real-time, high quality, noise-resistant, time-frequency analysis [Source code] <https://github.com/fastlib/fcwt>



OPEN

The fast continuous wavelet transformation (fCWT) for real-time, high-quality, noise-resistant time–frequency analysis

Lukas P. A. Arts and Egon. L. van den Broek

The spectral analysis of signals is currently either dominated by the speed–accuracy trade-off or ignores a signal's often non-stationary character. Here we introduce an open-source algorithm to calculate the fast continuous wavelet transform (fCWT). The parallel environment of fCWT separates scale-independent and scale-dependent operations, while utilizing optimized fast Fourier transforms that exploit downsampled wavelets. fCWT is benchmarked for speed against eight competitive algorithms, tested on noise resistance and validated on synthetic electroencephalography and in vivo extracellular local field potential data. fCWT is shown to have the accuracy of CWT, to have 100 times higher spectral resolution than algorithms equal in speed, to be 122 times and 34 times faster than the reference and fastest state-of-the-art implementations and we demonstrate its real-time performance, as confirmed by the real-time analysis ratio. fCWT provides an improved balance between speed and accuracy, which enables real-time, wide-band, high-quality, time–frequency analysis of non-stationary noisy signals.

Signals are essential in both nature and (man-made) technology, because they enable communication^{1,2} (Fig. 1). Mathematically, a signal is a function of one (for example, speech) or more (for example, a two-dimensional (2D) image) dimensions that carries information about the properties (for example, state) of a physical system³. A source transmits a signal via a channel to a receiver, which delivers it to its destination. For example, a brain sends an oral message via vocal cords through the air, which is received by the listener's ear, which brings it to the listener's brain. When the same message is transmitted via a smartphone, the air is complemented by a chain of technology, leaving the rest of the chain untouched. Signals are omnipresent in society^{3,4} (Fig. 1).

Independent of its source, a signal needs to be processed to enable the generation, transformation, extraction and interpretation of the information it is carrying³. A widely used method to interpret (that is, extract and analyze) repeating patterns in signals is the Fourier transform (FT)^{3,4}. A FT transforms a function of time into a complex-valued function of frequency, representing the magnitudes of the frequencies. The FT assumes the signal is stationary. In other words, it is a stochastic process in which the marginal and joint density functions do not depend on the choice of time origin². However, in real-world practice, this assumption is often violated. Consequently, the FT is unable to process real-world non-stationary signals reliably⁵. To circumvent the problem of non-stationarity, advanced algorithms exist that analyze a signal based on their decomposition in elementary signals that are well localized (or boxed) in time and frequency⁴. These include the short-term Fourier transform (STFT), also known as the Gabor transform, and the wavelet transform (WT)⁶.

The STFT is very similar to the FT, but it uses a window function and short wavelets localized in both time and frequency, instead of pure waves, to extract temporal and spectral information. The drawback of the STFT is its use of a fixed-width window function, as a result of which frequency analysis is restricted to frequencies with a wavelength close to the window width⁷. Additionally, chopping up the signal in short, fixed-width windows scrambles the signal's properties. Accordingly, the frequency analysis is affected⁸.

The WT overcomes the drawback of the STFT by not relying on a window function. Instead, it uses a family of base functions that dilate and contract with frequency to represent the signal, thereby ensuring high resolution across the entire frequency spectrum. Consequently, the WT suffers from a high computational load. This prohibits its use with low-end hardware and for real-time applications⁹, as real-time computation requires an algorithmic computation time that is smaller than the signal's duration.

To reduce the computational burden of the WT, the discrete wavelet transform (DWT) has been proposed, which applies a coarse, logarithmic discretization. This makes DWT suitable for data compression, but simultaneously disqualifies it from use in detailed analysis, as it is not able to analyze intricate time–frequency details⁸ (as shown in Fig. 2). For this, a true WT—the computationally expensive continuous wavelet transform (CWT)—also called an integral wavelet transform (IWT), is needed. CWT offers a high-resolution representation of the time–frequency domain by using near-continuous discretization. Its continuous time and frequency scales better support intricate time–frequency analysis. Consequently, CWT is often described as the mathematical microscope of data analysis¹⁰ (Fig. 2).

In this Resource paper we introduce the open-source fast continuous wavelet transform (fCWT), which brings real-time, high-resolution CWT to real-world practice (for example, biosignals^{11–13}, cybersecurity^{14,15} and renewable energy management^{16,17}; Fig. 1). Next, we assess the performance of fCWT in a benchmark study and then validate the use of fCWT on synthetic, electroencephalography (EEG) and in vivo electrophysiological data. We end with a concise discussion.

Results

The performance of fCWT was benchmarked against six widely used CWT implementations, then it was subjected to a threefold validation on accuracy, resolution and throughput using, respectively, synthetic data, human EEG data and high-density in vivo extracellular rodent electrophysiology.

	Science and engineering	Industry	Health
Applications	Cybersecurity, seismology, microscope focusing, spectral analysis, signal compression and enhancement, gravity wave detection, mass spectrometry...	Oil and mineral prospecting, process monitoring and control, renewable energy management, leakage detection, circuit testing...	Diagnostic imaging, brain-computer interfaces (BCI), heart disease diagnosis, medical image classification...
Data type	Video, audio, imagery, accelerometer, fingerprints, ultrasonography, laser interferometer, various sensory data...	Ground-penetrating radar, sonar, audio, ultrasonic, power consumption, various sensory data...	Electrocardiography (ECG), electroencephalography (EEG), magnetic resonance imaging (MRI), ultrasonography...
fCWT's advantage	Real-time high-resolution signal processing (e.g. classical music), stress reduction via enhanced noise canceling.	Reducing risk by real-time machine monitoring and increasing prospecting efficiency 34–120 x.	Real-time BCI and saving lives by continuous, remote, real-time monitoring of the cardiovascular system.
Example	<p>Hanford, Washington (H1)</p> <p>CWT used in gravity wave detection by the Laser Interferometer Gravitational-Wave Observatory (LIGO)</p>	<p>Colorbar shows power (dB)</p> <p>CWT used to measure ground deformation above gas reservoir</p>	<p>Movement</p> <p>High spectrum power: 17.0 Low spectrum power: 64.0</p> <p>CWT used to measure overall body health remotely during nocturnal body movements</p>

Fig. 1 | The impact of time–frequency analysis across society. In both nature and technology, signals enable communication, and processing techniques such as the CWT (also called IWT) are applied throughout. CWT was the primary processing method used in the Laser Interferometer Gravitational-wave Observatory (LIGO) experiment to detect gravity waves in highly non-stationary gravitational wave data. In industry, CWT has been applied to enhance mineral detection and speech segmentation. CWT also allows the detailed analysis of biosignals such as an electrocardiogram in the medical domain. BCI, brain-computer interface; BPM, beats per minute. Image credits: (left) adapted with permission from ref. ⁸², Caltech/MIT/LIGO Laboratory; (center) adapted from ref. ⁸³ under a CC BY license.




Fig. 2 | Comparison of DWT and CWT. A time-varying pulse signal of a sonar device is analyzed in the range 0–60 kHz using the DWT and the CWT. The DWT uses a coarse time–frequency discretization to favor speed. By contrast, the CWT uses a time-consuming near-continuous discretization of the time and frequency scales to favor resolution.

Benchmark. To benchmark the performance of fCWT we compared fCWT to the six widely used CWT implementations shown in Fig. 3. Because of its widespread use across research, the complex Morlet wavelet ($\sigma=6$) was used to calculate the CWT of three signals, all containing $N=100,000$ samples. The Morlet wavelet is defined as a plane wave modulated by a Gaussian envelope. The parameter σ controls the time–frequency resolution trade-off⁸. The first signal was generated to be non-stationary using a sine wave whose frequency changed linearly from $f_{\text{start}}=1$ Hz to $f_{\text{end}}=7$ Hz. The second and third signals contained uniformly random noise and a

stationary piecewise defined function, respectively. Three different signals were used to prove fCWT’s flexibility and signal independence. Nevertheless, the signal content and wavelet choice are irrelevant to the performance of fCWT (see Methods for details).

All CWT implementations, including fCWT, use a near-continuous frequency scale containing 3,000 frequencies (range, $f_0=1$ Hz to $f_1=32$ Hz), evenly spaced in exponential space. fCWT thus features a high-frequency resolution in the low-frequency spectrum and a lower frequency resolution in the high-frequency spectrum.

PyWavelet¹⁹ and SciPy²⁰ execution times were measured in a Python 3.8.6 environment, using the Timeit library inside the code to exclude compile time. The overhead resulting from the translation between C and Python was removed by estimating the intersection factor of the linear relationship between signal size and execution time. MATLAB v2019b and Mathematica 12.0.0.0 execution times were measured using the program-specific timing functions that measure the exact kernel execution times.

Wavelib²¹ was used as the benchmark’s baseline algorithm as it is the reference CWT C/C++ library⁹, and most microcontrollers are programmed using C/C++. Wavelib²¹ thus serves as a baseline for the reported speed-ups (Fig. 3). The reported execution times were obtained from an eight-core 2.30-GHz central processing unit (CPU) via 100 successive runs, which removed the influence of caching behavior. A 10-s pause between runs was implemented to prevent the CPU from overheating. Outliers that deviated by more than 3 s.d. from the mean were removed. Wavelib and SciPy had three outliers, leaving $N=97$ samples for all algorithms to ensure equal group sizes. A repeated-measures analysis of variance (ANOVA) revealed that the algorithms differed significantly, $F(4,93)=2,474,778.911$, $P\ll 0.001$, $\eta^2=1.000$, where F denotes the ANOVA statistic based on the ratio of mean squares, which indicates the ratio between the explained and unexplained variance or, in other words, the between- and within-group variability. P is the probability that an observed difference occurred by chance, and η^2 indicates the proportion of variance accounted for (that is, a generalization of r/r^2 and R/R^2 in correlation/regression analysis)¹³. Also, all pairwise comparisons were highly significant ($P\ll 0.001$,




Fig. 3 | Benchmarking with fCWT and six state-of-the-art time-frequency methods. **a**, The average speed-up of fCWT and six publicly available implementations after 100 runs on a signal of length $N=100,000$ with accompanying statistics (in seconds). The signal was analyzed using 3,000 frequencies ranging from $f_0=1$ Hz to $f_1=32$ Hz. **b**, The RAR (equation (1)) of fCWT (600 frequencies, $\sigma=6$), the fastest CWT available (PyWavelet’s CWT, 600 frequencies, $\sigma=6$), STFT (500-ms Blackman with 400-ms overlap) and DWT (four-order Debauchie 20 levels) versus sampling frequency on a 10-s synthetic signal. Parameters were chosen to reflect actual usage in real-world applications. Jumps in the performance of fCWT are explained in the Methods.

Bonferroni-corrected), with fCWT being, respectively, 122 times and 34 times faster than the reference Wavelib²¹ and the fastest available algorithm, PyWavelet¹⁹. Figure 3 presents descriptive statistics for all distributions.

The fast running time of fCWT was also compared to two other fast time–frequency estimation algorithms: the STFT and DWT. In this benchmark, STFT uses a Blackman window of 500 ms with 400-ms overlap, and DWT uses 20 dyadic (that is, $a^l=2^l$) scales of Debauchie decomposition. The parameters were chosen to reflect actual usage in real-world applications (Fig. 1). Both algorithms are implemented and benchmarked in MATLAB using the in-program timing functions. CWT implementations use 600 frequencies, evenly spaced in exponential space. Fewer frequencies are used to reduce memory usage.

