

GGIR: A Research Community–Driven Open Source R Package for Generating Physical Activity and Sleep Outcomes From Multi-Day Raw Accelerometer Data

Jairo H. Migueles

PROFITH Research Group, University of Granada University of Leicester and NIHR Biomedical Research Centre

Alex V. Rowlands

Florian Huber

Netherlands eScience Center

S  verine Sabia

Inserm, University College London

Vincent T. van Hees

Netherlands eScience Center

Recent technological advances have transformed the research on physical activity initially based on questionnaire data to the most recent objective data from accelerometers. The shift to availability of raw accelerations has increased measurement accuracy, transparency, and the potential for data harmonization. However, it has also shifted the need for considerable processing expertise to the researcher. Many users do not have this expertise. The R package GGIR has been made available to all as a tool to convert multi-day high resolution raw accelerometer data from wearable movement sensors into meaningful evidence-based outcomes and insightful reports for the study of human daily physical activity and sleep. This paper aims to provide a one-stop overview of GGIR package, the papers underpinning the theory of GGIR, and how research contributes to the continued growth of the GGIR package. The package includes a range of literature-supported methods to clean the data and provide day-by-day, as well as full recording, weekly, weekend, and weekday estimates of physical activity and sleep parameters. In addition, the package also comes with a shell function that enables the user to process a set of input files and produce csv summary reports with a single function call, ideal for users less proficient in R. GGIR has been used in over 90 peer-reviewed scientific publications to date. The evolution of GGIR over time and widespread use across a range of research areas highlights the importance of open source software development for the research community and advancing methods in physical behavior research.

Keywords: accelerometry, open source, physical behaviors, sleep, software

Human physical activity and sleep are popular areas of research because of their important role in health outcomes (He, Zhang, Li, Dai, & Shi, 2017; Lee et al., 2012). Physical activity and sleep have traditionally been quantified with diaries and questionnaires, but wearable sensors have gained momentum since the 1990s. In the beginning, wearable movement sensors (i.e., accelerometers) typically performed onboard signal processing and only stored derived output to reduce battery consumption and memory requirements. However, following a general movement towards

more transparent and open science, and thanks to technological evolution towards smaller, cheaper, and power efficient sensors, accelerometers now tend to store ‘raw’ data for offline processing and analysis. The data recorded are typically expressed in gravitational acceleration (g) because this is the reference point for acceleration value calibration, reflecting both the movement and gravitational component (van Hees et al., 2013). However, this technological advance is counterbalanced by the large amount of data collected per measurement (typically $2 \cdot 10^8$ data points per week of measurement) and the necessity to process the data to obtain meaningful variables that can be used in standard statistical analysis and software. Many physical activity and sleep researchers do not have the expertise necessary to process and analyze raw accelerometer data. The GGIR software presented in this paper facilitates the processing and extraction of insightful physical activity and sleep variables of the data collected with these so called raw data accelerometers from three widely used sensor brands (Wijndaele et al., 2015). The individual algorithms that are embedded in GGIR have been described across a number of published papers. The fast advances in wearable sensor technology over the last decades comes with the price of mandatory development of scientific software to ensure a good valorization of the newly available sensors (Seinstra, Wallom, & Keahey, 2015).

Migueles is with the PROFITH (PROmoting FITness and Health through Physical Activity) Research Group, Sport and Health University Research Institute (iMUDS), Dept. of Physical Education and Sports, Faculty of Sport Sciences, University of Granada, Spain. Rowlands is with the Diabetes Research Centre, University of Leicester, Leicester, United Kingdom; NIHR Biomedical Research Centre, Leicester, UK; and Alliance for Research in Exercise, Nutrition, and Activity (ARENA), Sansom Institute for Health Research, Division of Health Sciences, University of South Australia, Adelaide, Australia. Huber and van Hees are with the Netherlands eScience Center, Amsterdam, The Netherlands. Sabia is with the Centre of Research in Epidemiology and Statistics—Inserm U1153, EpiAgeing, Universit   de Paris, Paris, France; and the Dept. of Epidemiology and Public Health, University College London, London, United Kingdom. Migueles (jairohm@ugr.es) is corresponding author.

However, scientific software instruments need to be subjected to a peer review process as it is the case for other methodological components (e.g., algorithms and study protocols). Previous publications related to GGIR focused on specific algorithms, such as sleep detection, and their validity. However, those publications did not focus on GGIR as a generic piece of scientific software that connects all these algorithms and adds a range of essential extra functionalities (e.g., time zone and daylight-saving time handling) to provide value to an entire research community far beyond the specificities of those studies. Therefore, this paper aims to provide a one-stop overview of the GGIR package, the papers underpinning the theory of GGIR, and how research contributes to the continued growth of the GGIR package.

How Open Source Software Can Contribute to Advances in the Field of Physical Activity Measurement

GGIR contributes to scientific discovery by enabling researchers to study (human) physical activity and sleep using accelerometers without pre-required programming expertise. GGIR is appropriate for use across a wide variety of study designs (e.g., variations in measurement duration, in sample frequency, instructions given to the participant) and study populations.

Applications of GGIR have been reported in over 90 peer reviewed journal publications since its first release in 2013, with 24 in 2017 and 48 in 2018 based on looking up the citations to the key journal publications underlying GGIR (see also: <https://github.com/wadpac/GGIR/wiki/Publication-list>). Additionally, nine methodological papers were written to motivate and evaluate parts of GGIR. Previous accelerometer software has been commercial (e.g., Actilife) and/or restricted to one brand of accelerometer (e.g., Actilife, GENEActiv PC combined with Excel macros). GGIR facilitates the processing and analysis of data from three of the most widely used brands of research-grade movement sensors (GENEActiv by ActivInsights Ltd, ActiGraph by Actigraph LLC, and Axivity by Axivity Ltd) using open source generic brand agnostic methods, potentially providing a means for harmonization of data from large surveys globally (Rowlands, Mirkes, et al., 2018). Further, GGIR is continually updated to include innovations developed by the software team and/or users expediting the application of novel analytics in research (Rowlands, Edwardson, et al., 2018; Rowlands, Mirkes, et al., 2018; van Hees et al., 2014, 2015, 2018).

GGIR is very much a community-driven development as testified by the multiple contributors from both health research and technological backgrounds, the fact that many publications that use GGIR are not co-authored by the development team, and the existence of a support and maintenance service by V. van Hees where GGIR users can hire van Hees' time as freelancer to help address specific needs from the user-community (www.movementdata.nl). This service fueled a range of package upgrades in 2018, and is one of the possible ways to sustain open source software like GGIR. The availability of a user forum makes it possible for users to reach out to each other (see Table 1).

