



Energy-efficient recurrence quantification analysis

Norbert Marwan^{1,2,3,a} 

¹ Potsdam Institute for Climate Impact Research (PIK), Member of the Leibniz Association, Telegrafenberg A31, 14473 Potsdam, Germany

² Institute of Geoscience, University of Potsdam, Karl-Liebknecht-Straße 32, 14476 Potsdam, Germany

³ Institute of Physics and Astronomy, University of Potsdam, Karl-Liebknecht-Straße 32, 14476 Potsdam, Germany

Received 18 November 2025 / Accepted 18 December 2025

© The Author(s) 2026, modified publication 2026

Abstract Recurrence quantification analysis (RQA) is a widely used tool for studying complex dynamical systems, but its standard implementation requires computationally expensive calculations of recurrence plots (RPs) and line length histograms. This study introduces strategies to compute RQA measures directly from time series or phase space vectors, avoiding the need to construct RPs. The calculations can be further accelerated and optimised by applying a random sampling procedure, in which only a subset of line structures is evaluated. These modifications result in shorter run times, less memory use and access, and lower overall energy consumption during analysis whilst maintaining accuracy. This makes them especially appealing for large-scale data analysis and machine learning applications. The ideas are not limited to diagonal line measures, but can likewise be applied to vertical line-based measures and to recurrence network measures. By lowering computational costs, the proposed strategies contribute to energy saving and sustainable data analysis, and broaden the applicability of recurrence-based methods in modern research contexts.

1 Introduction

In the light of sustainable energy production and climate change, energy efficiency in modelling and data analysis is an important and growing topic [1, 2]. As the demand for data processing grows, it is crucial to focus on reducing related energy consumption [3–5]. This problem is increasingly attracting attention across various fields, but it remains in its early stages. Amongst the many goals and actions that help address energy-efficient computations [3] and are especially feasible in daily work and short-term contexts, the key measures include educating about this topic and enhancing implementations [6].

In the context of sustainable data analysis, energy efficiency is becoming increasingly critical. This is particularly relevant for computationally intensive methods such as recurrence quantification analysis (RQA), a powerful tool in nonlinear data analysis [7–10]. This framework has shown high potential in investigating diverse research questions across many disciplines [11] and its integration with machine learning (ML) has recently sparked considerable interest [12].

However, this potential comes at a cost: RQA stands out as a particularly energy-intensive method in the ML toolkit. Unlike linear techniques, RQA relies on pairwise comparisons in high-dimensional data, resulting in computational complexity that scales quadratically with data size. In ML workflows, this challenge is amplified: RQA is often embedded in iterative processes—such as feature extraction, hyperparameter tuning, or real-time inference—where its energy demands compound with those of the broader pipeline. For example, extracting RQA-based features from thousands of time series segments or combining RQA with neural networks in hybrid models can result in high energy consumption costs [13, 14]. Moreover, the nonlinear nature of RQA resists simple optimisations, making efficient implementations not just a performance concern, but a necessity for sustainable ML. Addressing this gap is critical to unlocking RQA’s potential in data-intensive fields like climate science, biomedicine, and IoT—where energy-efficient algorithms are crucial for scalable, real-world applications [14]. Beyond reducing energy consumption, optimised algorithms generally provide faster computations, enabling real-time applications and large-scale deployments [2, 13, 15].

^a e-mail: marwan@pik-potsdam.de (corresponding author)

The standard RQA algorithm begins with calculating the RP, which has a computational complexity of $\mathcal{O}(N^2)$ [8]. The lines formed by consecutive points in the RP are then detected, and their lengths are measured. As we consider diagonal and vertical lines, the computational complexity for each of these steps is also of order $\mathcal{O}(N^2)$ (also in the case of symmetric RPs, where we would consider only one triangle of the RP). This complexity requires many calculation steps and, therefore, also increases computation times. In the last decade, various approaches have been proposed to speed-up RQA calculations. Classical approaches utilise parallelisation and use multiple graphics processing unit (GPU) devices [15, 16]. The latter distributes the calculations of RPs and line lengths on different GPU devices using the divide and recombine paradigm [17]. Although the acceleration of the computations is remarkable, the complexity of the calculations has not changed. An approach that can indeed reduce the computational complexity extremely is a numerical and geometrical approximative ansatz, which can heavily accelerate the calculations by several magnitudes, but with the costs of increasing uncertainties [18, 19]. More recently, the introduction of microstate recurrence analysis [20] has offered another (random sampling-based) approximative but fast approach, with much less uncertainty in the RQA measures [21, 22]. Here, the line length distributions, required for the RQA measures, are estimated from randomly sampled sub-matrices from the RP.