To assess whether or not the algorithms perform in real time (that is, an algorithmic computation time less than the signal’s duration), we define the real-time analysis ratio (RAR):

$$RAR = \frac{\Delta t_{\text{computation}}}{\Delta t_{\text{signal}}}, \tag{1}$$

with $\Delta t_{\text{computation}}$ and Δt_{signal} being the duration of the computation and signal, respectively. In the case of $RAR > 1$, an algorithm does not operate in real time. In the case of RAR just shy of 1, the algorithm is unlikely to run in real time as the time–frequency calculation is merely one step in a processing pipeline. When $RAR \ll 1$, real-time operation is likely to be achieved or within reach. For all six CWT implementations and two traditional time–frequency techniques (that is, STFT and DWT), Fig. 3b shows RAR versus sampling frequency. The RARs were obtained by averaging 100 successive runs on 10-s signals with varying sampling frequencies (range, $f_{s0}=1$ kHz

to $f_{s1}=200$ kHz). fCWT and CWT used 5-s signals to fit memory constraints. Small fluctuations in RAR are caused by the stochastic nature of benchmarks performed under real-world conditions. It should be noted that the sampling frequency is directly related to the number of samples. Therefore, we test fCWT’s performance for different signal lengths.

STFT and DWT exhibit superior real-time behavior on signals with sampling frequencies up to 200 kHz and beyond. However, they achieve these very high speeds because of their considerable drop in precision, as shown in Fig. 2. Therefore, STFT and DWT are not suitable for wide-band high-resolution time–frequency estimation. In these cases, CWT is favored. However, even the fastest CWT implementation available tends to be extremely slow compared to STFT and DWT. fCWT merges the best of both worlds, yielding real-time behavior on signals with sampling frequencies up to 200 kHz. This has brought CWT’s execution time close to that of STFT and DWT, while having 25 times to 100 times the spectral resolution of DWT throughout the spectral domain. As such, fCWT is a truly competitive real-time, high-resolution alternative for STFT and DWT.

fCWT allows signals with 34 to 122 times the sampling frequency of existing CWT implementations. Figure 3 shows fCWT’s capability of analyzing signals up to 200 kHz in real time, whereas the fastest implementation of CWT fails at $f_s=30$ kHz. Consequently, fCWT enables real-time analysis of high-frequency signal dynamics, as exist in audio (for example, loudspeaker characterization²², full band speech coding²³ and paralinguistic analysis²⁴), biosignals (for example, brain–computer interfaces¹² and peripheral signals such as ECG, electromyography, electrodermal activity and respiration^{11,13}), image and video (for example, distance transforms^{25,26}), sonar and radar^{27,28}, network analysis (for example, renewable

energy management^{16,17} and cybersecurity^{14,15}) and machine fault diagnosis^{29,30} (Fig. 1).

Synthetic data. fCWT's spectral resolution is equal to that of CWT. In contrast to many other CWT optimization studies, we do not compromise precision. To demonstrate this, we compared fCWT to CWT on both clean and noisy synthetic datasets (see Data availability statement for details). Each dataset consists of three wavepackets that validate an algorithm on spectral and temporal resolution and bandwidth size. A noisy dataset was generated to mimic realistic conditions and assess noise resilience.

Quantitative assessment of each algorithm's performance is carried out by calculating the per-wavepacket mean absolute percentage error (MAPE) scores of 100 runs on both datasets between actual frequencies and the time–frequency ridges extracted from the spectra (see Methods for details). The MAPE scores of the clean data are based on one run, as they are completely deterministic. We used a relative error measure to weight errors at all frequencies evenly.

Next to fCWT and CWT, STFT and DWT were also included, allowing us to show the speed–accuracy trade-off that currently dominates the time–frequency landscape. STFT is based on calculating multiple traditional FTs with overlapping fixed-sized windows. The STFT is very fast and efficient as it relies on the fast Fourier transform (FFT). However, the use of fixed-sized windows requires the wavelengths to be close to the window size. Hence, frequency resolution changes drastically over the spectrum, and only a small frequency band can be analyzed at the same time. DWT does not have this drawback. It does not rely on a window function. Similar to CWT, it uses wavelets that dilate and contract with frequency to represent the signal. However, in contrast to CWT, it uses far fewer wavelets to represent the signal. This makes DWT a very fast time–frequency estimator. Finally, to complete the time–frequency landscape and allow a thorough comparison on accuracy, we added the high-resolution Wigner–Ville distribution (WVD)⁴, the advanced Hilbert–Huang transform (HHT)³¹ and the more recent empirical wavelet transform (EWT)³². WVD has the highest time–frequency resolution mathematically possible and HHT and EWT improve the resolution by using a slow but accurate adaptive iterative process to decompose a signal into fundamental functions that are not necessarily sine functions (for example, FFT). Manual tuning obtained the following parameters for optimal time–frequency sharpness. fCWT and CWT use the complex Morlet wavelet ($\sigma=6$) and a frequency scale of 480 frequencies (range, $f_0=0.25$ Hz to $f_1=250$ Hz), evenly spaced in exponential space (cf. the 111Benchmark section). STFT uses a 500-ms Blackman window with 400-ms overlap, DWT uses 11 dyadic (that is, $a^l=2^l$) scales of 15-order Daubechie wavelet decomposition, and WVD does not take parameters. HHT and EWT use a frequency resolution of 0.25 Hz. HHT uses seven intrinsic modes that were extracted using a maximum signal-to-residual ratio of 20 as a stopping criterion. EWT decomposes the signal using a peak threshold of 5%. Outliers that deviated more than 3 s.d. from the mean were removed. The HHT had four outliers, which resulted in $N=96$ for all algorithms to ensure equal group sizes.

Overall, the per-wavepacket MAPE scores differed significantly on both the clean and noisy datasets between the algorithms ($F(6, 90)=112, 243.890, P\ll 0.001, \eta^2=1.000$; Fig. 4). Within each algorithm, the per-wavepacket MAPE scores also differed significantly between each other ($F(2, 94)=399.044, P\ll 0.001, \eta^2=0.895$) However, fCWT and CWT generated similar, low MAPE scores on both the clean and noisy datasets for all three wavepackets. This was confirmed by a correlation analysis per wavepacket, respectively $r(94)=0.996, P<0.001, r(94)=1.000, P<0.001$ and $r(94)=0.997, P<0.001$. The low MAPE scores can be explained by CWT's and fCWT's wavelet convolution, which averages fluctuations of a signal at different scales³³, and its redundancy (that is, wavelets are

not orthogonal at different scales), which reduces noise by canceling out the random signal components³⁴. Hence, both can separate frequency bands and their details across the full frequency range. When compared to the slow CWT, fCWT's accuracy and noise-handling capabilities are not compromised by its highly efficient implementation. Small differences in the time–frequency spectrum can be seen at the edges. However, these differences are caused by MATLAB's mitigation of edge artifacts (2020Implementation of fCWT section in the Methods).

STFT cannot extract details of the lower frequency bands present in the first and third wavepackets. The wavelengths of these waves are too long for the 500-ms window we used, whereas a larger window cannot distinguish the complex non-stationary behavior of the first packet. Nevertheless, STFT shows strong noise-handling capabilities that result from the averaging effect of FFT's inherited convolution. DWT is powerful in denoising, but not suitable for time–frequency analysis. WVD suffers from its well-known artifacts, which are only made worse by the additive noise⁴. HHT and EWT are very good at separating the frequency bands of the clean dataset. Unfortunately, HHT's frequency estimations, and to a lesser extent those from EWT, fluctuate heavily, leading to high MAPE values. These distortions are caused by the interference between the multiple wavefunctions in each wavepacket. This effect increases dramatically for both algorithms in the noisy dataset⁴.

EEG. Owing to its ease of measurement and high temporal resolution, the vast majority of neuroscience studies are based on EEG measurements³⁵. As EEG measures brain activity via electrodes on the skull, no medical procedures are needed. However, such external measurements do suffer from increased noise. Fluctuations in EEG caused by brain activity are orders of magnitude smaller than the disturbances caused by eye, face and body movements³⁶. Therefore, studies average the recordings of many trials to cancel random fluctuations. Unfortunately, the use of repeated trials removes the temporal advantage of EEG and prevents its applicability in real-time implementations, which rely on single-trial estimation.

The often-used FFT cannot handle the highly non-stationary character of EEG signals. Additionally, EEG sampling frequencies are often 1 kHz, and the simultaneous recording of 64 electrodes is standard. Hence, high-speed, non-stationary, time–frequency analysis is essential to have any chance of success in single-trial estimation. This is a criterion that current time–frequency techniques are unable to meet. Techniques like STFT and DWT⁸ are fast but lack the desired resolution in representation, whereas methods like CWT⁶ are precise but lack speed. fCWT fuses the best of both worlds by accelerating the high-resolution CWT by 34 to 122 times. So, we can improve the resolution by ≥ 34 times or handle ≥ 34 times as many data than the fastest CWT implementation available in the same time frame. To demonstrate the impact of real-time super-resolution on neuroscience, fCWT was thus benchmarked against full-resolution CWT and fast STFT, and DWT on a single-trial EEG dataset of subjects performing mental arithmetic tasks³⁷.

Because active concentration is known to be most visible in the frontal region of the brain³⁶, the signals of three frontal electrodes (pre-frontal 1, pre-frontal 2 and mid-frontal in the 10–20 system³⁶) were averaged to reduce local fluctuations. We analyzed the resulting signal in the δ (delta), θ (theta), α (alpha), β (beta) and γ (gamma) frequency bands, using a frequency range that spans five octaves ($f_0=2$ Hz to $f_1=64$ Hz). Simultaneous analysis of all these frequency bands is vital for cognitive task experiments, with pre-frontal δ frequencies (2–4 Hz) being associated with attention and motivation³⁸, and the power of θ oscillations (4–7 Hz) reflecting memory encoding and retrieval³⁹. Lower α -desynchronization (8–13 Hz) relates to task-unspecific attentional demands and β -band (13–30 Hz) power increases with demanding cognitive tasks³⁶. The γ oscillations




Fig. 4 | Benchmark results for synthetic data. **a**, Synthetic data composed of wavepackets WP1, WP2 and WP3 (see Methods for details). Seven time-frequency estimation techniques that cover a frequency range from $f_0 = 0.25$ Hz to the Nyquist frequency $f_1 = 250$ Hz are shown. fCWT and CWT use the Morlet wavelet ($\sigma = 6$) and 480 frequencies to divide the spectrum, DWT uses 11 levels of 15-order Daubechie wavelet decomposition, and STFT uses a 500-ms Blackman window with 400-ms overlap to obtain optimal time-frequency resolution. WVD takes no parameters. HHT and EWT have a frequency resolution of 0.25 Hz and rely on an adaptive iterating process. HHT uses seven intrinsic modes that were extracted using a maximum signal-to-residual ratio stopping criterion. A close-up of the time-frequency estimation of the third wavepacket is also shown for comparison. As relative intensity is of primary interest, the spectra are normalized to a [0, 1] range. **b**, As in **a**, but 0-dB white Gaussian noise is added to the synthetic data. The parameters remained the same. **c**, MAPE scores for the clean and noisy data. Boxes show the median and 25th to 75th percentile range; whiskers show minima and maxima. In the top plot only medians are visible as results on the clean dataset are deterministic and, hence, contain no variance. See Supplementary Table 1 for the distribution statistics.

(~30–100 Hz) indicate complex cognitive thinking (for example, object recognition and sensory processing⁴⁰). Consequently, full-range, high-resolution frequency analysis is vital.