Previously, accelerometers were typically either used to assess waking physical activity or sleep and circadian rhythms, but rarely tailored for both research areas. The new generation of accelerometers, worn night and day, allows the measurement of both physical activity and sleep using the same tool. Moreover, in several studies, such as UK Biobank (Doherty, Jackson, et al., 2017), a sleep diary was not used to define a sleep window as commonly done in sleep research. The GGIR package allows differentiation of waking from sleep periods and provides sleep quality parameters. Using the GGIR package, the Colaus study reported better sleep efficiency among those more physically active (Gubelmann, Heinzer, Haba-Rubio, Vollenweider, & Marques-Vidal, 2018). Several papers using UK Biobank data have now been published and show the advantage of using accelerometer instead of questionnaire data to identify, for example, the genetics of sleep (Jones, Lane, et al., 2019; Jones, van Hees, et al., 2019; Lane et al., 2018).

There are further examples of research studies that were facilitated by the use of GGIR. Results from the Whitehall II study showed 1) the association between physical activity and body mass index was much stronger when using accelerometer data rather than questionnaire data (Sabia et al., 2015); 2) obese people with an unfavorable metabolic profile had a lower level of physical activity than those with a healthy metabolic profile, which was not evident when using questionnaire data (Bell et al., 2015); 3) the association between moderate and vigorous physical activity and healthy aging was evident whether the activity was performed in short (<10 minutes) or long (≥10 minutes) bouts (Menai et al., 2017). Results from the UK Biobank also showed 1) people with cardiometabolic disease are less physically active and tend to engage more in sedentary activities that last more than 30 minutes (Cassidy et al., 2018), and 2) short bursts of very high intensity physical activity are associated with better bone health in pre- and post-menopausal women (Stiles, Metcalf, Knapp, & Rowlands, 2017).

Table 1 Key Metadata on the Current Version of GGIR

Current code version	1.6-7
Permanent link to code/repository used of this code version	https://github.com/wadpac/GGIR
Legal code license	LGPL (≥ 2.0, < 3)
Code versioning system used	Git
Software code languages, tools, and services used	R, C++, Travis-CI
Compilation requirements, operating environments & dependencies	64-bit operating system & R environment version 3.2.3 and up (64-bit) & R packages: MASS, signal, zoo, mmap, bitops, matlab, GENEActiv, tuneR, testthat, covr, knitr, rmarkdown, data.table, Rcpp
Link to developer documentation/manual	https://cran.r-project.org/web/packages/GGIR/GGIR.pdf
User forum	https://groups.google.com/forum/#!forum/rpackageggir
Works on Operating Systems	Windows, Linux, and OSx

Note. LGPL = Lesser General Public License.

A Typical Experimental Setting

Studies vary in size from a few dozen participants in clinical or methodological studies (Bachasson, Landon-Cardinal, Benveniste, Hogrel, & Allenbach, 2017), a few thousand in epidemiological cohorts (da Silva et al., 2014; Sabia et al., 2014), and a hundred thousand in biobanks like UK Biobank (Doherty, Jackson, et al., 2017). The researcher configures an accelerometer with the desktop software supplied by the accelerometer manufacturer. Next, the accelerometer is given or sent to the participant who wears it on their wrist or other body location (depending on instructions) day and night for usually seven consecutive days, although different measurement periods have also been reported. After the period of wear, the participant returns the accelerometer either in person or by post. The data file is then downloaded with the same desktop software from the accelerometer manufacturer. The file size typically ranges between 0.5 and 1.5 GB depending on specific format, sensor brand, and recording duration. GGIR can either run on a local computer or be integrated in parallel processing tasks on a computing cluster when working with large numbers of datafiles. GGIR generates reports in csv-format that can then be loaded in the researcher's preferred statistical software (e.g., R, SAS, SPSS, or Stata).

Related Work

Software. A few other software tools exist to work with multi-day raw accelerometer data (Table 2). Actilife (ActiGraph, Pensacola, FL, USA) is a closed-source commercial software designed for the accelerometer hardware developed by the same company. The GENEActiv PC software by the developer of the GENEActiv (ActivInsights Ltd, Kimbolton, UK), in combination with freely available Excel macros (available from <https://open.geneactiv.org>), is designed for the accelerometer hardware developed by the same company. GENEAcclassify is an open source R package primarily aimed at facilitating the segmentation and classification of accelerometer data produced by the GENEActiv accelerometer (Campbell, Gott, Langford, & Sweetland, 2018). OMGUI (Open Movement, Newcastle University, UK) by Dan Jackson and colleagues is an open source GUI implemented in C# and developed for the open source hardware AX3 (Jackson, 2018b). In addition to the monitor's own binary .CWA format, OMGUI can also create and load brand agnostic csv and audio WAVE file formats (Jackson, 2018a). Pampro by Tom White (Cambridge, UK) is implemented in Python, available on GitHub,

and designed to process data from multiple wearable sensor brands and data formats using methodologies replicated from publications, including the publications describing parts of GGIR (White, 2018). Pampro does not include an open source license and we have been informed by Pampro's developer that this is intentional. The lack of an Open Source license prohibits using a software without explicit permission from the copyright holder. BiobankAccelerometerAnalysis (Doherty, Hollowell, Willetts, Jackson, & Hammerla, 2017) by Aiden Doherty (Oxford, UK) and colleagues is an open source tool designed to provide a minimum set of basic outcomes for the UK Biobank accelerometer data, utilizing some of the key algorithms from GGIR (Doherty, Jackson, et al., 2017). The BiobankAccelerometerAnalysis software has recently been enhanced with activity type classification for Axivity AX3 .cwa accelerometer files (Willetts, Hollowell, Aslett, Holmes, & Doherty, 2018).

Literature. The main publications underlying GGIR describe 1) the potential of using wrist-worn raw data accelerometers for estimating human daily energy expenditure (van Hees et al., 2011), 2) how raw acceleration data can be meaningfully aggregated (van Hees et al., 2013), 3) a method to calibrate the acceleration signals based on the recording itself with no need for additional data collection (van Hees et al., 2014), 4) methods for sleep detection when the accelerometer is worn on the wrist with (van Hees et al., 2015) and without (van Hees et al., 2018) use of a sleep diary, and 5) comparisons with other research methods (Innerd, Harrison, & Coulson, 2018; Rowlands, Cliff, et al., 2016; Rowlands, Yates, Davies, Khunti, & Edwardson, 2016; Sabia et al., 2014; Stiles et al., 2017).