In the following, we will consider the original approach of RQA calculations, without parallelisation (although it would be possible) and without numerical/geometrical approximation, but using the diagonal line structures of interest and developing some optimisation strategies. This optimisation follows the action numbers M2 (“Reduce and compress data having the anticipated scientific value of the retained information and the resource requirements in mind”) and M7 (“Design software for optimised energy consumption and provide tools to measure it”) [3]. These actions are part of the “Call-to-action in digital transformation”, which outline concrete measures developed by members of the ErUM-Data community during a dedicated workshop, aiming to reduce greenhouse gas emissions and promote the use of renewable energy within data-intensive research [3]. For the sake of simplicity, only the diagonal lines are considered here; the same approach can be applied for calculating the measures based on the vertical lines or even more complex patterns, such as triangle motifs (corresponding to the recurrence configuration that two neighbours of a state vector are also neighbours, which is required for recurrence network measures [23, 24]). The first optimisation will be very simple and has surely already been applied in several implementations. The second one is inspired by the microstates approach by Ref. [21]. These approaches are compared to a simple, straightforward standard implementation without any optimisations, as someone lacking knowledge of algorithmic and language-specific optimisations (such as those available in Julia) would perform it. The algorithms are implemented using the Julia language because it enables easy measurement of the performance of the implementations.

2 Recurrence plots and recurrence quantification analysis

A recurrence plot (RP) is a square matrix \mathbf{R} indicating the similarity of two states $\vec{x}(i)$ and $\vec{x}(j)$ at different points in time i and j :

$$R(i, j) = \Theta(\varepsilon - \|\vec{x}(i) - \vec{x}(j)\|), \quad (1)$$

with states $\vec{x} \in \mathbb{R}^d$ (d the dimension of the phase space), $i, j = 1, \dots, N$ the time indices, Θ the Heaviside function, and ε the recurrence threshold [8]. The RP exhibits typical large-scale appearances, depending on the system’s dynamics [7, 8]. It further contains small-scale line structures that represent temporally close evolution of the phase space trajectory pairs (diagonal lines) or trapped states (vertical lines). The distributions of these line lengths provide insights into the dynamics of the system and are used to define measures of complexity within recurrence quantification analysis (RQA).

A typical RQA measure is the determinism measure (DET), which quantifies the ratio of recurrence points ($R(i, j) = 1$) which form diagonal lines and the total number of recurrence points, and is based on the histogram $P(\ell)$ of the occurrences of lines of length ℓ ,

$$\text{DET} = \frac{\sum_{\ell=2}^N \ell P(\ell)}{\sum_{\ell=1}^N \ell P(\ell)}. \quad (2)$$

This measure is related to predictability, with large values for highly predictable dynamics and low values for non-predictable dynamics. It can be used, e.g. to identify regime changes or classify different regimes or systems [25–28].

For a more general overview of the method and its application potential, the reader is referred to Ref. [11, 29].

3 Optimisation strategies

The following implementations have been tested on a single core of the high-performance cluster “Eunice Newton Foote” at the Potsdam Institute of Climate Impact Research (AMD EPYC 9554 processors with scalar frequencies of up to 3.75 GHz and 6 GByte DDR5 memory per core). The performance of an implementation was tested using the Julia macros `@time` and `@timed`. The calculations were performed 100 times to ensure reliable statistical estimates of the performance measurements.

As test data, the three components of the Rössler system (using standard parameters $a = 1.2$, $b = 0.2$, and $c = 5.7$) are used, with length 25,000 and sampling time $\Delta t = 0.2$ [30]. The Rössler system was integrated using a Tsitouras 5/4 Runge–Kutta–Solver as implemented by the DifferentialEquations.jl package in Julia.

Recurrence, Eq. (1), is calculated for all three components of the Rössler system, using a threshold ε of 10% of the range of the values of all three components, i.e. $\varepsilon = 0.1(\max(\vec{x}) - \min(\vec{x}))$.

3.1 Histogram estimation without RP (RQA_woRP)

A line structure in an RP is a sequence of pairs of time indices. A diagonal line in the RP of length ℓ with the coordinates $\{(i, j), (i + 1, j + 1), \dots, (i + \ell - 1, j + \ell - 1)\}$ represents that the sequences of states $\{\vec{x}(i), \vec{x}(i + 1), \dots, \vec{x}(i + \ell - 1)\}$ and $\{\vec{x}(j), \vec{x}(j + 1), \dots, \vec{x}(j + \ell - 1)\}$ are similar within the recurrence uncertainty defined by the threshold ε (Fig. 1).