The analysis of CWT, fCWT, STFT and DWT was complemented with 3.0% CWT (that is, CWT with fCWT's RAR; Fig. 5). 3.0% CWT enables a fair comparison between the real-time




Fig. 5 | Benchmark results of human EEG data. **a**, The Fp1 and Fp2 pre-frontal and Fz mid-frontal EEG electrodes, which were averaged to assess mental workload. Credit: Imagewriter/Alamy. **b**, Full fCWT and CWT, 3.0% CWT, STFT and DWT of EEG, recorded during 30 s of rest and 30 s of mental arithmetic. Full fCWT and 3.0% CWT analyze the signal using the Morlet wavelet ($\sigma=20$) at 650 and 20 scales, evenly spaced in exponential space, respectively. STFT uses a 500-ms Blackman window with 400-ms overlap and DWT uses 11 levels of 15-order Daubechie wavelet decomposition. Spectra are normalized to [0, 1], except for a few spectra that are amplified to enhance visibility. **c**, Zoomed view during the arithmetic task to show each algorithm’s ability to extract the intricate time–frequency details of the β frequency band (13–30 Hz). **d**, The RAR (equation (1)) of full fCWT and CWT, 3.0% CWT, STFT and DWT versus the number of electrodes with a 1-kHz EEG signal.

resolution of CWT and full fCWT using 650 frequencies and 3.0% CWT using 20 frequencies. The three CWTs use the complex-valued Morlet wavelet ($\sigma=20$), tuned for optimal time–frequency resolution. Based on manual tuning we set a 500-ms Blackman window with 400-ms overlap for STFT and 11 dyadic (that is, $a^l=2^l$) scales of 15-order Daubechie wavelet decomposition for DWT, enabling maximal time–frequency sharpness. RAR versus the number of 1-kHz channels was calculated for full-resolution CWT and fCWT, STFT and DWT.

The resolution difference between the equally fast full fCWT and 3.0% CWT is most prominent during the mental arithmetic task. Real-time fCWT distinguishes different EEG frequency bands much better than real-time CWT. The sheer amount of subdivisions in the frequency spectrum allows fCWT to show the small chaotic β -frequency variations often seen during active concentration³⁶ and the slow oscillating δ -band power associated with motivation³⁸, in real time. Having the same runtime, the fastest CWT implementation fails. Although STFT can separate frequencies in the β -frequency (13–30 Hz) and γ -frequency (~ 30 –100 Hz) bands, it suffers from low spectral resolution in the δ -frequency (<4 Hz) and θ -frequency (4–7 Hz) bands. Hence, STFT makes wide-band EEG analysis impractical. Again, DWT was shown to be unsuitable for detailed time–frequency analysis.

fCWT’s power excels when an entire array of EEG electrodes is analyzed in real time. Although the use of EEG is gaining popularity, its low spatial resolution remains a huge drawback. Figure 5 shows that the fastest CWT implementation available can only handle ~ 20 –24 electrodes (or streams of data) simultaneously at full resolution in real time. By contrast, fCWT is easily capable of

calculating real-time, high-resolution time–frequency representations of state-of-the-art EEG set-ups with up to 512 electrodes.

In vivo electrophysiology. Using depth electrodes, local field potentials (LFPs) measure local voltage changes inside the brain caused by the activity of neuron clusters. LFPs are recorded in vivo and, consequently, they do not suffer from the skull’s high-frequency mask behavior. Consequently, the γ -frequency (~ 30 –100 Hz) and high γ -frequency (>100 Hz) bands can be reliably recorded, these being bands that highly correlate with single neuron firing and reflect aspects of movement (in the motor cortex⁴¹) and vision (in the visual cortex⁴²). Recording these frequencies requires sampling rates that are several times those used for EEGs (that is, 2–3 kHz). Furthermore, in vivo electrophysiology techniques⁴³ use huge amounts of electrodes⁴⁴. LFPs are often recorded simultaneously at 100–300 channels, or even more⁴⁵. In the future, data bandwidth is expected to increase even more than its recent tremendous increases. Neuropixels⁴³, Utah arrays⁴⁴ and Michigan probes⁴⁶ are currently able to measure hundreds of LFPs and thousands of neurons simultaneously. Real-time LFP time–frequency analysis could lead to next-generation prosthetics⁴¹. Unfortunately, current implementations are unable to handle these bandwidths without compromising resolution. fCWT shows that super-resolution can be maintained when analyzing hundreds of high-bandwidth LFP data streams simultaneously.

Rodent in vivo electrophysiology data from the Allen Brain Observatory data collection⁴⁷ were analyzed. During randomly alternating full-field, high- and low-contrast flashes, six Neuropixel probes⁴³ with 374 electrodes (Neuropixel 3a; 20 μ m vertical




Fig. 6 | Benchmark results of in vivo electrophysiology data. **a**, In vivo electrophysiology measurements were obtained by the insertion of a Neuropixel⁴³ inside the anteromedial area of a rodent's visual cortex. Mouse drawing adapted from ref. ⁸⁴ under a CC BY license. **b**, Time–frequency estimations by fCWT, CWT, STFT and DWT during 9 s of four 250-ms full-field, high- and low-contrast flashes. The LFP shows exclusive activation after the black stimuli. Full fCWT and 3.0% CWT analyze the signal using the Morlet wavelet ($\sigma=16$) at 520 and 16 scales evenly spaced in exponential space, respectively. STFT uses a 500-ms Blackman window with 400-ms overlap and DWT uses 11 levels of 15-order Daubechie wavelet decomposition. Spectra are normalized to [0, 1], except for a few spectra that are amplified to enhance visibility. **c**, Zoom-in of the β - (15–30 Hz), γ - (32–100 Hz) and high γ -frequency bands (>100 Hz), immediately after a black stimulus. Three frequency components in the β -frequency band and two γ bursts are present. Plot scales are aligned as well as possible, despite differences in exponential scale (fCWT and CWT) and linear scale (STFT). **d**, The RAR (equation (1)) of full-resolution fCWT and CWT, 3.0% CWT, STFT and DWT versus the number of channels in a 2.5-kHz electrophysiology signal.

electrode separation) each recorded a mouse visual cortex's responses. LFPs were obtained by downsampling the data to 1.25 kHz and filtering using a 1,000-Hz low-pass filter. Full fCWT and CWT, 3.0% CWT (EEG section), STFT and DWT time–frequency estimations were performed on 9 s of raw single-trial LFP data containing four stimuli.

We compared CWT and fCWT to STFT and DWT, as the latter two are used in situations where speed is key. Other time–frequency algorithms offer much higher resolution but are orders of magnitude slower, making them impractical for LFP analysis.

The analysis covers a frequency range from $f_0=8$ Hz to $f_1=128$ Hz, allowing simultaneous analysis of both low frequency (that is, α and β bands) and high frequency (that is, γ and high γ bands), which is very important as they reflect different aspects of task performance. Low-frequency LFPs unveil long-distance communication, whereas high-frequency activity reflects local neural processing⁴⁸. As the interplay between these frequency ranges discloses the coordination at the inter- and intra-cortical level⁴⁹, real-time, wide-band time–frequency estimation is key in the LFP analysis of complex brain mechanics.

The three CWTs use the complex-valued Morlet wavelet ($\sigma=16$), tuned for optimal time–frequency resolution. Based on manual tuning we set a 500-ms Blackman window with 400-ms overlap for STFT and 11 dyadic (that is, $a^i=2^i$) scales of 15-order Debauchie wavelet decomposition for DWT, enabling maximal time–frequency sharpness. The RAR versus number of channels was also

calculated for fCWT and CWT at full resolution and STFT and DWT for a 2.5-kHz input signal.

The subfigures of Fig. 6c show the ability of real-time, full fCWT to separate multiple β -frequency components (16, 20 and 25 Hz), locate four γ bursts and reveal the overall γ -frequency dynamics, all at the same time. By contrast, real-time 3.0% CWT misses two out of four γ bursts, cannot separate low-frequency β components, and loses higher γ -frequency dynamics. With STFT, the resolution is on par in the mid-frequency range, but the high- and low-frequency ranges suffer from low resolution. Despite their very high speeds, both STFT and DWT are unsuitable for broadband, high-resolution, time–frequency estimations.

Electrode density is set to increase dramatically; for example, 5,000-electrode Neuropixels have already been announced⁵⁰. Figure 6d shows RAR (equation (1)) versus the number of channels per algorithm. Full CWT can hardly process 15 LFP channels (or data streams) in real time. By contrast, fCWT offers a real-time, full-resolution performance for up to 350–400 channels. Considering the Allen Brain Observatory dataset, fCWT supports real-time analysis and feature extraction of three to four entire Neuropixel probes, whereas the fastest CWT implementation available supports only one-tenth of a single probe.

Discussion

One of WT's most powerful features is the possibility to use custom wavelets. However, not all wavelet types are suitable for existing fast

approximate CWT implementations, which rely on finite impulse response filters⁴. fCWT does not suffer from this setback, as it calculates wavelets starting directly from its definition. With custom wavelets, fCWT performance can be improved even further⁵¹. As such, fCWT enables the real-time analysis of high-frequency non-stationary signals, such as in audio^{22–24,52}, biosignals (for example, brain–computer interfaces¹² and ECG^{11,13}), image and video^{25,26}, sonar and radar^{27,28}, renewable energy management^{16,17}, cybersecurity^{14,15} and machine fault diagnosis^{29,30,53} (Fig. 1).

The implementation of fCWT could be extended to other time–frequency methods as well. The synchrosqueezed transform (SST)⁵⁴ uses reassignment to sharpen the CWT spectrum, and the chirplet transform (CT)⁵⁵, superlets (SL)⁶ and the noiselet transform (NT)⁵⁶ use atoms to describe a signal, sharing a wavelet-like implementation. Future research could explore speed-ups of these algorithms and bring them to real-time applications. Hence, fCWT’s impact is broader than CWT-based applications alone. Consequently, we did not include the SST, CT, SL and NT in the benchmark study, as these rely on the CWT in their core. These second-order techniques as well as modifications of the included first-order techniques (for example, smoothed WVD⁶) are by definition slower than the already expensive CWT.

fCWT shares its mathematical definition with CWT and, hence, without compromise, inherits both all its benefits¹⁰ and all its limitations (for example, its degrading spectral resolution⁵⁷ and increasing redundancy in higher frequency ranges⁵). Fortunately, these are well-known limitations that have solutions^{4,54}. Moreover, the time–frequency landscape keeps growing, including new CWT implementations⁵⁸. We therefore invite everyone to compare their implementations against fCWT’s open source⁵⁹, and, to extend its validity, we invite all to apply fCWT on more extensive and different specimens that fall outside this paper’s scope.

fCWT allows an acceleration in the developments of science and engineering, industry and health (Fig. 1). Although maintaining CWT’s full resolution and supporting customization, fCWT enables real-time time–frequency analysis of non-stationary signals. As such, fCWT can bring offline research that is hindered by the low resolution of DWT, the limited range of STFT and/or the computational burden of CWT into real-time practice.

Methods

Datasets. In this Resource paper, three types of data were used: synthetic, EEG and in vivo electrophysiological data. Details on each dataset are described in the following subsections.

Synthetic data. Two synthetic datasets were generated for this paper, both composed of the same three time-varying wavepackets with a sampling frequency of 500 Hz:

1. Three 5-s sine waves, the frequencies of which gradually change between 100 and 110 Hz, 20 and 22 Hz and 5 and 6 Hz, respectively, with a periodicity of 1 Hz.
2. Two 5-s sine waves with linearly changing frequencies between [5, 50] and [100, 50] Hz.
3. Three 10-s low-frequency waves of 2, 1 and 0.5 Hz. All wavepackets are separated by 0.5 s and are multiplied by a Gaussian window function to mitigate discontinuities at the boundaries.