Software Description

Software Architecture

R package GGIR has been released with an Open Source LGPL 2 license on CRAN since 2013, the central repository for R packages, and on GitHub [www.github.com/wadpac/GGIR] since December 2016. GGIR can run on Windows, Linux, and OSx (this and additional key facts about GGIR are summarized in Table 1). The package comes with the following core functionalities: load data; extract signal metrics (also called signal features); detect when the sensor was not worn and impute these periods if requested by the researcher; detect the sleep period time window and sleep episodes within it; segment the data according to conventional heuristic threshold techniques; specify which parts of the recording

Table 2 Related Software

	GGIR	Actilife	GENEActiv PC + macros	GENEAcclassify	OMGUI	Pampro	Biobank- Accelerometer- Analysis
Executable version online	YES	NO	YES	YES	YES	YES	YES
Source code available online	YES	NO	NO, but YES for macros	YES	YES	YES	YES
Open source software (OSS) license	LGPL (copyleft) OSS license	No OSS license	NO OSS license	GPL (copyleft) OSS license	Custom permissive OSS license	No OSS license	BSD 2-Clause License (permissive)
Primary programming language	R	Unknown	Unknown	R	C#	Python	Java and Python

Note. LGPL = Lesser General Public License; GPL = General Public License.

should be considered based on the researchers knowledge about the study design (e.g., participant started wearing the accelerometer for a certain number of hours after starting the record); and finally to store convenient summary reports in csv and pdf format (Figure 1).

The package is split in five parts that group functionalities in logical processing order and in line with how the package historically evolved. The parts are numbered from 1 to 5 and the corresponding function names refer to the part number: *g.part1*, *g.part2*, *g.part3*, *g.part4*, and *g.part5*. The parts should be executed sequentially with milestone data automatically being stored in between parts to facilitate re-processing parts without having to go back to the original raw input data. One shell function allows for interacting with the five parts and all underlying functionalities from one central point: *g.shell.GGIR*. The shell function takes all arguments from the five parts mentioned above. In this way, the users can interact with GGIR from one single function call to function *g.shell.GGIR* and easily share their call to allow for replicating the analysis on a different dataset. The duration of GGIR depends on computer specifications, input arguments, and data characteristics. Part 1 (*g.part1*) is the most time consuming taking up at least 80% of the processing time and lasting around 10 minutes when applied to a seven-day data file using GGIR's default argument settings.

The code builds a folder structure with a depth of two to store the milestone data per participant in .RData files (including collections of data.frame and vector objects) and the analysis reports per dataset (see Figure 1). The analysis reports are in .csv format to give the user the flexibility to process their quantitative results in other statistical or data processing environments. There are many variables stored in the reports, an introduction to these variables can be found in the package vignette (<https://cran.r-project.org/web/packages/GGIR/vignettes/GGIR.html>), while a more elaborate discussion is found in the package manual (<https://cran.r-project.org/web/packages/GGIR/GGIR.pdf>).

In GitHub the code is supported by continuous integration with Travis-CI (<https://travis-ci.org/wadpac/GGIR>). Starting in 2017 we have adopted the habit of writing unit tests, but as it was not done before 2017 not all parts of the code are covered by tests yet and current test coverage is 63%. In addition to unit tests, the development has typically gone in close collaboration with GGIR end-users who checked code validity by close examination of the package output. We have created one video to introduce GGIR (<https://youtu.be/RuFBCAqFJ2M>) and a second video to provide a visual summary of how GGIR is typically used (<https://youtu.be/S8YPTYNWdU>).

Software Dependencies

Most of the code is written in R, with a small part in C++ needed for reading the binary data from the AX3 accelerometer brand using R package *Rcpp* as a dependency (Eddelbuettel, 2013). R packages *bitops* and *matlab* are used to enable reading the binary data from the Genea accelerometer (the non-commercial precursor to the GENEActiv, Unilever Discover, UK). Dependencies are R packages *GENEAread* and *mmap*, which are used for reading GENEActiv accelerometer binary data. Finally, R package *tuneR* (Ligges, Krey, Mersmann, & Schnackenberg, 2018) is used to read wav format data, which is an optional export format for AX3 accelerometer data. The R package *signal* is used for frequency filtering (Ligges et al., 2014), R package *zoo* is used for calculating a rolling median, sum, and mean (Zeileis & Grothendieck, 2005),

and R package *data.table* is used at some steps to efficiently handle large amounts of tabular data.

Software Functionalities

Part 1 (g.part1)

The *g.part1* function searches the specified data directory recursively for files that could possibly represent acceleration data. Next, it automatically detects which accelerometer brand the file belongs to, the data format in which it is stored (.csv, .bin, .cwa, .wav), and extracts the file header using appropriate functions. Next, function *g.calibrate* is used to investigate calibration error which results in proposed correction coefficients as motivated and described in van Hees et al., 2014 (van Hees et al., 2014). Then, metrics essential for sensor wear detection, physical activity and sleep analysis are extracted from the raw data. Here, the user can choose one or multiple aggregation metrics out of a collection of most common metrics (e.g. ENMO) and control the window size over which the metrics are calculated. Additionally, a standard set of metrics is extracted per long time window (default 15 minutes) which are needed for the detection of accelerometer non-wear. The data loading and metric extraction takes place in approximately 24-hour blocks since putting the full file content in computer memory may not always be possible. The data window size is modifiable in case the 24-hour blocks are still too large. Additionally, the code evaluates available memory throughout the processing and shortens the window by 20% when available memory is getting too low. At the end of part 1 the signal metrics are stored as milestone data in an RData file (using a filename corresponding to the input accelerometer data file). For most of the analysis we use POSIX format for timestamps, which is the default in R, but for the exported time series we transform these to ISO8601 format to facilitate usage in other software environments. Further, GGIR takes into account day saving time and time zone. The user is expected to explicitly provide time zone of the recording, this to avoid confusion about where the experiments took place, which may not always be the time zone of the machine on which the analyses are run, or the default time zone assumed by GGIR (UTC: Europe/London).