It is clear that we can find line lengths in the RP without the RP. We can test the condition

$$\|\vec{x}(i) - \vec{x}(j)\| \leq \varepsilon \tag{3}$$

directly at the series of the phase space vectors (or the time series) for increasing indices i and j . Considering the condition

$$\|\vec{x}(i - 1) - \vec{x}(j - 1)\| > \varepsilon \tag{4}$$

for the start of a line and

$$\|\vec{x}(i + \ell) - \vec{x}(j + \ell)\| > \varepsilon \tag{5}$$

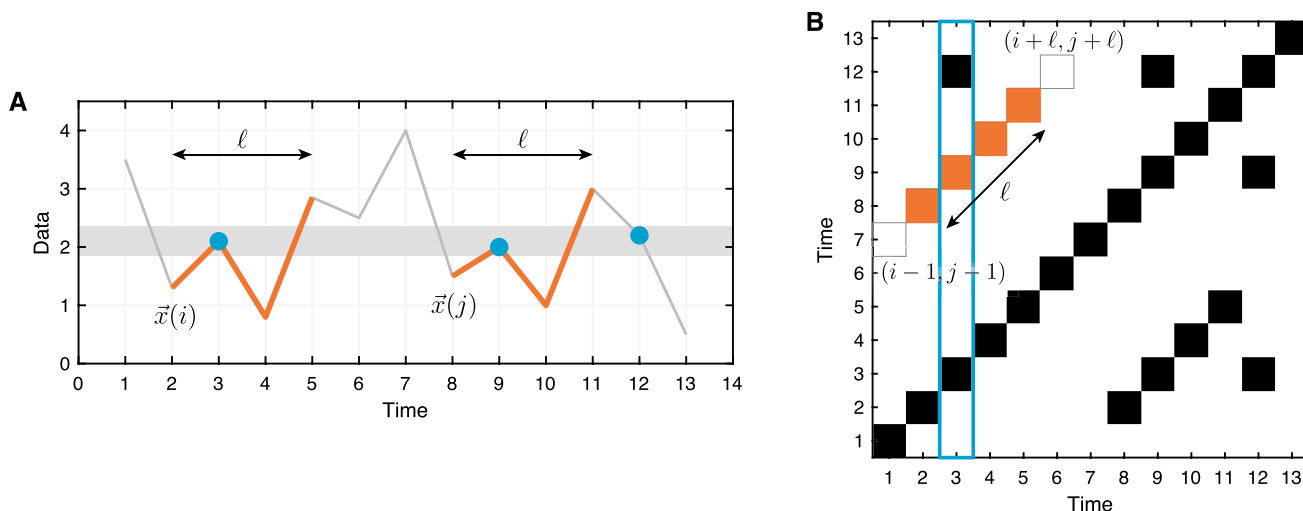


Fig. 1 (A) Time-series sequence and (B) corresponding recurrence plot. The sequence at time points 2 to 5 (orange) repeats four times within a small error ε (here $\varepsilon = 0.25$) during the interval 8–11, forming a diagonal line with points $(i = 2, j = 8)$, $(i + 1 = 3, j + 1 = 9)$, $(i + 2 = 4, j + 2 = 10)$, and $(i + 3 = 5, j + 3 = 11)$ in the RP (orange points). The line is preceded by condition Eq. (4), $\|\vec{x}(i - 1) - \vec{x}(j - 1)\| > \varepsilon$, and followed by condition Eq. (5), $\|\vec{x}(i + \ell) - \vec{x}(j + \ell)\| > \varepsilon$, (grey boxes in panel (B)), ensuring the lines start and end points. The length of the diagonal line ($\ell = 4$) can be measured in the RP or directly from the time series. The state at time 3 recurs (within the ε uncertainty, grey horizontal bar in panel (A)) at time 9 and 12 (blue dots), indicated by the points in the column 3 (marked by blue box)

Table 1 General performance values of the different implementations when calculating the line length histograms required for RQA and using the Rössler system with all three components and of length $N = 25,000$

Implementation	Computation time	Speedup	Allocations	Energy
RQA from RP (RQA_RP)	2.01 s	1	7	27.9 mWh
RQA without RP (RQA_woRP)	0.88 s	2.3	3	12.2 mWh
RQA without RP ^a (RQA_woRP ¹)	0.48 s	4.2	3	6.7 mWh
RQA from sampling (RQA_Samp)	0.25 s	8.0	3	3.5 mWh
RQA from sampling ^b (RQA_Samp ²)	0.012 s	168	3	0.17 mWh
RQA from microstates	0.017 s	1148	29	0.24 mWh