One set contained clean data and the other was contaminated with white Gaussian noise with a 1:1 signal-to-noise ratio (SNR) across the whole signal, with the SNR being determined by the average power. Both datasets have a total duration of 21.0 s and are available in the fCWT CodeOcean repository⁵⁹.

EEG. The EEG mental arithmetic dataset by Zyma et al.³⁷ was obtained from PhysioNet⁶⁰ and loaded into MATLAB R2021a. EEG data were recorded monopolarly at 500 Hz, using Ag/Ag electrodes and the Neurocom EEG 23-channel system (Ukraine, XAI-MEDICA). The International 10/20 scheme was used for electrode placement. Electrodes were referenced to the interconnected ear reference electrodes. Data were preprocessed using a 30-Hz high-pass filter and a 50-Hz power line notch filter. Common EEG artifacts were removed using independent component analysis. All participants had normal or corrected-to-normal vision and had no mental or cognitive impairment.

In this paper we use the data of subject 13, a 24-year old male who excelled in mental arithmetic by performing 34 subtractions between four-digit and two-digit numbers in 4 min. Subject 13 was chosen to ensure task compliance. We used the last 30 s of EEG during rest and the first 30 s of EEG during the arithmetic task.

In vivo electrophysiology. In vivo electrophysiology data were collected from The Visual Coding—Neuropixels project⁴⁷. LFP data from female specimen 738651054 from stimuli IDs 3861–3864 were used. Six Neuropixel version 3a probes were inserted into the mouse visual cortex. In this study, LFP data from fifth probe (Probe ‘e’) channel 63 were used. The 250-ms high-contrast stimuli, 2,000 ms apart, alternate in random order. Mice were shown a neutral gray screen between stimuli. Additional technical, experimental and medical details about the dataset can be found in ref. ⁴⁷.

Mathematical preliminaries. *The Fourier transform.* With its core idea that a function, often a signal, can always be decomposed into pure sine and cosine functions, the FT is foundational in spectral pattern analysis^{3,4,8,61}. However, not all functions $f(t)$ can be decomposed—only those that live in the Lebesgue space $L^2(0, 2\pi)$. This space includes all functions that are (1) finite in energy, (2) 2π -periodic and (3) square-integrable, formally

$$\int_0^{2\pi} |f(t)|^2 dt < \infty \quad t \in (0, 2\pi) \quad (2)$$

$$f(t) = f(t - 2\pi) \quad t \in \mathbb{R}, \quad (3)$$

which allows $f(t)$ to be represented as a weighted sum of complex wavefunctions:

$$f(t) = \sum_{-\infty}^{\infty} c_n e^{2\pi i n t}, \quad (4)$$

with the Fourier coefficients c_n given by the amount of overlap between the conjugated complex wavefunction and the function $f(t)$:

$$c_n = \frac{1}{2\pi} \int_{-\infty}^{\infty} f(t) e^{-2\pi i n t} \quad (5)$$

or in discrete form when used on actual digital samples in a sequence f having length N :

$$x_k = \sum_{n=0}^{N-1} f[n] e^{-i2\pi k n / N}. \quad (6)$$

In other words, any 2π -periodic, square-integrable function $f(t)$ can be represented by this superposition of complex-valued sinusoidal waves that are translated in the frequency domain. However, this is precisely Fourier’s pitfall; not all functions, or signals for that matter, are 2π -periodic. FTs cannot decompose the wide variety of non-stationary functions that are not 2π -periodic. Unfortunately, this constraint is often misunderstood, and FT are still used to analyze signals with varying frequencies.

The mathematical reason behind FT’s constraint becomes apparent when we consider the Lebesgue space $L^2(\mathbb{R})$ containing all square-integrable functions that have finite energy along the entire real axis:

$$\int_{-\infty}^{\infty} |f(t)|^2 dt < \infty. \quad (7)$$

The reason why equation (4) cannot represent these functions is that pure sine waves extend to infinity and therefore do not have finite energy. Pure waves do not lie in $L^2(\mathbb{R})$ and, as such, they cannot represent its functions.

Wavelets. We can define a set of functions other than equation (4) that do have finite energy. The result is the set of short periodic functions $\psi(t)$ called wavelets that are well localized in both the time and frequency domains^{5,6,8,33,57,62}. Consequently, wavelets need to be able to translate in both domains as well:

$$\psi_{jk}(t) = 2^{-j/2} \psi(2^j t - k), \quad (8)$$

where ψ_{jk} is a daughter wavelet function, defined as the mother wavelet $\psi(t)$ scaled in the frequency domain by j and translated in the time domain by k . So, the WT outputs a 2D time–frequency matrix, where the FT gives a 1D frequency spectrum.

Similar to equation (4), the superposition of these wavelets can represent any function

$$f(t) = \sum_{j,k=-\infty}^{\infty} c_{jk} \psi_{jk}(t), \quad (9)$$

where, like with the FT, the wavelet coefficients c_{jk} are given by the amount of overlap between the wavelet and the function $f(t)$. This definition also shows us

that wavelets, similar to Fourier’s pure wavefunctions, live in Hilbert space as multiplication between two functions is defined as an inner product:

$$W_{\psi}f(j, k) = \langle f, \psi_{jk} \rangle = c_{jk} = \int_{-\infty}^{\infty} f(t)\overline{\psi_{jk}(t)}dt, \tag{10}$$

in which $\overline{\psi_{jk}}$ corresponds to the conjugate of ψ_{jk} . However, as j and k can be any real number, we have to define both variables’ optimal discretization such that the resulting time–frequency matrix does not under- or overdetermine the function $f(t)$. So, the variables should be discretized such that the wavelets form an orthogonal basis in Hilbert space^{63,64}—in other words, such that the wavelet functions have zero overlap.

Wavelets are orthogonal in Hilbert space if

$$\langle \psi_{jk}, \psi_{lm} \rangle = \delta_{jk}\delta_{lm}, \tag{11}$$

from which it follows that equation (8) is indeed logarithmic orthogonal. The WT that uses this type of discretization is called the DWT^{8,65,66}. In this context, ‘discrete’ refers to the use of its wavelets, not to the type of data it processes. As all DWT’s wavelets are orthogonal, it describes a function by the minimal number of wavelet coefficients possible. However, as stated at the beginning of this paper, a redundant, overcomplete representation is often much more favorable for signal analysis. Therefore, it is also possible to define a WT with arbitrary wavelet discretization. Such a wavelet transformation is called the CWT⁶⁷. Again, ‘continuous’ does not refer to the type of data it can handle. CWT features continuously scalable and translatable wavelets that allow a much more precise analysis of a signal’s spectrum:

$$W_{\psi}f(a, b) = |a|^{-1} \int_{-\infty}^{\infty} f(t)\overline{\psi\left(\frac{t-b}{a}\right)}dt, \tag{12}$$

which comes with considerable computational complexity. When implemented digitally, its discrete form is used:

$$W_{\psi}f[a, b] = |a|^{-1} \sum_{n=0}^{N-1} f[n]\overline{\psi\left[\frac{n-b}{a}\right]}, \tag{13}$$

which is mathematically equivalent to passing the input signal through a series of wavelet filters of different lengths. Care is required at the boundaries of the signal. As the discrete form assumes signals of finite length, wavelet coefficients near the boundaries become increasingly meaningless. Instantaneous frequency at the first or last sample is impossible to calculate as one should know how the signal continues. There are several strategies to solve this uncertainty. For more details about this topic, see the Boundary effects section.

Equation (10)’s computational complexity can be estimated using the trapezoidal rule for integral solving and assuming a signal of length $N = 2^l$. Furthermore, we assume J wavelets at $a_j = 2^j$ discrete scales, and a wavelet length of L samples at unit scale. Starting at unit scale $a_0 = 1$, we then have $O(a_0NL)$ complexity, with the cost of all scales resulting in

$$NL + 2NL + 4NL + \dots + 2^lNL = O(LN^2). \tag{14}$$

In other words, a naïve approach to DWT calculation would result in a polynomial complexity of $O(N^2)$. CWT would be even worse, as the discretization of the time and frequency domains is much finer. Fortunately, scientists quickly realized a considerable reduction in computational complexity could be achieved using Parseval’s theorem.

Fourier-based wavelet transform. Applying Parseval’s theorem to equation (12), a reduction in CWT’s complexity can be achieved:

$$W_{\psi}f(a, b) = \frac{1}{2\pi} \int \hat{f}(\xi)\overline{\widehat{\psi}_{a,b}(\xi)}d\xi. \tag{15}$$

Subsequently, we define $\widehat{\psi}_{a,b}(\xi)$ in terms of the FT of the mother wavelet function $\psi(t)$, using its basic time-shifting and time-scaling properties:

$$\widehat{\psi}_{a,b}(\xi) = \frac{1}{a}\widehat{\psi}(\xi)e^{-ib\xi} \quad (\text{time shifting}) \tag{16}$$

$$= \widehat{\psi}(a\xi)e^{-ib\xi} \quad (\text{time scaling}). \tag{17}$$

Substitution gives

$$W_{\psi}f(a, b) = \frac{1}{2\pi} \int \hat{f}(\xi)\overline{\widehat{\psi}(a\xi)}e^{ib\xi}d\xi \tag{18}$$

or in its discrete form

$$W_{\psi}f[a, b] = \frac{1}{K} \sum_{k=0}^{K-1} \hat{f}[k]\overline{\widehat{\psi}[ak]}e^{i2\pi bk/K}, \tag{19}$$

which describes $W_{\psi}f[a, b]$ as an inverse FT of $\hat{f}[k]\overline{\widehat{\psi}[ak]}$. So, WT’s computational complexity no longer depends on the time–offset parameter b . As $\hat{f}[k]$ can be calculated beforehand, it is reduced to three distinct steps per scale:

1. Generate $\overline{\widehat{\psi}[ak]}$
2. Calculate $\hat{f}[k]\overline{\widehat{\psi}[ak]}$ and
3. Evaluate the inverse FT and obtain $W_{\psi}f[a, b]$,

with the first two steps evaluated in $O(N)$ and the last one requiring at least $O(N\log_2N)$ when using a fast FT implementation^{68,69}. This results in $O(N\log_2N)$ complexity, a considerable reduction compared to $O(N^2)$, which is needed for the naïve approach. Additionally, the constant factor of this complexity can be reduced even more, as we will see in the next section.

Implementation of fCWT. Fourier-based wavelet transformation’s computational complexity is mainly determined by the inverse FT. Consequently, equation (12) has been rewritten regularly to use spline interpolation of the wavelet and circumvent the FT entirely^{70,71}. Spline interpolation, also known as polynomial interpolation, defines a wavelet by only a few evenly spaced sampling points across the domain. Because the number of points is independent of the wavelet’s scale, the theoretical complexity of equation (12) is reduced to linear time. However, while complexity is lowered, the constant factor that equals the number of sampling points has been increased tremendously. In turn, this yields a trade-off between speed and accuracy: more interpolation points leads to increases in both precision and computation time. Additionally, the spline interpolation only works for specific wavelet types. To avoid the trade-off, we optimize the Fourier-based wavelet transformation by reducing the constant factor of its computational complexity. In this way, we maintain WT’s ability to use custom wavelet types⁵¹ and can exploit optimized FFT libraries^{72–74}.

fCWT separates scale-independent and scale-dependent operations, which have to be performed separately for each wavelet’s scale. A detailed schematic of fCWT’s algorithmic implementation is provided in Extended Data Fig. 1. With CWTs, the frequency scale is often divided into hundreds of scales. We thus focused the optimization on the fCWT’s scale-dependent part by exploiting its repeated nature and high parallelizability. The scale-independent operations are performed first as their result forms the input for the scale-dependent steps. We pre-calculate two functions: (1) the input signal’s FFT and (2) the FFT of the mother wavelet function at scale $a_0 = 2$. Both functions are independent of the scale factor a , so they can be pre-calculated and used as look-up tables in the processing pipeline.