Part 2 (g.part2)

As part of *g.part2* function *g.impute* takes the milestone output from *g.part1*, optionally identifies unreliable signal sections (e.g., monitor not worn or signal clips near its extreme), and replaces these sections by imputed values (average of same time point on all other days of the measurement). Then, function *g.analyse* performs a descriptive analysis of the output and summarizes it per measurement, per day of measurement and conveniently per day type (i.e., weekdays and weekend days separately) as well as per segment(s) of each day (specified by user with argument *qwindow*). For example, if *qwindow* has value 'c(0,9,12,17,24)' the summaries will be derived for the time windows 0:00–24:00, 0:00–9:00, 9:00–12:00, 12:00–17:00, and 17:00–24:00. This could for instance be useful in research on the level of physical activity during specific segments of the day in children and workplace interventions in adults. Examples of summaries generated per time segment are time spent in specific acceleration magnitude ranges (bins), average acceleration metric value, and the timing of the least and most active five-hour window of the day (number of hours can be set by user). The user can also

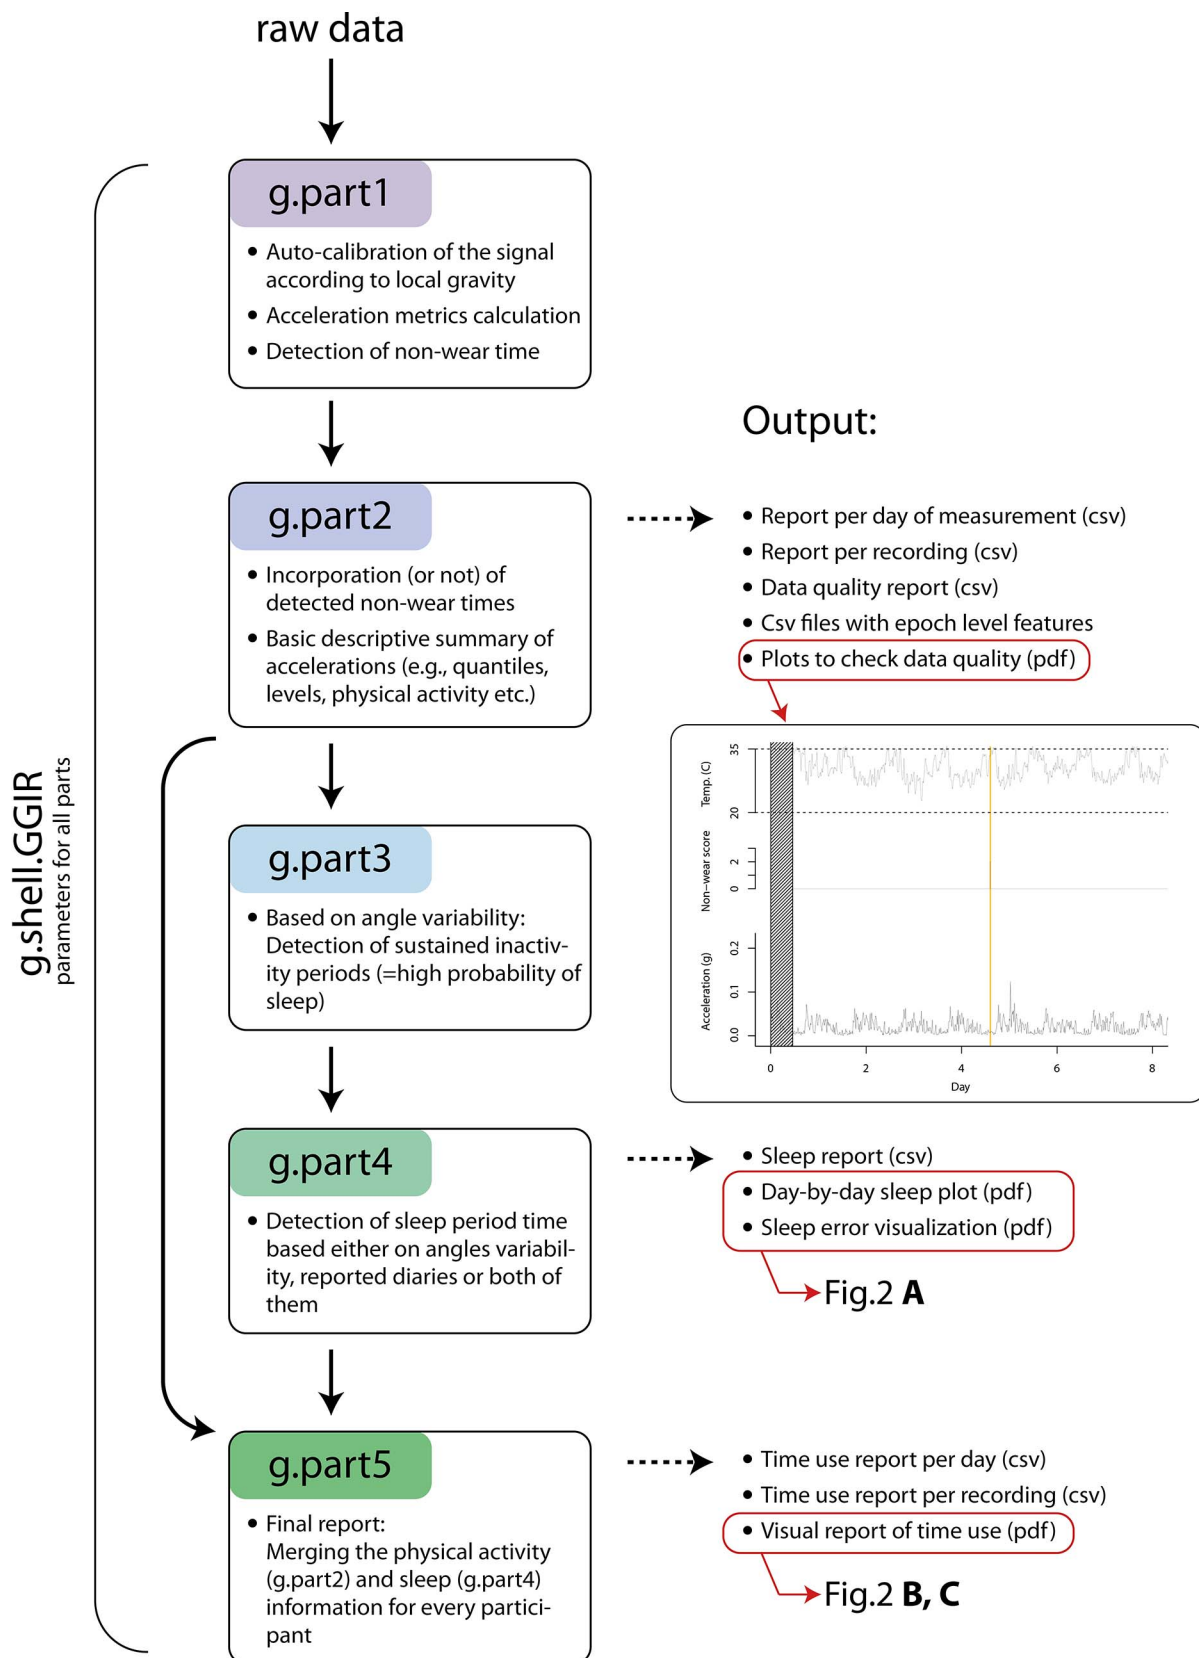


Figure 1 — Overview of main steps and output in GGIR workflow.

provide important knowledge about the experimental design, which is accounted for when generating the descriptive summaries (e.g. instruct GGIR to ignore the first hour of the measurement or all data before the first and after the last midnight of the recording). In addition, the descriptive summaries take into account the quality of the measurement. For example, the user can specify the necessary minimum number of valid hours per day. Days with fewer valid hours of data will be ignored in the person level descriptive summaries.