Number of samplings for the sampling version and microstates $M = 4N = 100,000$, except for (RQA_Samp²)

^aThe calculation directly from the time series (RQA_woRP) has been performed additionally with the Julia instruction `@inbounds` which can accelerate the code further

^bCalculation of RQA_Samp using a smaller number of samplings $M = 0.2N = 5000$

as the end of a line, we can get the distribution of line lengths $P(\ell)$. This is a very simple and straightforward calculation of $P(\ell)$ and allows us to get the RQA measures without previous calculating the RP beforehand. The resulting histogram of the diagonal lines will be exactly the same as the one obtained from the RP. It will not reduce the numerical complexity, but the number of calculation steps, and should, therefore, result in more efficient and faster RQA calculations when the RP is not required. It also reduces the number of memory allocations, as we do not need to store and calculate the matrix \mathbf{R} .

This kind of optimising the histogram calculation without calculating the RP beforehand is surely not novel. It has been used by the author before [31] and also have been implemented by others (e.g. Ref. [32]). For example, such implementation is available in Ref. [31] and was used in the runtime comparisons in the studies by Refs. [16, 17] and [19].

The comparison with the standard implementation using RP-based line histogram calculation shows a speed-up by a factor of more than 2 (Table 1). The number of memory allocations dropped from 7 to 3. The just-in-time compiler (JIT) in Julia allows some internal optimisation steps, e.g. by applying the macro `@inbounds` to the for-loops [33]. Using it can further accelerate the implementation. Here, it finally speeds up the implementation by a factor of more than 4. In the following analysis, we will use the `@inbounds` optimised version.

3.2 Sampling-based histogram estimation (RQA_Samp)

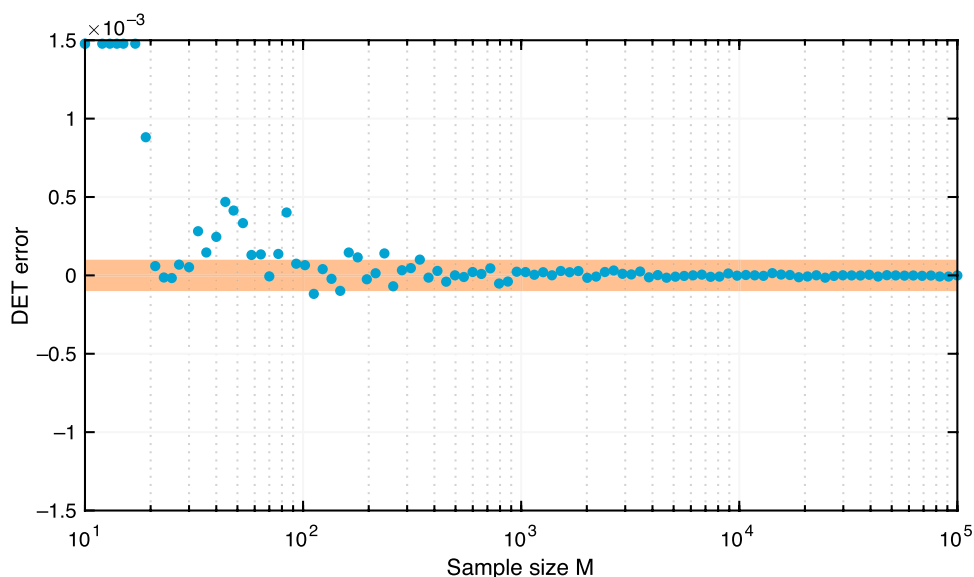
Based on the histogram estimation directly from the data, and without RP, we can further reduce the calculation costs by applying a random sampling scheme. Instead of sequentially testing every index j from 1 to N for every i , we can randomly draw indices i and j and test whether this pair fulfils the conditions given by Eqs. (3) and (4), i.e. correspond to the beginning of a diagonal line. If not, we randomly draw a new pair i and j . If yes, we can measure the line length as is done in RQA_woRP, but stop after the end of a line and randomly sample a new line. By this sampling schema, we get only a subset of lines in the histogram $P(\ell)$, depending on the number M of samplings. Using a smaller number of M , e.g. $M = 4N$, we can significantly reduce the number of calculations to estimate $P(\ell)$. The resulting histogram $P(\ell)$ is a direct approximation of the histogram as it would be obtained using the RP_woRP algorithm described in Sect. 3.1. Importantly, this approximation concerns the same underlying distribution, and the standard RQA measures can be computed from $P(\ell)$ without any modification. Even for small values of M , this distributional approximation is already quite accurate.