FFT. Using the float- and AVX2-enabled Fastest Fourier Transform in the West (FFTW) library⁷³, the input signal’s FFT is calculated. FFTW has superior performance in various benchmarks⁷⁵ and has the ability to dynamically optimize its algorithmic implementation. FFTW determines the most efficient way to calculate the signal’s FFT with length N on hardware set-up X . This requires considerable time, which makes it only useful in situations where many FFTs are calculated with the same N and X . This is the case with fCWT, as its scale-dependent part evaluates a fixed-length inverse FFT for every scale factor a . Other high-performance FFT libraries include the Fastest Fourier Transform in the South⁷² and Intel’s Math Kernel Library⁷⁴. However, as Fastest Fourier Transform in the South lacks important optimization techniques and Intel’s Math Kernel Library is limited to Intel processors only, FFTW is currently the most flexible and versatile high-performance FFT library available.

Before a signal’s FFT is calculated, it is first zero-padded to the nearest power of two, which allows more time-efficient calculations than with other signal lengths. Zero padding lets all signals that map to the same nearest power of two use the same FFTW optimization. Hence, the flexibility of fCWT as a tool is preserved while still enjoying the benefit of FFTW’s optimization plans. However, it will result in step-like performance behavior as seen in Fig. 3. After FFT calculation, we let FFTW write the complex-valued FT to memory in an interleaving format (Extended Data Fig. 2). Using this, we exploit the CPU’s predictive caching behavior and hence reduce memory access in the next steps. Because a CPU works with chunks of memory instead of single values, it always caches adjacent memory next to a requested value as well^{76,76}. While we access the real part of a value, interleaving takes advantage of this behavior as the complex part is cached. Consequently, accessing the complex part after the real part does not require an additional memory request, which reduces memory accesses by 50%.

Scale-independent mother wavelet generation. The FFT of the mother wavelet function $\widehat{\psi}[k]$ is generated once during the scale-independent step. Because wavelets in the frequency domain uniformly contract as their scale increases, daughter wavelet functions can be generated by downsampling a pre-generated mother wavelet function. Because scales must be at least $a_{\min} = 2$, we generate the mother wavelet function at $a_0 = 2$ to save memory. It is important to note that the mother wavelet function is generated directly from its analytical Fourier-transformed definition. Consequently, we create $\widehat{\psi}[k]$ such that its length always matches that of $\hat{f}[k]$. This ensures fCWT’s independence of wavelet length and achieves the highest wavelet resolution possible.

After calculation of the FFT signal and the generation of the FFT mother wavelet, the scale-independent step is complete. fCWT proceeds to the scale-dependent phase (Extended Data Fig. 1). This phase is repeated $m = |a|$ times. Using m steps with step size Δa , the scale factors a are defined by discretizing the frequency spectrum evenly on a logarithmic scale:

$$a = \{2^{x\Delta a} | x \in N \wedge 0 < x \leq m\}. \quad (20)$$

This generates the wavelet coefficient matrix $W_w[f(a, b)]$ one row at a time.

Scale-dependent downsampling. Each iteration of the scale-dependent step first generates the Fourier-transformed daughter wavelet function $\hat{\psi}_a[k]$ by downsampling the mother wavelet function generated in the scale-independent step. This optimization is realized by using the mother wavelet as a look-up table (Extended Data Fig. 3). Hence, as explained earlier, the expensive Gaussian calculations involved in wavelet calculation are removed from the scale-dependent step. The daughter wavelet is generated by only performing a systematic look-up that accesses every a th value of the mother wavelet function. A schematic overview of this process is shown in Extended Data Fig. 3.

Scale-independent multiplication. Using the single instruction, multiple data (SIMD) model, another acceleration is achieved. By using the CPU's full power, eight multiplications are executed at once^{77,78}, which is used to exploit the elemental-wise multiplication between $\hat{\psi}_a[k]$ and $\hat{f}[k]$. In our case, SIMD performs four complex-valued multiplications in parallel, because the multiplication between the real-valued daughter wavelet and complex-valued Fourier-transformed input signal takes two multiplications per element. By exploiting the parallelizable nature of this step, an additional speed-up of 4× is achieved. Extended Data Fig. 4 shows this process graphically.

An additional acceleration is achieved by merging the generation of the daughter wavelet (Scale-dependent downsampling section) and the multiplication with $\hat{f}[k]$ (Scale-independent multiplication section) in one loop. Consequently, no intermediate results are stored in memory, which largely eliminates memory access.

Scale-independent inverse FFT. Finally, using FFTW's inverse FFT function, the result is transferred back to the time domain. Similar to the FFT calculation in the scale-independent step, the inverse FFT uses a pre-calculated optimization based on the input signal's zero-padded length. The complex-valued time–frequency matrix is stored in row-major order as an array of $2NM$ floats, where N is the signal length and M the number of scales. Each value is stored as two floats as the matrix is complex-valued.

Boundary effects. Because CWT uses convolution to calculate the wavelet coefficients, a wavelet is eventually close enough to the beginning or end of the signal to be multiplied with undefined data outside the boundaries of the signal. In these situations, frequency becomes a meaningless construct as one does not know how a signal would proceed beyond these limits. As this effect becomes more evident with larger wavelets (that is, lower frequencies) one can speak of a cone of influence³³ caused by the edges that affect the entire spectrum. Several strategies exist to handle these so-called boundary effects^{4,79,80}.

One could extend a signal by adding zeros at the beginning and the end to define data outside the boundaries. Because convolution relies on the element-wise multiplication between the signal and the wavelet, this strategy is similar to stopping the convolution at the boundary. Other strategies rely on making assumptions about the signal outside its bounds. For example, the signal could be extended by mirroring or repeating the signal at its boundaries⁴. With fCWT, we decided to let the users decide themselves. As fCWT is designed to be independent of signal content, we assume an unbiased zero extension. Consequently, users can choose their own boundary strategy by extending the signal manually before the fCWT is applied.

The direct result of fCWT's strategy can be seen in the Synthetic data section. At both edges, fCWT shows a strong cone of influence effect of the boundary. MATLAB, by default, performs signal extension, which mitigates these artifacts. However, MATLAB's default extension strategy sometimes leads to an increase in artifacts instead of a reduction. An extreme example can be seen in a visual comparison between both techniques in Extended Data Fig. 5. Consequently, with fCWT we went for an unbiased zero extension strategy aiming for maximal transparency and flexibility.

Time–frequency ridge extraction. To perform quantitative assessment of time–frequency spectra on the synthetic data, a time–frequency ridge extraction methodology is used. This allows a comparison between the ridges (that is, frequency components) in the time–frequency spectra and the actual frequency components used to generate the dataset.

The synthetic dataset (see the Data availability statement for details) consists of three distinct wavepackets. The time–frequency ridge extraction is performed on each wavepacket separately. Each segment is defined such that it trims the first 0.5 s and last 0.5 s of each wavepacket to remove the Gaussian window function influence. In the third wavepacket, 3.0 s is trimmed from the end to remove the influence of boundary effects (Boundary effects section).

MATLAB's tfridge() is used to extract, respectively, three, two and three ridges from the first, second and third wavepacket, as it is the most used approach to ridge extraction. To do so, tfridge() needs pre-defined penalty coefficients (P_{coef}), which determine the stability of the ridge estimation. As wavepackets and time–frequency algorithms largely differ in their characteristics, these penalty coefficients need to be optimized manually for each combination. Manual optimization is performed greedily by first testing different orders of magnitude (P_{mag}): $P_{\text{mag}} \in \{10^{-3}, 10^{-2}, 10^{-1}, 10^0, 10^1, 10^2, 10^3\}$. When the optimal magnitude range $[P_{\text{mag},1}, P_{\text{mag},2}]$ is selected, ten equally distanced coefficients are tested within that range $P_{\text{coef}} \in P_{\text{mag},1} \cdot \{0, 1, 2, \dots, 8, 9\}$. The resulting penalty coefficient is chosen for the benchmark. All penalty coefficients are provided in the source data for Fig. 4.

A fair comparison among the algorithms was secured, as the same stable ridge extraction was applied on all included algorithms. However, in future work, alternate open-source algorithms could be worth exploring³¹. These might yield highly accurate ridge extractions and/or remove the need for manually tuned parameters. As such, this might result in an even more fine-grained comparison among the algorithms.

Data availability

The generated synthetic dataset used in Fig. 4 is provided under 'data' in the CodeOcean fCWT capsule³⁹. The 'EEG During Mental Arithmetic Tasks v1.0.0' used in Fig. 5 is available at <https://physionet.org/content/eeegmat/1.0.0/>. The in vivo electrophysiology data collected by The Visual Coding—Neuropixels project⁴⁷ and used in Fig. 6 is available in the Neurodata Without Borders (NWB) format via AllenSDK (<https://allensdk.readthedocs.io>). An example Jupyter Notebook for accessing the LFP data is available at https://allensdk.readthedocs.io/en/latest/_static/examples/nb/ecephys_lfp_analysis.html. Source data are provided with this paper.

Code availability

fCWT is released under Apache License Version 2.0 and will be maintained in a public GitHub repository available at <https://github.com/fastlib/fCWT>. fCWT is available for Mac OSX, Linux and Windows systems and features a detailed description on how to obtain Visual Studio project-files and generate cpp-based MEX-packages for easy MATLAB integration. For benchmark reproducibility, a version of fCWT at the time of publication is available on CodeOcean³⁹.

Received: 4 January 2021; Accepted: 13 December 2021;

Published online: 27 January 2022

References

- Gabor, D. Theory of communication. Part 1: the analysis of information. *J. Inst. Electr. Eng.* **93**, 429–441 (1946).
- Fano, R. M. *Transmission of Information: A Statistical Theory of Communications* (MIT Press, 1961).
- Boukouvola, E., Miridakis, N. & Veloni, A. *Digital and Statistical Signal Processing* (CRC Press, 2019).
- Boashash, B. *Time–Frequency Signal Analysis and Processing: A Comprehensive Reference* 2nd edn (Academic, 2016).
- Addison, P. S. Introduction to redundancy rules: the Continuous Wavelet Transform comes of age. *Philos. Trans. R. Soc. A Math. Phys. Eng. Sci.* **376**, 20170258 (2018).
- Moca, V. V., Bârzan, H., Nagy-Dabâcan, A. & Mureşan, R. C. Time–frequency super-resolution with superlets. *Nat. Commun.* **12**, 337 (2021).
- Stockwell, R. G., Mansinha, L. & Lowe, R. P. Localization of the complex spectrum: the S transform. *IEEE Trans. Signal Process.* **44**, 998–1001 (1996).
- Cohen, L. *Time–Frequency Analysis* (Prentice Hall, 1995).
- Smeets, H., Ceriotti, M. & Marrón, P. J. Adapting recursive sinusoidal software oscillators for low-power fixed-point processors. *ACM Trans. Embedded Comput. Syst.* **19**, 1–26 (2020).
- Adeli, H., Zhou, Z. & Dadmehr, N. Analysis of EEG records in an epileptic patient using wavelet transform. *J. Neurosci. Methods* **123**, 69–87 (2003).
- Sharma, K., Castellini, C., van den Broek, E. L., Albu-Schaeffer, A. & Schwenker, F. A dataset of continuous affect annotations and physiological signals for emotion analysis. *Sci. Data* **6**, 196 (2019).
- Nason, S. R. et al. A low-power band of neuronal spiking activity dominated by local single units improves the performance of brain–machine interfaces. *Nat. Biomed. Eng.* **4**, 973–983 (2020).
- van den Broek, E. L. *Affective Signal Processing (ASP): Unraveling the Mystery of Emotions*. PhD thesis, Univ. Twente (2011); <https://doi.org/10.3990/1.9789036532433>
- Dalal, M. & Juneja, M. Steganography and steganalysis (in digital forensics): a cybersecurity guide. *Multimed. Tools Appl.* **80**, 5723–5771 (2021).
- Derbeko, P., Dolev, S. & Gudes, E. Wavelet-based dynamic and privacy-preserving similitude data models for edge computing. *Wirel. Netw.* **27**, 351–366 (2021).
- Eroğlu, H., Cuce, E., Cuce, P. M., Gul, F. & Iskenderoğlu, A. Harmonic problems in renewable and sustainable energy systems: a comprehensive review. *Sustain. Energy Technol. Assess.* **48**, 101566 (2021).