The reason why *g.part1* and *g.part2* are not merged is because *g.part1* takes much longer to run and involves only minor decisions of interest to the physical activity or sleep researcher. Function *g.part2* on the other hand is relatively fast and comes with all the decisions that directly impact on the variables that are of interest to the researcher. Therefore, the user may want to run *g.part1* overnight or on a computing cluster, while *g.part2* can then be the main playground for research.

Part 3 and 4 (*g.part3* and *g.part4*)

The functions *g.part3* and *g.part4* provide functionality for estimating sleep when the accelerometer was worn on the wrist as described in van Hees et al. (van Hees et al., 2015, 2018). Although these functions can be applied to accelerometer data from any wear location, currently no scientific literature exists to support the interpretation for attachment locations other than the wrist. In short, *g.part3* detects time episodes with a sustained lack of change in arm angle, which are referred to as sustained inactivity bouts. The user can configure the parameters used for this and can specify multiple parameter values to facilitate comparisons. The *g.part3* function only generates milestone data to facilitate the work done in *g.part4*. Output from *g.part3* for example includes the start and end times of the sustained inactivity bouts, and estimated start and end of the sleep period time window. Then, the *g.part4* function gives the user the option to either use a sleep diary or to rely on the estimated sleep period time window from *g.part3* (van Hees et al., 2018). Another option is to specify a general time window when individuals within the study population are assumed to be in bed, which is probably more naïve and primarily implemented for reference purposes. Sustained inactivity bouts that occur within the sleep period time window are considered sleep episodes, and sustained inactivity bouts outside the sleep period time window are considered rest, potential nap period, or undetected short episodes of monitor non-wear time. Further, *g.part4* offers the user the option to exclude the first and last night. Both *g.part3* and *g.part4* store a record of the amount of available valid and invalid data per night, and an indicator of whether sleep diary data was available for each specific night used. To facilitate a relatively quick inspection, sleep detection plots of the classification are exported as pdf files. To facilitate identifying obvious mistakes in sleep diary data *g.part4* provides the option to visualize the differences between accelerometer-based estimates and sleep diary; this was for example used in van Hees et al. (van Hees et al., 2015) to identify a dozen of problematic nights out of 27000 nights (Figure 2, panel a).

Part 5 (*g.part5*)

The *g.part5* function takes the output from *g.part2* and *g.part4* to describe time spent in 16 time use classes, composed of night-time or daytime behaviour; sustained inactive or sleep, other inactive, light, moderate, or vigorous physical activity behavior; and

un-bouted, short, medium, or long bouts of behavior. The number of 16 classes is the default, this can be adjusted. Next, the time series of epoch level classes are segmented by day based on one or both available definitions of a day: 1) From waking up to waking up the next day, which means that a full night of sleep is included per day and that the duration of days vary; or 2) From midnight to midnight which equals a calendar day, meaning that a night of sleep is likely to be split between days. In the latter definition days are by default assumed to end at midnight, but the user also has the option to change the hour on which the day ends. From these daily segments *g.part5* calculates the time spent in these classes, the number of occurrences of these classes, and the average acceleration within the classes.

The thresholds to describe level of physical activity level intensity (light, moderate, and vigorous) have their origin in energy expenditure research, but in GGIR these levels should be interpreted as thresholds to mark ranges in acceleration values. However, if the user specifies the thresholds wisely he/she may be able to interpret the levels as indicators of energy metabolism. The concept of behavioral bouts may also need clarification: *g.part5* calculates the total time spent in behavioral bouts, which is a time segment that meets user-specified criteria on the range of acceleration, the percentage of time during which those criteria on the range need to be met, and the maximum duration of breaks allowed to define a bout. The user selects one of four metrics for bout calculation using argument *bout.metric*. This functionality overlaps somewhat with *g.part2* which also calculates time spent in bouts, but *g.part5* does it in a much more advanced way with also attention to inactivity bouts, light activity bouts, as well as un-bouted behavior and sleep. All this information is stored in csv files and visual summaries are stored in pdf files (Figure 2, panels b and c). The *g.part5* function offers the user a large freedom to explore multiple parameters simultaneously (thresholds, minimal and maximal bout durations and percentage criteria to define a bout) resulting in potentially hundreds or even thousands of output variables to enable the study of physical behavior.

Impact

The main value of GGIR is that it offers a broad set of functionalities ranging from data quality handling to 24 hours/7 days time use characterization of physical activity and sleep utilizing literature supported methods, and is study population agnostic making it suitable for a wide range of research areas. GGIR can be operated without significant prior programming expertise.

At the same time, the user experiences freedom in the specification of input arguments and the selection of output variables. Being fully open source, GGIR can be adapted and extended according to the needs of the respective research project. GGIR is currently being used by the physical activity and sleep research communities, with over 90 peer-reviewed journal papers published using the software since 2013, with the number published per year increasing rapidly.

By facilitating analysis of raw accelerometer data, GGIR enables the analysis of objective measures of physical activity and sleep and encourages raw data accelerometers to be used more readily in research studies that aim to understand the importance of physical activity and sleep for human health. Moving away from closed commercial accelerometer software and self-report questionnaires to assess physical activity and sleep will improve knowledge, allowing better translation in public health recommendations.