In general, this sampling approach is, somehow, similar to estimating RQA measures using randomly sampled microstates as explained in Ref. [21], but with the difference that RQA measures (such as DET) can be directly estimated from the histogram $P(\ell)$, i.e. they are not reconstructed from a different (the microstated) distribution, and, thus, will result in higher accuracy when using the same sampling size M .

Comparing to the standard estimation algorithm (RQA_RP) and the direct one (RP_woRP), we find a further speedup depending on the number M (Table 1). The memory allocations are the same as for the direct histogram estimation without RPs. A large number of samples $M = 4N$, the computation time is halved; further reducing M to 20% of the data length speeds up the RQA estimation by two orders of magnitude.

Therefore, it might be desired to have a small M , but it obviously causes worse estimates of the line length measures. Considering the error $\text{DET}_{\text{sampled}} - \text{DET}_{\text{true}}$ shows that increasing M reduces the error of the DET estimation when obtained with this sampling schema (Fig. 2). For the considered data with $N = 25,000$, a sampling size of $M = 500$ would already result in excellent estimations of DET with errors of $< 10^{-4}$, using $M = 5000$ the estimation errors drop below 10^{-5} .

Fig. 2 Estimation error $DET_{\text{sampled}} - DET_{\text{true}}$ derived from sampled line length histogram for the Rössler system with 25,000 data points and for sampling size ranging from $M = 10$ to 100,000. The orange region indicates the error range of $\pm 10^{-4}$. A random sampling of only 500 line segments already results in minor errors smaller than 10^{-4}



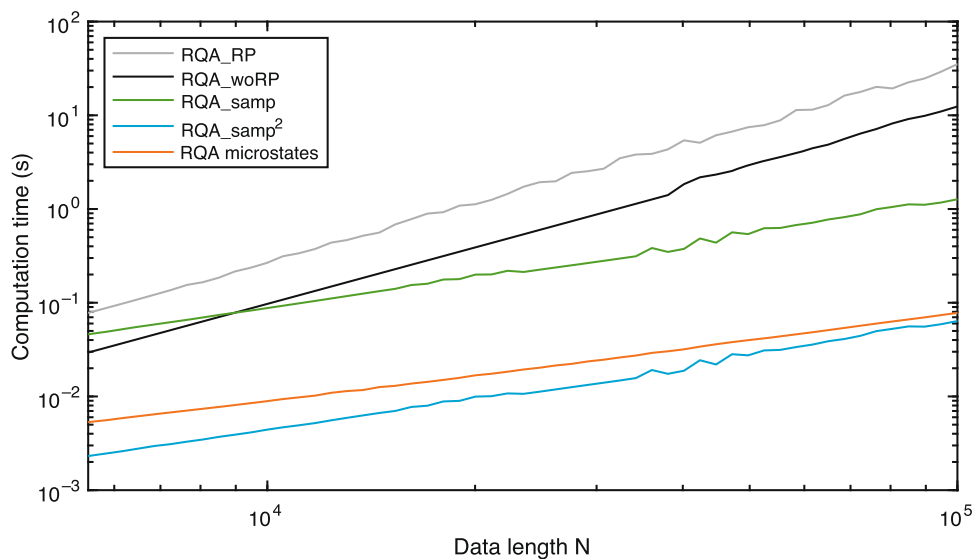
3.3 Efficiency by computation time

One aspect of efficient implementations is the reduced number of calculation steps, saving energy and also computation time. The first approach of calculating the RQA directly from the data, without the RP (RP_woRP), shows already a significant reduction of computation time of more than 50%, because calculating the RP is not required (Fig. 3). Further optimisation using the sampling schema (RP_Samp) reduces the computation time remarkably. This reduction depends on the sampling size M . However, it is not necessary to scale the sampling size M with the squared size N of the data. We find high accuracy already when we scale M linearly with the size N , e.g. by $M = 4N$ or $M = 0.2N$, as it is used here in the comparison. This results in a lower scaling exponent of the computation time as data length increases (Fig. 3). This means, the longer the data, the more efficient this method becomes.

We also note that for small data ($N < 10,000$), the sampling schema is not necessarily faster than RQA_woRP. In addition, uncertainties in the RQA results increase for shorter time series when using sampling-based approaches, affecting both RQA_Samp and the microstates method (see Appendix, Fig. 5), suggesting that RQA_woRP is preferable in such cases.