17. Adebayo, T. S. & Kirikkaleli, D. Impact of renewable energy consumption, globalization and technological innovation on environmental degradation in Japan: application of wavelet tools. *Environ. Dev. Sustain.* **23**, 16057–16082 (2021).
18. Carmona, R., Hwang, W.-L. & Torresani, B. *Practical Time–Frequency Analysis: Gabor and Wavelet Transforms, with an Implementation in S* (Academic, 1998).
19. Lee, G. R., Gommers, R., Waselewski, F., Wohlfahrt, K. & O Leary, A. PyWavelets: a Python package for wavelet analysis. *J. Open Source Softw.* **4**, 1237 (2019).
20. Virtanen, P. et al. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nat. Methods* **17**, 261–272 (2020).
21. Hussain, R. Wavelib (GitHub, 2015); <https://github.com/rafat/wavelib>
22. Beerends, J. G., van Nieuwenhuizen, K. & van den Broek, E. L. Quantifying sound quality in loudspeaker reproduction. *J. Audio Eng. Soc.* **64**, 784–799 (2016).
23. Beerends, J. G. et al. Subjective and objective assessment of full bandwidth speech quality. *IEEE/ACM Trans. Audio Speech Lang. Process.* **28**, 440–449 (2020).
24. van den Broek, E. L., van der Sluis, F. & Dijkstra, T. in *Philips Research Book Series* Vol. 12, Ch. 10 (eds Westerink, J. H. D. M., Krans, M. & Ouwerkerk, M.) 153–180 (Springer Science+Business Media B.V., 2011).
25. Karam, C., Sugimoto, K. & Hirakawa, K. Fast convolutional distance transform. *IEEE Signal Process. Lett.* **26**, 853–857 (2019).
26. Schouten, T. E. & van den Broek, E. L. Fast exact Euclidean distance (FEED): a new class of adaptable distance transforms. *IEEE Trans. Pattern Anal. Mach. Intell.* **36**, 2159–2172 (2014).
27. Huang, Z., Shi, Y., Tang, B. & Zhang, J. Unimodular multiple-input-multiple-output radar wave-form design with desired correlation properties. *IET Radar Sonar Navigation*. (2021); <https://doi.org/10.1049/rsn2.12192>
28. Gao, C., Wang, D., Qiao, K. & Jiang, X. Optical system design of fully symmetrical Fourier transform lens. *Optical Rev.* **28**, 349–357 (2021).
29. Hannon, B., Sergeant, P., Dupré, L. & Pfister, P.-D. Two-dimensional Fourier-based modeling of electric machines—an overview. *IEEE Trans. Magn.* **55**, 8107217 (2019).
30. Jalayer, M., Orsenigo, C. & Vercellis, C. Fault detection and diagnosis for rotating machinery: a model based on convolutional LSTM, fast Fourier and continuous wavelet transforms. *Comput. Ind.* **125**, 103378 (2021).
31. Huang, N. E. & Shen, S. S. P. *Hilbert Huang Transform and its Applications* 2nd edn, Vol. 16 (World Scientific, 2014).
32. Gilles, J. Empirical wavelet transform. *IEEE Trans. Signal Process.* **61**, 3999–4010 (2013).
33. Dremim, I. M., Ivanov, O. V. & Nechitailo, V. A. Wavelets and their uses. *Phys. Uspekhi* **44**, 447–478 (2001).
34. Huang, N. E. et al. The empirical mode decomposition and the Hilbert spectrum for nonlinear and non-stationary time series analysis. *Proc. R. Soc. A Math. Phys. Eng. Sci.* **454**, 903–995 (1998).
35. Valdes-Sosa, P. A. et al. The Cuban Human Brain Mapping Project, a young and middle age population-based EEG, MRI and cognition dataset. *Sci. Data* **8**, 45 (2021).
36. Cacioppo, J. T., Tassinary, L. G. & Berntson, G. G. *Handbook of Psychophysiology* 4th edn (Cambridge Univ. Press, 2017).
37. Zyma, I. et al. Electroencephalograms during mental arithmetic task performance. *Data* **4**, 14 (2019).
38. Knyazev, G. G. EEG delta oscillations as a correlate of basic homeostatic and motivational processes. *Neurosci. Biobehav. Rev.* **36**, 677–695 (2012).
39. Kahana, M. J., Sekuler, R., Caplan, J. B., Kirschen, M. & Madsen, J. R. Human theta oscillations exhibit task dependence during virtual maze navigation. *Nature* **399**, 781–784 (1999).
40. Engel, A. K., Fries, P. & Singer, W. Dynamic predictions: oscillations and synchrony in top-down processing. *Nat. Rev. Neurosci.* **2**, 704–716 (2001).
41. Mehring, C. et al. Inference of hand movements from local field potentials in monkey motor cortex. *Nat. Neurosci.* **6**, 1253–1254 (2003).
42. Womelsdorf, T., Fries, P., Mitra, P. P. & Desimone, R. Gamma-band synchronization in visual cortex predicts speed of change detection. *Nature* **439**, 733–736 (2006).
43. Jun, J. J. et al. Fully integrated silicon probes for high-density recording of neural activity. *Nature* **551**, 232–236 (2017).
44. Nordhausen, C. T., Maynard, E. M. & Normann, R. A. Single unit recording capabilities of a 100 microelectrode array. *Brain Res.* **726**, 129–140 (1996).
45. Frank, J. A., Antonini, M.-J. & Anikeeva, P. Next-generation interfaces for studying neural function. *Nat. Biotechnol.* **37**, 1013–1023 (2019).
46. Campbell, P. K., Jones, K. E., Huber, R. J., Horch, K. W. & Normann, R. A. A silicon-based, three-dimensional neural interface: manufacturing processes for an intracortical electrode array. *IEEE Trans. Biomed. Eng.* **38**, 758–768 (1991).
47. Siegle, J. H. et al. Survey of spiking in the mouse visual system reveals functional hierarchy. *Nature* **592**, 86–92 (2021).
48. Einevoll, G. T., Kayser, C., Logothetis, N. K. & Panzeri, S. Modelling and analysis of local field potentials for studying the function of cortical circuits. *Nat. Rev. Neurosci.* **14**, 770–785 (2013).
49. Fries, P. A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. *Trends Cogn. Sci.* **9**, 474–480 (2005).
50. Steinmetz, N. A. et al. Neuropixels 2.0: a miniaturized high-density probe for stable and long-term brain recordings. *Science* **372**, eabf4588 (2021).
51. Unser, M. & Chenouard, N. A unifying parametric framework for 2D steerable wavelet transforms. *SIAM J. Imaging Sci.* **6**, 102–135 (2013).
52. Qiu, Z., Lee, C.-M., Xu, Z. & Sui, L. A multi-resolution filtered-x LMS algorithm based on discrete wavelet transform for active noise control. *Mech. Syst. Signal Process.* **66**, 458–469 (2016).
53. Huo, Z., Zhang, Y., Francq, P., Shu, L. & Huang, J. Incipient fault diagnosis of roller bearing using optimized wavelet transform based multi-speed vibration signatures. *IEEE Access* **5**, 19442–19456 (2017).
54. Daubechies, I., Lu, J. & Wu, H.-T. Synchrosqueezed wavelet transforms: an Empirical Mode Decomposition-like tool. *Appl. Comput. Harmonic Anal.* **30**, 243–261 (2011).
55. Mann, S. & Haykin, S. in *Vision Interface*, Vol. 91, 205–212 (Citeseer, 1991).
56. Coifman, R., Geshwind, F. & Meyer, Y. Noiselets. *Appl. Comput. Harmonic Anal.* **10**, 27–44 (2001).
57. Strang, G. Wavelets. *Am. Sci.* **82**, 250–255 (1994).
58. Chu, J. P. & Kemere, C. T. Ghostipy: an efficient signal processing and spectral analysis toolbox for large data. *eNeuro* **8**, ENEURO.0202-21.2021 (2021); <https://doi.org/10.1523/ENEURO.0202-21.2021>
59. Arts, L. P. A. & van den Broek, E. L. The fast continuous wavelet transformation (fCWT): real-time, high quality, noise-resistant, time-frequency analysis (Code Ocean, 2021); <https://doi.org/10.24433/CO.8389373.v1>
60. Goldberger, A. L. et al. PhysioBank, PhysioToolkit and PhysioNet: components of a new research resource for complex physiologic signals. *Circulation* **101**, e215–e220 (2000).
61. Bracewell, R. N. *The Fourier Transform and its Applications* Vol. 31999 (McGraw-Hill, 1986).
62. Daubechies, I. *The Wavelet Transform, Time–Frequency Localization and Signal Analysis* (Princeton Univ. Press, 2009).
63. Mallat, S. G. Multiresolution approximations and wavelet orthonormal bases of $L^2(\mathbb{R})$. *Trans. Am. Math. Soc.* **315**, 69–87 (1989).
64. Halmos, P. R. *Introduction to Hilbert Space and the Theory of Spectral Multiplicity* (Dover, 2017).
65. Heil, C., & Walnut, D. F. *Fundamental Papers in Wavelet Theory* (Princeton Univ. Press, 2009).
66. Shensa, M. J. The discrete wavelet transform: wedding the à trous and Mallat algorithms. *IEEE Trans. Signal Process.* **40**, 2464–2482 (1992).
67. Aguiar-Conraria, L. & Soares, M. J. The continuous wavelet transform: moving beyond uni- and bivariate analysis. *J. Econ. Surv.* **28**, 344–375 (2014).
68. Johnson, S. G. & Frigo, M. A modified split-radix FFT with fewer arithmetic operations. *IEEE Trans. Signal Process.* **55**, 111–119 (2006).
69. Lundy, T. & Van Buskirk, J. A new matrix approach to real FFTs and convolutions of length 2^k . *Computing* **80**, 23–45 (2007).
70. Muñoz, A., Ertlé, R. & Unser, M. Continuous wavelet transform with arbitrary scales and $O(N)$ complexity. *Signal Process.* **82**, 749–757 (2002).
71. Unser, M., Aldroubi, A. & Schiff, S. J. Fast implementation of the continuous wavelet transform with integer scales. *IEEE Trans. Signal Process.* **42**, 3519–3523 (1994).
72. Blake, A. M., Witten, I. H. & Cree, M. J. The fastest Fourier transform in the south. *IEEE Trans. Signal Process.* **61**, 4707–4716 (2013).
73. Frigo, M. & Johnson, S. G. FFTW: an adaptive software architecture for the FFT. In *Proc. 1998 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP'98* 3, (eds Atlas, L., Huang, X., Hermansky, H. & Hwang, J.) 1381–1384 (IEEE, 1998).
74. Wang, E. et al. in *High-Performance Computing on the Intel® Xeon Phi™* (eds Wang, E. et al.) 167–188 (Springer, 2014).
75. Frigo, M. & Johnson, S. G. BenchFFT (2017); <http://www.fftw.org/benchfft/>
76. Stallings, W. *Computer Organization and Architecture: Designing for Performance* (Pearson Education India, 2003).
77. Chakrabarti, C. & Vishwanath, M. Efficient realizations of the discrete and continuous wavelet transforms: from single chip implementations to mappings on SIMD array computers. *IEEE Trans. Signal Process.* **43**, 759–771 (1995).
78. Jamieson, L. H., Mueller, P. T. Jr & Siegel, H. J. FFT algorithms for SIMD parallel processing systems. *J. Parallel Distrib. Comput.* **3**, 48–71 (1986).
79. Mallat, S. A *Wavelet Tour of Signal Processing* (Elsevier, 1999).
80. Mertins, A. & Mertins, D. A. *Signal Analysis: Wavelets, Filter Banks, Time–Frequency Transforms and Applications* (Wiley, 1999).
81. Iatsenko, D., McClintock, P. V. E. & Stefanovska, A. Extraction of instantaneous frequencies from ridges in time–frequency representations of signals. *Signal Process.* **125**, 290–303 (2016).
82. Abbott, B. P. et al. Observation of gravitational waves from a binary black hole merger. *Phys. Rev. Lett.* **116**, 061102 (2016).