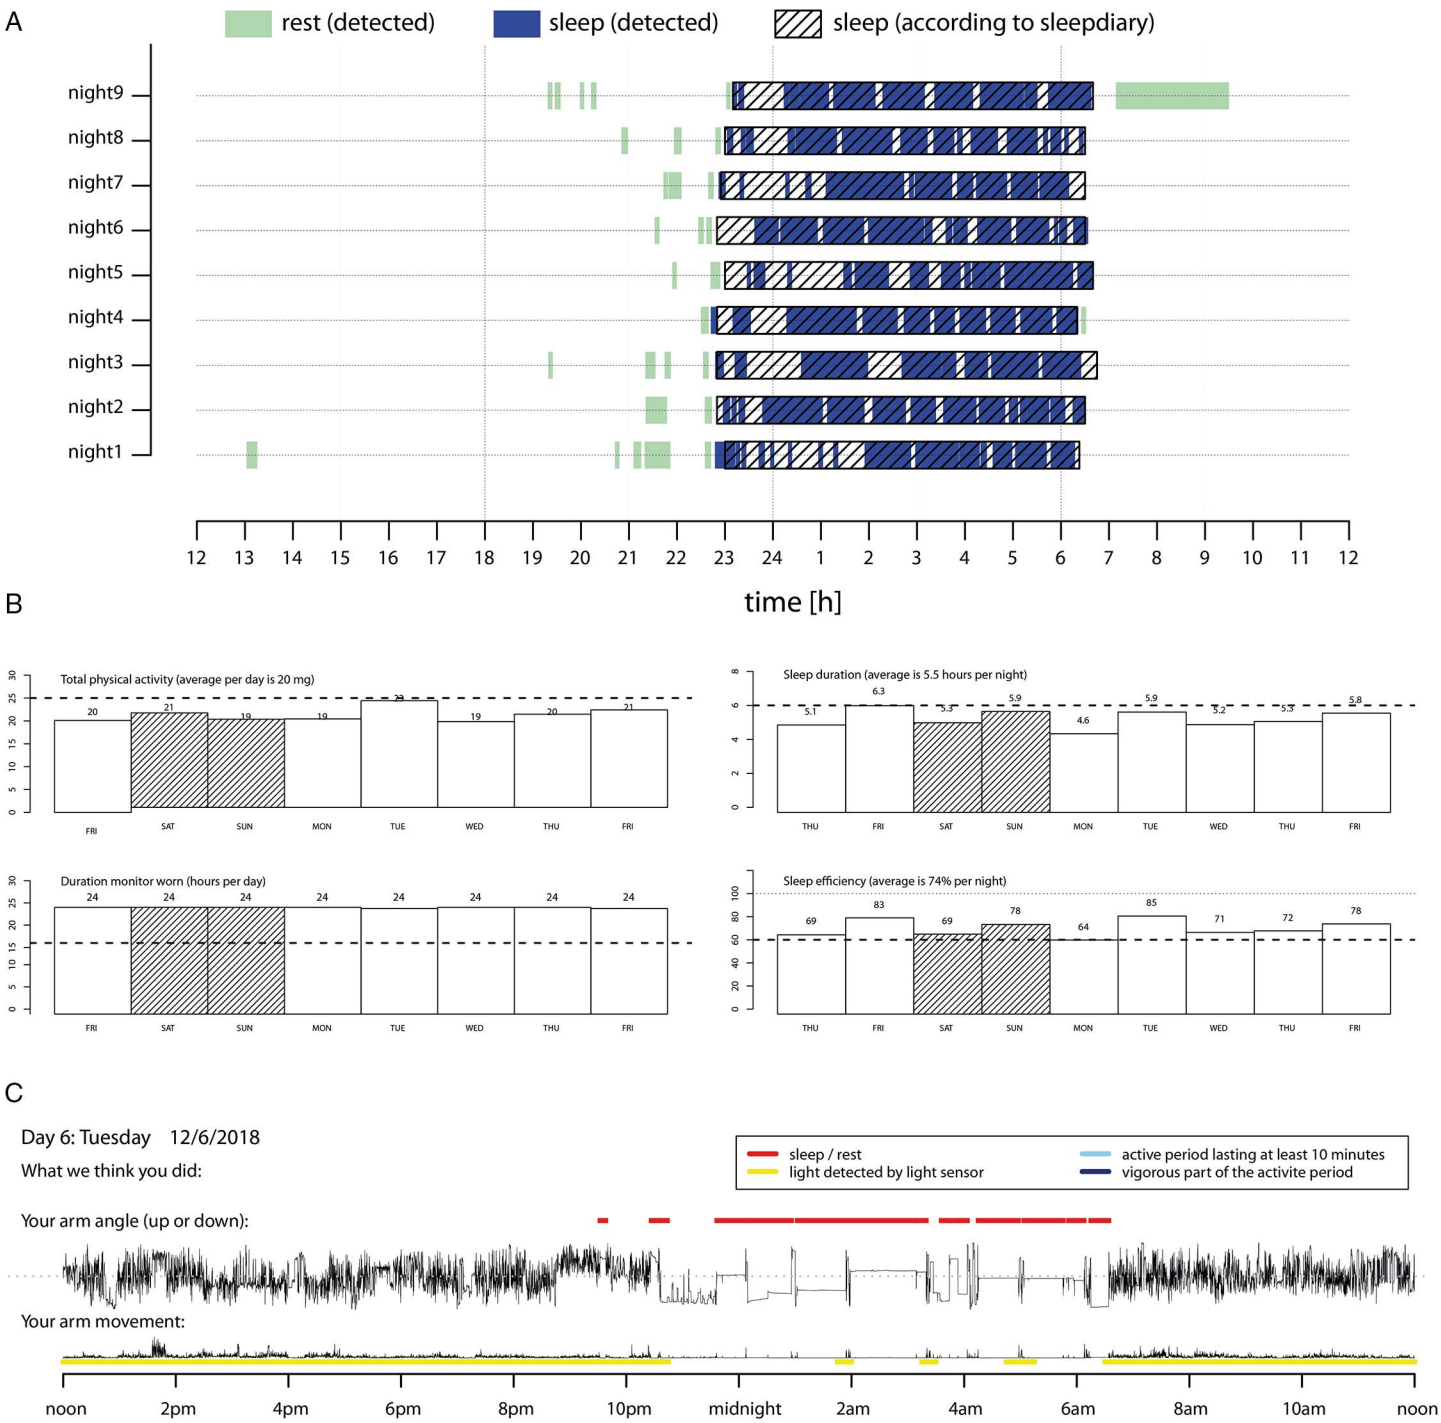


Figure 2 — Visual output reported by GGIR. a) Rest and sleep detected from acceleration features and sleep diary (g.part4); b) Bar plots with information on key physical activity and sleep variables (g.part5); c) Visual summary of the physical activity and sleep patterns (g.part5).

Conclusion

Most of the knowledge on physical activity and sleep comes from studies using questionnaire data or, more recently, proprietary physical activity metrics generated using commercial software. With the rapid advancement in technology over recent decades, raw acceleration movement sensors can now be used in large scale studies. However, the data they generate are not straightforward to

analyze. GGIR provides a tool for researchers to derive variables that characterize physical activity and sleep assessed in an objective manner. In addition, as the code is open source, it can be used in part or as a whole making it flexible to research needs. It further facilitates a reproducible analysis of the raw data which is key to generating conclusions in clinical and observational research settings. Previously, widely used software packages for analyzing accelerometer data have been proprietary and there has been little

opportunity for most researchers to feed into how they can be developed. This demonstrates the importance of open source software development for the research community and for the advancement of methods in physical behavior research.