Using the same number M in the recurrence microstates approach [22], the computation time is, in general, smaller (although it might not be directly comparable, because the microstates calculation is based on compiled C code highly optimised for Julia [34]), but with the cost of less accurate results (see Appendix, Figs. 4 and 5).

Fig. 3 Computation times for the RQA implementations using RP (RQA_RP), without RP (RP_woRP), with a sampling scheme (RP_Samp), and using a recurrence microstates approximation for increasing data length N (Rössler system, same parameters as in Fig. 2). The sampling size M for RP_Samp and microstates is $M = 4N = 100,000$, and for RP_Samp², it is $M = 0.2N = 5000$



The reduced computation time directly translates into lower energy consumption. Estimating a specific power draw per CPU core for RQA calculations is not easy, as it depends on the workloads, architecture, computer system (e.g. HPC, laptop), underlying operating systems and other factors [4, 5, 35]. However, assuming an average power draw of ≈ 50 W per CPU core as a realistic baseline, replacing RQA_RP (2.06 s for data with length $N = 25,000$) with RQA_Samp (0.26 s) reduces the energy usage from approximately 0.0290.0036 Wh per run, corresponding to an energy saving of about 0.025 Wh, or almost 90%. Using the even faster RQA_Samp² (0.013 s) further lowers the consumption to 0.00018 Wh, i.e. a saving of roughly 0.028 Wh per run. Whilst these savings are small for a single computation, they accumulate substantially when processing large data sets, performing repeated analyses, or running parallel workflows on HPC systems. For example, 100,000 runs of a typical RQA of data with length $N = 25,000$ accumulate to an energy saving of 2.8 kWh, equivalent to ≈ 1.0 kg of CO₂ emissions (assuming 363 g CO₂ per kWh based on the German electricity mix in 2024 [36]).

4 Discussion and conclusion

Efficiency gains in data analysis can be achieved at multiple levels. Low-level optimisations, such as compiler hints (e.g. `@inbounds` in Julia's JIT compiler) or vectorisation directives, can reduce execution time without changing the underlying algorithm. However, more substantial improvements are typically obtained through algorithmic modifications, such as avoiding intermediate data structures and calculation steps, or applying sampling schemes, like the RQA_woRP and RQA_Samp approaches.

These discussed approaches demonstrate that recurrence-based complexity measures can be obtained in a direct and computationally efficient way without the need for pre-computed RP. When this idea is combined with a sampling strategy, in which only a fraction of the possible structures are analysed, the computational costs can be reduced even further. This leads not only to faster execution, but also to lower memory requirements and energy consumption, which are relevant considerations in modern large-scale data analysis pipelines.

Such efficiency gains are particularly important in applications where RQA is employed in combination with machine learning techniques, as efficient computation allows more extensive training and validation procedures and facilitates the systematic exploration of the model architectures. They are equally relevant in the context of large-scale (big) data analysis, where reduced memory and runtime requirements make it possible to apply RQA to very long time series, large ensembles of signals, or even continuous data streams. By strengthening both domains, the proposed approach enhances the accessibility and practical utility of RQA in modern data-driven research. Furthermore, the reduced computational overhead may facilitate the integration of RQA methods into real-time or embedded systems, where resources are limited.

The choice of which approach to use depends on the research questions, data characteristics and computational constraints. For small datasets ($N < 10,000$) and when RPs are not needed, the RQA_woRP approach is the best choice, ensuring high accuracy and reasonable computation time. RQA_Samp becomes advantageous for longer time series, due to its linear sampling complexity ($M \propto N$), which scales more favourably than the $\mathcal{O}(N^2)$ complexity of RQA_woRP, i.e. the computational savings grow with data length.

The results presented here demonstrate typical behaviour for a continuous chaotic system (Rössler attractor). Different dynamical regimes may affect the performance of these approaches differently. For instance, discrete maps like the logistic map generate RPs with shorter, more heterogeneous line structures. This leads to faster sampling (as valid line starts are found more frequently) but potentially larger uncertainties, particularly when M is insufficient to adequately represent the line length distribution (Fig. 6A, C). In contrast, periodic systems have more homogeneous RPs with very long, non-interrupted, but less frequent diagonal lines. This reduces the probability of finding lines through random sampling, increasing computation time (can be even worse than the standard RQA_RP approach). However, the limited diversity of line lengths ensures accurate results even for small M (Fig. 6B, D).