83. Burnol, A. et al. Wavelet-based analysis of ground deformation coupling satellite acquisitions (Sentinel-1, SMOS) and data from shallow and deep wells in Southwestern France. *Sci. Rep.* **9**, 8812 (2019).
84. Petrucco, L. Mouse head schema (Zenodo, 2020); <https://doi.org/10.5281/zenodo.3925903>

Acknowledgements

Both authors (L.P.A.A. and E.L.v.d.B.) have received funding for this work from the European Union's Horizon 2020 research and innovation program under grant agreement no. 952095, from the Intrinsic Motivations to Transitional Wearable INtelligent companions for autism spectrum disorder (IM-TWIN) project. The funders had no role in study design, data collection and analysis, decision to publish or preparation of the manuscript.

Author contributions

L.P.A.A. carried out all programming, performed the empirical validation experiments and created the figures. E.L.v.d.B. and L.P.A.A. composed the final paper together. All authors reviewed and agreed the source code and manuscript. Figure 1 includes two subfigures from other sources, which are used with permission of respectively Caltech/MIT/Ligo Laboratory and Nature.

Competing interests

The authors declare no competing interests.

Additional information

Extended data is available for this paper at <https://doi.org/10.1038/s43588-021-00183-z>.

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s43588-021-00183-z>.

Correspondence and requests for materials should be addressed to Lukas P. A. Arts or Egon. L. van den Broek.

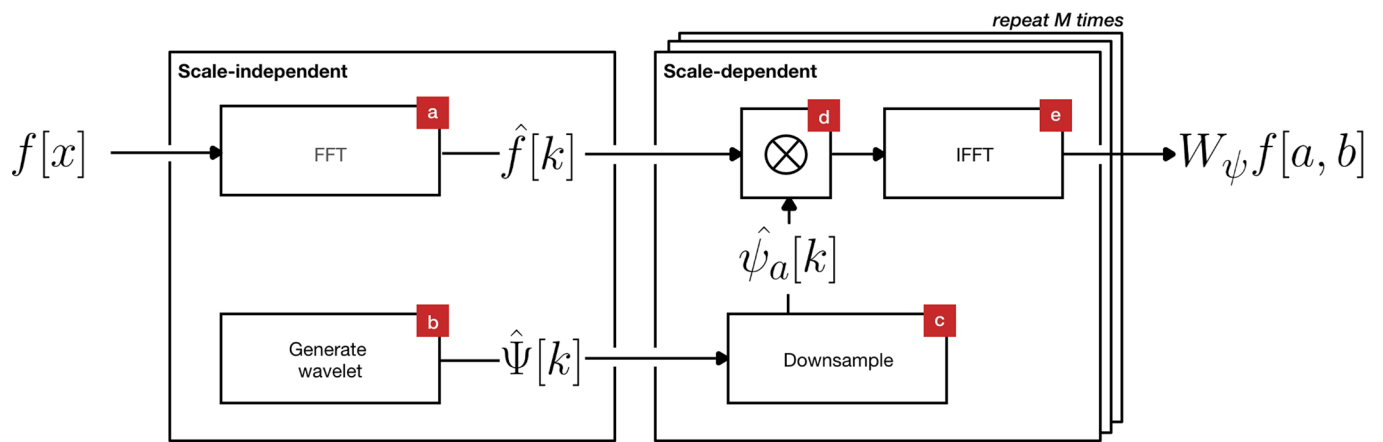
Peer review information *Nature Computational Science* thanks Leontios Hadjileontiadis, Madhur Srivastava and the other, anonymous, reviewer(s) for their contribution to the peer review of this work. Handling editor: Jie Pan, in collaboration with the *Nature Computational Science* team.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.




Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.
© The Author(s) 2022




Extended Data Fig. 1 | Algorithmic implementation of fCWT The algorithmic implementation behind fCWT can be divided into: i) scale-independent and ii) scale-dependent operations. The scale-dependent operations each calculate the wavelet coefficients of a single scale-factor in the final time-frequency matrix. By repeating the scale-dependent part $m = |a|$ times, the time-frequency matrix is build up one row at a time.




Extended Data Fig. 2 | FFTW's interleaving storing format Using an interleaving value format, the Fastest Fourier Transform in the West (FFTW) writes a complex-valued Fourier transform to memory. As the CPU caches adjacent values when accessing memory, accessing the complex and real part only requires single memory access instead of two.



Extended Data Fig. 3 | From mother to daughter wavelet The generation of the daughter wavelet $\hat{\psi}_a[k]$ is done efficiently by downsampling the mother wavelet $\hat{\Psi}[k]$. This eliminates the need for expensive Gaussian calculations in the scale-dependent step. The mother wavelet is only calculated once in the scale-independent step.



Extended Data Fig. 4 | SIMD multiplication fCWT combines the generation of the daughter wavelet and its multiplication with the Fourier transformed input signal together in one Single Instruction, Multiple Data (SIMD) multiplication. As the Fourier transformed input signal is complex-valued, the real daughter wavelet values are copied twice such that SIMD can perform an element-wise multiplication between both buffers. In this example a scale-factor of $a=3$ is used.



Extended Data Fig. 5 | Boundary effects in fCWT and MATLAB With fCWT we perform zero extension to mitigate boundary effects. In contrast, by default MATLAB uses a content dependent mirror extension. In some cases, such an extension strategy can increase boundary effect severity instead of decreasing it as can be seen here.




News

27 January 2022

Publication in Nature Computational Science

Hundred times faster signal processing with new calculation technique



A graphic summary of how the new technique works, translating the signal into a more informative representation. "In short, we look at the signal in a different light!"

By re-implementing an existing calculation technique for signal processing, Utrecht computer scientists have succeeded in speeding up the technique by a

factor of one hundred, without loss of quality. This can enable considerable improvements in countless applications that work with signals or data flows from sensors, from MRI scanners to systems that predict earthquakes. The researchers, [Lukas Arts](#) and [Egon van den Broek](#), are publishing their results today in Nature Computational Science.

“There are two commonly used techniques for signal processing,” Van den Broek explains. “One works very precisely, but is so slow that it is hardly useful. This is why the other technique is often used in practical applications: it is lightning fast, but loses a lot of quality during processing.” The researchers managed to speed up the slow technique by a factor of one hundred, so that it combines the best properties of both techniques: high speed as well as high precision.

Brain signals

The new calculation method can be used in all sorts of applications, says Van den Broek. “Take, for example, a brain-computer interface that allows a fully paralysed person to control a wheelchair. The wheelchair needs to respond immediately to steering signals from the brain, so you have to use the fast algorithm. Consequently, the signal quality is so low that such wheelchairs can only distinguish a few signals: left, right, accelerate and decelerate. With this new calculation technique, you can make the wheelchair recognise many more different signals without compromising on speed.”

Classical computer science


The researchers have made the slower signal processing technique a factor of one hundred faster by combining the underlying mathematics with the latest insights into hardware and software. “This is actually a perfect example of classical computer science,” says Van den Broek. “You take the mathematical basis and make it run on a device in the fastest possible way.” As a result, the new implementation is also a factor of one hundred more energy efficient.

Interactive cuddly toy

“The technology is ready to be used in all kinds of devices in our daily lives,” says Van den Broek. “We’ve made our new algorithm [available open source](#) [↗](#), so anyone can get started right away.” In an [interdisciplinary European project](#) [↗](#), Van den Broek

and Arts are working on an interactive cuddly toy for children with autism, to support diagnosis and therapy. “A perfect application of the technology: the cuddly toy uses sensors to measure very precisely what is happening, and can then respond in real time. This makes it possible to better understand and facilitate social interaction between people, animals and robots.”

Publication

[The fast continuous wavelet transformation \(fCWT\) for real-time, high-quality, noise-resistant time-frequency analysis](#) 

Lukas P. A. Arts and Egon L. van den Broek

Nature Computational Science, 27 January 2022, DOI 10.1038/s43588-021-00183-z

Share



Contact

Egon van den Broek

Researcher

E.L.vandenBroek@uu.nl

Nieske Vergunst

Press officer

N.L.Vergunst@uu.nl

[+31 6 2490 2801](tel:+31624902801)

See also

[Faculty of Science](#)

[Research Institute of Information and Computing Sciences](#)

Go to

Choose a study

[Student information](#)

[Research areas](#)

[Jobs](#)

[Giving](#)

[Services](#)

[Student portal](#)

[Alumni information](#)

[Intranet for employees](#)

[Course planner](#)

[My Timetable](#)

[About UU](#)

[Contact](#)

[Directory](#)

[Disclaimer](#)

[Follow Utrecht University](#)



Utrecht University
Heidelberglaan 8
3584 CS Utrecht
The Netherlands
Tel. +31 (0)30 253 35 50

Door onze redacteur
Laura Wismans

TECHNIEK

Uit dit signaal is veel meer detail te halen

Het kost rekenkracht om een signaal om te zetten in informatie. Als dit in real time moet, gaat veel detail verloren. Maar het kan sneller.



Bij het meten van de elektrische activiteit van de hersenen (EEG) bevat het signaal veel ruis.

AMSTERDAM. Verwerking van elektrische signalen kan wel honderd keer sneller. Dat hebben informatici van de Universiteit Utrecht laten zien in *Nature Computational Science*. Ze hebben daartoe de wiskunde achter een bestaand algoritme herschreven. Het maakt zowel snellere als preciezere analyses mogelijk.

De onderzoekers passen hun vernieuwde rekenmethode toe op signalen uit hersensensoren, maar de techniek is toepasbaar op allerlei signalen, zoals seismische trillingen voor aardbevingsvoorspellingen, netwerkactiviteit voor cybersecurity en audioanalyse voor geluidsfiltering.

Signaalverwerking is een onderzoeksveld op het snijvlak van elektrotechniek en wiskunde, het gaat over verwerking en analyse van informatie uit signalen in de breedste zin van het woord, zowel digitaal als analoog. „Er zijn twee gangbare technieken voor signaalverwerking”, zegt Egon van den Broek, associate professor computational science aan de Universiteit Utrecht en auteur van het onderzoek. „De ene is heel precies, maar zo langzaam dat hij nauwelijks te gebruiken is, zeker niet voor realtimeoepassing. De andere is snel, maar verliest veel kwaliteit bij de analyse. Toch wordt die vaak verkozen, omdat snelheid zo van belang is.”