Acknowledgments

We would like to thank Zhou Fang for his help in designing the auto-calibration algorithm, Jing Hua Zhao for his help in writing the code to read the binary data collected with the non-commercial GENE accelerometer, Joe Heywood for his help with fixing a bug in the timestamp conversion, and Evgeny Mirkes for his help in writing the code to read the binary data collected with the Axivity AX3 accelerometer. Further we would like to acknowledge funding from: Medical Research Council (MRC) grants MR/P012167/1 and MC_UU_12015/3, Economic and Social Research Council (ESRC) grant ES/K005987 for 'Millennium Cohort Study Sweep 6 (Age 14) Survey' (MCS6), Biotechnology and Biological Sciences Research Council (BBSRC) and Unilever Discover Ltd funded a CASE studentship (2008-2012), and NIH grants HL-094307 (AIP). This is part of a PhD Thesis conducted in the Biomedicine Doctoral Studies at the University of Granada, Spain. JHM is supported by a Grant from the Spanish Ministry of Education, Culture and Sport (FPU15/02645). AR is with the NIHR Leicester Biomedical Research Centre (2017–2022) and the Collaboration for leadership in Applied Health Research and Care (CLAHRC) East Midlands. This study has been partially funded by the University of Granada, Plan Propio de Investigación 2016, Excellence actions: Units of Excellence; Unit of Excellence on Exercise and Health (UCEES), and by the Junta de Andalucía, Consejería de Conocimiento, Investigación y Universidades and European Regional Development Fund (ERDF), ref. SOMM17/6107/UGR. VvT provides paid consultancy for GGIR users in need of support. The rest of the authors declare no conflict of interests.

References

- Bachasson, D., Landon-Cardinal, O., Benveniste, O., Hogrel, J.-Y., & Allenbach, Y. (2017). Physical activity monitoring: A promising outcome measure in idiopathic inflammatory myopathies. *Neurology*, 89(1), 101–103. PubMed ID: 28566549 doi:10.1212/WNL.0000000000004061
- Bell, J.A., Hamer, M., van Hees, V.T., Singh-Manoux, A., Kivimäki, M., & Sabia, S. (2015). Healthy obesity and objective physical activity. *The American Journal of Clinical Nutrition*, 102(2), 268–275. PubMed ID: 26156738 doi: 10.3945/ajcn.115.110924
- Campbell, C., Gott, A., Langford, J., & Sweetland, C. (2018). *GENEAclassif*. Retrieved from <https://cran.r-project.org/package=GENEAclassif>
- Cassidy, S., Fuller, H., Chau, J., Catt, M., Bauman, A., & Trenell, M.I. (2018). Accelerometer-derived physical activity in those with cardio-metabolic disease compared to healthy adults: A UK Biobank study of 52,556 participants. *Acta Diabetologica*, 55(9), 975–979. PubMed ID: 29808390 doi:10.1007/s00592-018-1161-8
- da Silva, I.C., van Hees, V.T., Ramires, V.V., Knuth, A.G., Bielemann, R.M., Ekelund, U., . . . Hallal, P.C. (2014). Physical activity levels in three Brazilian birth cohorts as assessed with raw triaxial wrist accelerometry. *International Journal of Epidemiology*, 43(6), 1959–1968. PubMed ID: 25361583 doi:10.1093/ije/dyu203
- Doherty, A., Hollowell, S., Willetts, M., Jackson, D., & Hammerla, N. (2017). *BiobankAccelerometerAnalysis*. Retrieved from <https://github.com/activityMonitoring/biobankAccelerometerAnalysis>
- Doherty, A., Jackson, D., Hammerla, N., Plötz, T., Olivier, P., Granat, M.H., . . . Wareham, N.J. (2017). Large scale population assessment of physical activity using wrist worn accelerometers: The UK biobank study. *PLoS One*, 12(2), e0169649. PubMed ID: 28146576 doi:10.1371/journal.pone.0169649
- Eddelbuettel, D. (2013). *Seamless R and C++ integration with Rcpp*. New York, NY: Springer New York. doi:10.1007/978-1-4614-6868-4
- Gubelmann, C., Heinzer, R., Haba-Rubio, J., Vollenweider, P., & Marques-Vidal, P. (2018). Physical activity is associated with higher sleep efficiency in the general population: the CoLaus study. *Sleep*, 41(7). doi:10.1093/sleep/zsy070
- He, Q., Zhang, P., Li, G., Dai, H., & Shi, J. (2017). The association between insomnia symptoms and risk of cardio-cerebral vascular events: A meta-analysis of prospective cohort studies. *European Journal of Preventive Cardiology*, 24(10), 1071–1082. PubMed ID: 28359160 doi:10.1177/2047487317702043
- Innerd, P., Harrison, R., & Coulson, M. (2018). Using open source accelerometer analysis to assess physical activity and sedentary behaviour in overweight and obese adults. *BMC Public Health*, 18(1), 543. PubMed ID: 29685121 doi:10.1186/s12889-018-5215-1
- Jackson, D. (2018a). *Export Re-sampled WAV*. Retrieved from <https://github.com/digitalinteraction/openmovement/wiki/AX3-GUI#export-re-sampled-wav>
- Jackson, D. (2018b). *OMGUI*. Retrieved from <https://github.com/digitalinteraction/openmovement/wiki/AX3-GUI>
- Jones, S.E., Lane, J.M., Wood, A.R., van Hees, V.T., Tyrrell, J., Beaumont, R.N., . . . Weedon, M.N. (2019). Genome-wide association analyses of chronotype in 697, 828 individuals provides new insights into circadian rhythms in humans and links to disease. *Nature Communications*, 10(1). <https://doi.org/10.1038/s41467-018-08259-7>
- Jones, S.E., van Hees, V.T., Mazzotti, D.R., Marques-Vidal, P., Sabia, S., van der Spek, A., . . . Wood, A.R. (2019). Genetic studies of accelerometer-based sleep measures in 85, 670 individuals yield new insights into human sleep behaviour. *Nature Communications*, 10(1). <https://www.nature.com/articles/s41467-019-09576-1>
- Lane, J.M., Jones, S., Dashti, H., Wood, A., Aragam, K., van Hees, V.T., . . . Saxena, R. (2018). Biological and clinical insights from genetics of insomnia symptoms. *Nature Genetics*, 51, 387–393. Retrieved from <https://www.biorxiv.org/content/early/2018/02/02/257956>
- Lee, I.-M., Shiroma, E.J., Lobelo, F., Puska, P., Blair, S.N., Katzmarzyk, P.T., & Lancet Physical Activity Series Working Group. (2012). Effect of physical inactivity on major non-communicable diseases worldwide: An analysis of burden of disease and life expectancy. *Lancet (London, England)*, 380(9838), 219–229. doi:10.1016/S0140-6736(12)61031-9
- Ligges, U., Krey, S., Mersmann, O., & Schnackenberg, S. (2018). *tuneR: Analysis of music and speech*. Retrieved from <https://cran.r-project.org/package=tuneR>
- Ligges, U., Short, T., Kienzle, P., Schnackenberg, S., Billingham, D., Borchers, H.-W., . . . Weingessel, A. (2014). *signal: Signal processing*. Retrieved from <http://r-forge.r-project.org/projects/signal/>
- Menai, M., van Hees, V.T., Elbaz, A., Kivimäki, M., Singh-Manoux, A., & Sabia, S. (2017). Accelerometer assessed moderate-to-vigorous physical activity and successful ageing: results from the Whitehall II study. *Scientific Reports*, 8, 45772. PubMed ID: 28367987 doi:10.1038/srep45772
- Rowlands, A.V., Cliff, D.P., Fairclough, S.J., Boddy, L.M., Olds, T.S., Parfitt, G., . . . Beets, M.W. (2016). Moving forward with backward compatibility. *Medicine & Science in Sports & Exercise*, 48(11), 2142–2149. PubMed ID: 27327029 doi:10.1249/MSS.0000000000001015
- Rowlands, A.V., Edwardson, C.L., Davies, M.J., Khunti, K., Harrington, D.M., & Yates, T. (2018). Beyond cut points. *Medicine & Science in Sports & Exercise*, 50(6), 1323–1332. PubMed ID: 29360664 doi:10.1249/MSS.0000000000001561
- Rowlands, A.V., Mirkes, E.M., Yates, T., Clemes, S., Davies, M., Khunti, K., & Edwardson, C.L. (2018). Accelerometer-assessed physical activity in epidemiology: Are monitors equivalent? *Medicine and*