Sampling approaches are most efficient for homogeneously structured RPs, characteristic of stationary dynamics. Non-stationary systems with heterogeneous RPs are expected to lead to larger uncertainties in RQA_Samp results, requiring either the use of RQA_woRP or increasing M . Systematic guidelines for selecting M in non-stationary cases warrant future investigation.

The presented approaches also come with potential trade-offs. The sampling of structures introduces an additional source of variability, and its influence on the robustness and reliability of different RQA measures still needs to be systematically evaluated. Future studies should investigate how sampling density, noise and parameter settings affect the stability of the obtained measures, particularly for non-stationary dynamics where line structures may be heterogeneously distributed. Another important aspect concerns the integration of established correction schemes, such as those for removing sojourn points [8, 37, 38] or mitigating border effects [38, 39], into the sampling framework.

Beyond methodological refinements, several directions could further enhance the computational efficiency and applicability of these approaches. Combining the proposed sampling strategy with parallel computing architectures

could improve scalability for very large datasets and real-time applications. In addition, the current implementations could benefit from further refinements. Replacing the fixed sampling size M with an adaptive convergence scheme—which stops sampling once results stabilise within a predefined threshold—could optimise the trade-off between accuracy and computation time. Combined with low-level compiler optimisations such as JIT compiler hints and vectorisation [33], these enhancements could yield additional modest speedups.

The strategies discussed here are not limited to the diagonal line measures. They can be analogously applied to RQA measures based on vertical line structures [40] and to recurrence network measures, such as clustering coefficient and transitivity [24], broadening the scope of energy-efficient recurrence analysis.

In summary, the combination of direct line detection and sampling represents a promising way forward in the efficient computation of RQA and related measures. Beyond the methodological advantages, the reduction in computational cost also contributes to the principles of green computing by lowering energy consumption, which makes the approach more sustainable for large-scale or long-term data analysis.

Acknowledgements The author thanks Felipe Eduardo Lopes da Cruz, Thiago de Lima Prado and Sergio Roberto Lopes for inspiring discussions on sampling microstates and for motivating this research. The author gratefully acknowledges support from the German Academic Exchange Service (DAAD) and Coordenação de Aperfeiçoamento de Pessoal de Nível Superior (CAPES) via the project “Recurrence quantifiers as features for machine-learning-decision-making processes” (Grant No. 57705568), as well as from the Ministry of Research, Science and Culture (MWFK) of Land Brandenburg for supporting this project by providing resources on the high-performance computer system “Eunice Newton Foote” at the Potsdam Institute for Climate Impact Research (Grant No. 22-Z105-05/002/001). For the RQA estimations using the microstates approach, the recently published Julia package `RecurrenceMicrostatesAnalysis.jl` [22] has been used.

Funding Open Access funding enabled and organized by Projekt DEAL.

Data Availability Code used for this analysis and to reproduce the results and figures of this study are available at Zenodo: <https://doi.org/10.5281/zenodo.17620044>.

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 licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence 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 licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendix

Energy consumption per run

Assuming a constant power draw P (e.g. $P = 50$ W) and a computation time t measured in seconds, the energy consumption per run is

$$E = P \frac{t}{3600} \text{ [Wh]}.$$

Accordingly, the energy saving obtained when replacing a reference method with runtime t_{ref} by a faster method with runtime t is

$$\Delta E = P \frac{t_{\text{ref}} - t}{3600} \text{ [Wh]}.$$

Performance of the algorithm and further dynamical systems

See Figs. 4, 5 and 6.

Recommendation for required data length

Whilst a more systematic study is necessary to get good recommendations on required time series lengths for the proposed `RQA_samp`² approach, we can already gather some clues.

Fig. 4 Estimation error of DET derived from recurrence microstates approximation for the Rössler system with 25,000 data points and sampling size ranging from $M = 10$ to 100,000. The orange region indicates the error range of $\pm 10^{-4}$. To achieve estimation errors below 10^{-4} , the sampling size must be at least 40,000. Up to $M = 5000$, the estimated DET remains at 1.00, resulting in the constant deviation