De 'langzame' techniek, de zogeheten *continuous wavelet transform* (CWT), heeft nu een snelle variant gekregen: *fcWT*, waar de *f* staat voor *fast*. „We hebben op een nieuwe manier gekeken naar de wiskunde die achter de methode zit”, zegt Lukas Arts, promovendus en eerste auteur van het onderzoek. „De formules hebben we opgedeeld in verschillende onderdelen. Het bleek dat berekeningen vaak opnieuw worden gedaan, en wij hebben gekeken waar we informatie uit voorgaande stappen kunnen hergebruiken in plaats van opnieuw berekenen.” Ook hebben ze de dertig jaar oude techniek geoptimaliseerd voor de huidige hardware.

Veel meer details

Ze demonstren hun algoritme met het verwerken van twee soorten hersensignalen. Bij EEG wordt een kap met elektroden op het hoofd geplaatst. Omdat de signalen door de schedel heen moeten, zit er veel ruis in. „We hebben ons algoritme even

snel laten draaien als de gangbare algoritmes voor hersensignaalverwerking, maar we konden eigenlijk veel sneller”, zegt Arts. „De rekenkracht die over was hebben we besteed aan het verhogen van de resolutie met een factor dertig. Er waren veel meer details te onderscheiden.”

„Ook hebben we gekeken naar hersenanalyse, waarbij kleine naaldjes in de visuele cortex worden geplaatst”, zegt Arts. „In één zo'n naaldje zitten vierhonderd sensoren. Voorheen kon een vijfde van de gegevens van een naald in real time bekeken worden, met onze methode konden we zes naalden tegelijkertijd analyseren.”

„Dat hersensignalen realtime geanalyseerd kunnen worden is heel interessant”, zegt Ciska Heida, associate professor biomedical signals and systems aan de Universiteit Twente, niet verbonden aan dit onderzoek. „Diepe hersenstimulatie wordt bijvoorbeeld bij parkinson ingezet. Alleen al tijdens de operatie is het nuttig als hersenfrequenties beter visueel gemaakt worden, zodat je weet dat je

De formules hebben we opgedeeld in onderdelen

Lukas Arts promovendus

in de juiste hersenkern zit en ook direct het effect van stimulatie kunt zien op de verschillende frequenties.” Ook ziet ze nut voor patiënten met het *locked-in* syndroom. „De snelheid waarmee deze patiënten kunnen communiceren is nu nog erg laag.”

Het nieuwe algoritme is ook in bredere zin toepasbaar. Het kan bijdragen aan de groeiende vraag naar uitlegbare algoritmes, de tegenhanger van zelflerende algoritmes. „Zelflerende algoritmes zijn nu hip. Daar gooi je ruwe data in en je krijgt er een mooie uitkomst uit. Maar wat het algoritme precies doet is ongewis, je moet maar hopen dat het bijvoorbeeld niet discrimineert”, zegt Van den Broek. „Wij zeggen nu: als je de voorverwerking veel gedetailleerder doet, kan de *machine learning* simpeler zijn, met misschien wel betere resultaten die ook nog eens uitlegbaar zijn. Bovendien werkt het als je relatief weinig data hebt ook. *Deep learning* is alleen effectief als je heel veel gegevens hebt, dit kan altijd.”



Commentaar

CORONABELEID

Wat Nederland kan leren van de Denen: je data op orde hebben

Met jaloersmakende zelfverzekering hief de Deense regering dinsdag in één klap alle coronamaatregelen op. De Denen wierpen hun mondkapjes af en vierden net als vorig jaar in september bevrijdingsdag. Het nachtleven mocht weer open, een coronapas hoeft nergens meer getoond te worden. Het zou kunnen dat de Denen over een tijdje op hun schreden terugkeren. Maar hun besluitvaardigheid steekt hoe dan ook scherp af bij de schroom waarmee de Nederlandse regering eind januari „na lang wikken en wegen” besloot „de grenzen van het mogelijke” op te zoeken, „op het gevaar af dat ons over een paar weken weer joboobeid wordt verweten”. Waar voor premier Rutte en zijn kabinet vooral „noodkreten”, „acties” en „emotie” redenen waren om horeca, theaters en andere locaties toestemming te geven weer - beperkt - open te gaan, stelde de Deense regering zonder enige aarzeling vast dat Covid-19 niet meer kon worden aangemerkt als „gevaar voor de samenleving” en dat er dus geen reden meer was om burgers beperkingen op te leggen.

De Denen zien het loslaten van alle maatregelen niet als een „risico”, zoals premier Rutte de Nederlandse versoepelingen betitelde, maar als een weloverwogen en rationele beslissing, gebaseerd op grondig verzamelde en bestudeerde data. In de week voor de heropening liep in de Deense ziekenhuizen het aantal patiënten met een positieve test terug tot duizend, op de IC's tot onder de 40. De meesten lagen er primair vanwege een andere aandoening dan corona.

In Nederland is het vertrouwen in het coronabeleid steeds verder gedaald, in Denemarken niet

Van naïviteit kunnen de Denen geenszins worden beticht, de manier waarop zij het coronavirus monitoren, dwingt respect af. Denemarken beschikt over drie keer zoveel teststraten als Nederland (433 versus 155), terwijl het slechts een derde van ons aantal inwoners heeft (5,8 mln versus 17,4 mln). Bepalingen worden vroegtijdig opgespoord en geïsoleerd, hun contacten opgespoord. Van elke positieve test wordt in een centraal lab uitgeplozen om welke virusvariant het gaat. Zo waren de Denen er vorig jaar razendsnel bij toen de Deltavariant een opmars maakte. Vergelijk dat eens met de haperende en overbelaste GGD's in Nederland, en de versplintering aan laboratoria. Wie de afgelopen weken een testafspraak wilde maken stuitte vaak op de mededeling dat er geen plek meer was. Van bron- en contactonderzoek lijkt niet veel meer te komen. Ook de registratie van de testresultaten bij het RIVM loopt spaak. De achterstand liep op tot boven de 100.000 positieve tests. Nederland is het zicht op het virus kwijt, zoveel is duidelijk.

Het is niet zo dat de Denen met hun grondige monitoring van het virus Covid-19 volledig hebben uitgebannen. In de week voor het land van het slot ging waren er nog ruim 43.000 nieuwe besmettingen per dag. Maar de Denen weten wel met meer zekerheid wanneer ze echt moeten ingrijpen en zijn daar ook beter op voorbereid. Denk aan de boostercampagne, die in Nederland pas tergend laat op gang kwam.

In Nederland is het vertrouwen in het coronabeleid steeds verder gedaald, in Denemarken niet. Ook als de Deense overheid „joboobeid” voert - regels laten varen en ze een paar maanden toch weer invoeren - volgen de meeste Denen op wat er van hen wordt gevraagd. Door het steeds vroegtijdig opsporen en isoleren van besmettingen waren er minder beperkingen nodig voor de rest van de bevolking. Als Nederland dat ook wil, vergt dat forse investeringen in een test-, lab- en vaccinatiestructuur die snel kan worden op- en afgeschaald. Maar dat kan wel het vertrouwen in het coronabeleid een boost geven, en dat is ook wat waard.

BRIEVEN



FOTO JOHANNES VANASSE

ABORTUSWET

Wachtijd is nodig vanwege de onomkeerbaarheid

De kop *„Wat ik vond, was blijkbaar niet goed genoeg voor de wet”*, boven het artikel (2/2) over abortus en bedenktijd, klopt niet. Wat mevrouw Chakh vond, haar wens een ongewenste zwangerschap te beëindigen, was wél goed genoeg voor de wet, maar niet met een dezelfde dag. We hebben in Nederland gelukkig een goede abortuswetgeving met een verplichte wachtijd van vijf dagen die impulsbeslissingen voorkomt. Ik werk ruim dertig jaar als huisarts, ben groot voorstander van het recht op abortus, maar ook een groot voorstander van deze wachtijd. Ik heb in de loop der jaren genoeg vrouwen gesproken die (achteraf) blij waren met die wachtijd. Het gaf ze tijd om te komen tot een meer weloverwogen beslissing, welke deze ook was.

Wanneer de spanning en angst door je keel gieren bij het ontdekken van een ongewenste zwangerschap lijkt een beslissing heel makkelijk. Het echte probleem ontstaat pas als er inderdaad dezelfde dag een abortus zou zijn verricht en de betrokken vrouw daar de volgende dag spijt van heeft. Dan is er geen weg terug. Daarom heeft de wachtijd naar mijn mening niets met betutting te maken maar alles met de onomkeerbaarheid van deze beslissing die te groot is om impulsief te nemen. Ik begrijp dat de wachtijd als overbodig kan voelen als je écht klaar bent met je denkwerk (en met het overleg met je partner); dit ongemak staat in geen verhouding tot de ramp die spijt kan geven. Kameriesden, bedenktijd voor u de wachtijd afschaft.
Ton Jansen Rijkevoort

OEKRAÏNE

De NAVO beslist zelf

Uitentreuren horen we dat de Oekraïners het recht hebben om zelf hun lot te bepalen. Dit argument wordt gebruikt tegen degenen die de toetreding van Oekraïne tot de NAVO een slecht idee vinden. Die zouden zwichten voor de agressie van Poetin. Het betreft een sterk staaltje van framing. Natuurlijk moeten de Oekraïners zelf kunnen beslissen wat ze doen, en natuurlijk mogen zij het NAVO-lidmaatschap aanvragen. Maar dat betekent niet dat de NAVO daarmee moet instemmen. We hebben hier te maken met twee geheel verschillende zaken. De NAVO heeft een andere verantwoordelijkheid en andere belangen dan de regering van Oekraïne. Wat doen we als Egypte NAVO-lid wil worden? Zijn het dan ook de Egyptenaren die hierover moeten beslissen?
Daan Brrouwer Amsterdam

MEDIA

Gevaarlijke concentratie

De vraag of het wel wenselijk is dat 90 procent van de Nederlandse dagbladenmarkt in handen is van slechts twee bedrijven, wordt door ceo Gert Ysebaert van een van die bedrijven, Mediahuis, ontwijkend beantwoord. Er zou veel journalistieke vrijheid zijn (*„We willen een Europese mediaspeler worden”*, 1/2). Ik weet het antwoord wel: het is zeer onwenselijk. Ik snap werkelijk niet waarom het Commissariaat voor de Media niet tegen eerdere overnames optreden heeft. Als morgen

De NAVO heeft andere belangen dan de regering van Oekraïne

twee aan Xi verbonden Chinezen of aan Poetin verbonden Russen deze bedrijven opkopen, dan zitten we met de gebakken peren. Deze schaalvergroting geeft vroeg of laat ellende.
Rienk Koopmans Brielle

MEDIA (2)

Verhoudingen zoek

Mediahuis-topman Gert Ysebaert is terughoudend in het sluiten van deals met Facebook en Google (*„We willen een Europese mediaspeler worden”*, 1/2). Dat is goed om te horen, maar toch weinig geruststellend. De winst van Mediahuis was in 2020 58,6 miljoen euro. Dat is absoluut een knappe prestatie, maar Googles moederbedrijf Alphabet verdiende in dat jaar 40 miljard dollar. De *terms of trade* zijn volstrekt zoek. Moet de (Europese anti-trust)wetgever hier niet een helpende hand bieden, zodat Mediahuis überhaupt nooit in de verleiding kan komen?
Tido van Wieringen Haarlem

ZORGOPLEIDINGEN

Er verandert wel wat
Folkert Jensma schreef over de

Advertentie

ONTDEK DE DIGITALE SCHATKIST VAN NRC

Doorzoek alle artikelen vanaf 1 oktober 1970 in het NRC-archief op onderwerp of op datum.



Ga naar nrc.nl/archief