- Science in Sports and Exercise*, 50(2), 257–265. PubMed ID: 28976493 doi:10.1249/MSS.0000000000001435
- Rowlands, A.V., Yates, T., Davies, M., Khunti, K., & Edwardson, C.L. (2016). Raw accelerometer data analysis with ggir r-package: Does accelerometer brand matter?," *Medicine and Science in Sports and Exercise*, 48(10), 1935–1941. PubMed ID: 27183118 doi:10.1249/MSS.0000000000000978
- Sabia, S., Cogranne, P., van Hees, V.T., Bell, J.A., Elbaz, A., Kivimaki, M., & Singh-Manoux, A. (2015). Physical activity and adiposity markers at older ages: Accelerometer vs questionnaire data. *Journal of the American Medical Directors Association*, 16(5), 438.e7–438.e13. doi:10.1016/j.jamda.2015.01.086
- Sabia, S., van Hees, V.T., Shipley, M.J., Trenell, M.I., Hagger-Johnson, G., Elbaz, A., . . . Singh-Manoux, A. (2014). Association between questionnaire- and accelerometer-assessed physical activity: The role of sociodemographic factors. *American Journal of Epidemiology*, 179(6), 781–790. PubMed ID: 24500862 doi:10.1093/aje/kwt330
- Seinstra, F., Wallom, D., & Keahey, K. (2015). Editorial. *SoftwareX*, 1–2, 1–2. doi:10.1016/j.softx.2015.08.001
- Stiles, V.H., Metcalf, B.S., Knapp, K.M., & Rowlands, A.V. (2017). A small amount of precisely measured high-intensity habitual physical activity predicts bone health in pre- and post-menopausal women in UK Biobank. *International Journal of Epidemiology*, 46(6), 1847–1856. PubMed ID: 29106579 doi:10.1093/ije/dyx080
- van Hees, V.T., Fang, Z., Langford, J., Assah, F., Mohammad, A., da Silva, I.C.M., . . . Brage, S. (2014). Autocalibration of accelerometer data for free-living physical activity assessment using local gravity and temperature: An evaluation on four continents. *Journal of Applied Physiology (Bethesda, Md. : 1985)*, 117(7), 738–744. PubMed ID: 26806874 doi:10.1152/japplphysiol.00421.2014
- van Hees, V.T., Gorzelniak, L., Dean León, E.C., Eder, M., Pias, M., Taherian, S., . . . Brage, S. (2013). Separating movement and gravity components in an acceleration signal and implications for the assessment of human daily physical activity. *PLoS One*, 8(4), e61691. PubMed ID: 23626718 doi:10.1371/journal.pone.0061691
- van Hees, V.T., Renström, F., Wright, A., Gradmark, A., Catt, M., Chen, K.Y., . . . Franks, P.W. (2011). Estimation of daily energy expenditure in pregnant and non-pregnant women using a wrist-worn tri-axial accelerometer. *PLoS One*, 6(7), e22922. PubMed ID: 21829556 doi:10.1371/journal.pone.0022922
- van Hees, V.T., Sabia, S., Anderson, K.N., Denton, S.J., Oliver, J., Catt, M., . . . Singh-Manoux, A. (2015). A novel, open access method to assess sleep duration using a wrist-worn accelerometer. *PLoS One*, 10(11), e0142533. PubMed ID: 26569414 doi:10.1371/journal.pone.0142533
- van Hees, V.T., Sabia, S., Jones, S.E., Wood, A.R., Anderson, K.N., Kivimaki, M., . . . Weedon, M.N. (2018). Estimating sleep parameters using an accelerometer without sleep diary. *Scientific Reports*, 8(1), 12975. PubMed ID: 30154500 doi:10.1038/s41598-018-31266-z
- White, T. (2018). *Pampro*. Cambridge, UK. doi: 10.5281/zenodo.1187043
- Wijndaele, K., Westgate, K., Stephens, S.K., Blair, S.N., Bull, F.C., Chastin, S.F.M., . . . Healy, G.N. (2015). Utilization and harmonization of adult accelerometry data: Review and expert consensus. *Medicine and Science in Sports and Exercise*, 47(10), 2129–2139. PubMed ID: 25785929 doi:10.1249/MSS.0000000000000661
- Willets, M., Hollowell, S., Aslett, L., Holmes, C., & Doherty, A. (2018). Statistical machine learning of sleep and physical activity phenotypes from sensor data in 96, 220 UK Biobank participants. *Scientific Reports*, 8(1), 7961. PubMed ID: 29784928 doi:10.1038/s41598-018-26174-1
- Zeileis, A., & Grothendieck, G. (2005). *zoo : S3 Infrastructure for regular and irregular time series*. *Journal of Statistical Software*, 14(6). doi:10.18637/jss.v014.i06