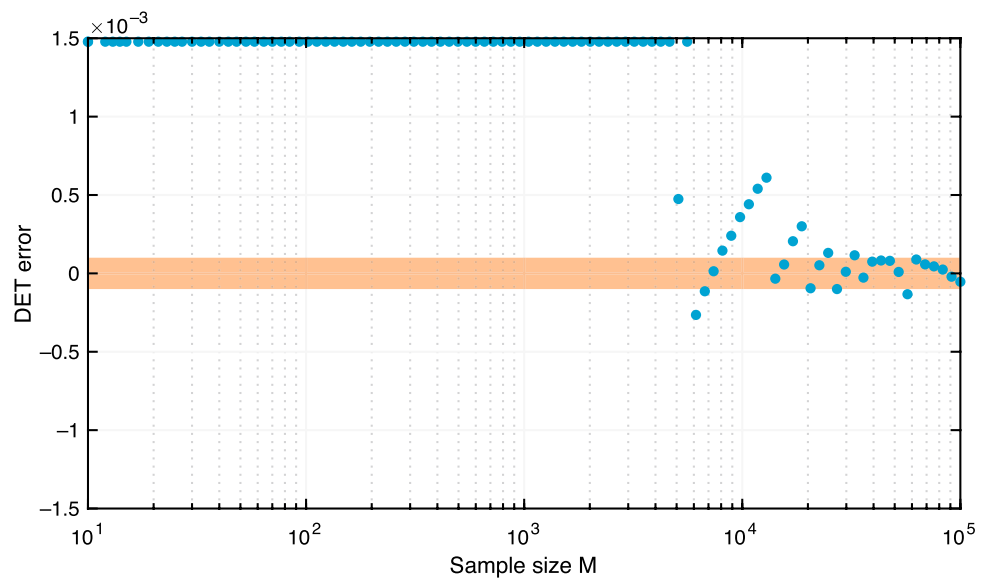
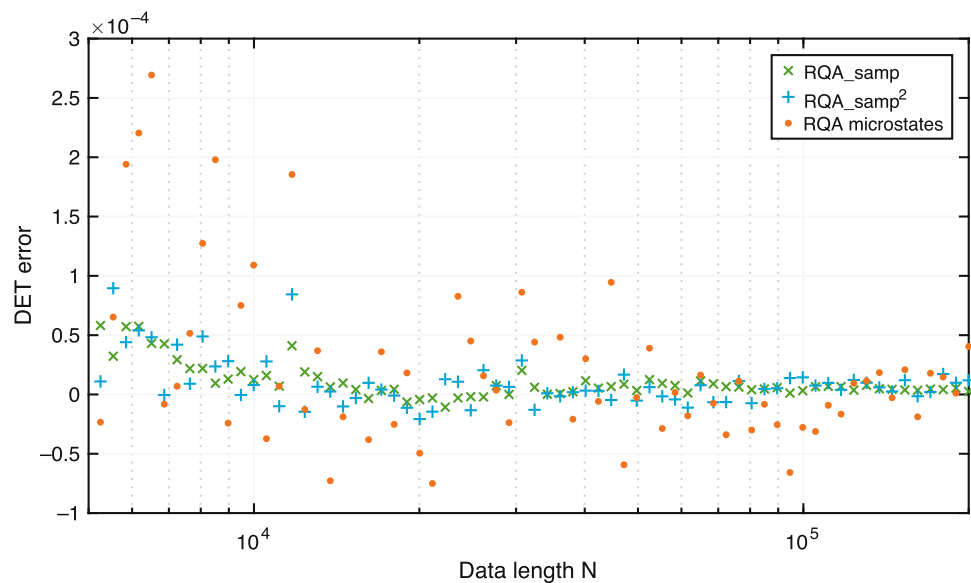


Fig. 5 Estimation error $\text{DET}_{\text{sampled}} - \text{DET}_{\text{true}}$ for DET estimations using RQA_Samp, RQA_Samp², and recurrence microstates approximation for the Rössler system for increasing data length N . The sampling was $M = 100,000$ for RQA_Samp and microstates RQA, and $M = 5000$ for RQA_Samp². Even for a very small sampling number, the RQA_Samp approach has small estimation errors $< 10^{-4}$



Errors in estimating the RQA measure and computation times with RQA_samp² depend on the characteristics of the data. Data that cause relatively homogeneously distributed long lines in an RP, such as for Rössler or Lorenz systems, show estimation errors that are already quite small for short data lengths, with $100 < N < 1000$. The speedup in computation time becomes significant for $N > 1000$ (Fig. 7A). For systems that generate many short lines and single points, as typical for maps, the errors are generally larger, but the computation time benefits already significantly for small N (Fig. 7B). To achieve reasonable results, at least $N > 2000$ are necessary, although this number can be lowered by increasing the sampling size M . In contrast, sparse RPs with consistent patterns cause numerous unsuccessful sampling tries, leading to limited speedup but very slight errors (Fig. 7C). Therefore, this approach might be less effective for periodic systems compared to signals with more complex or irregular dynamics.

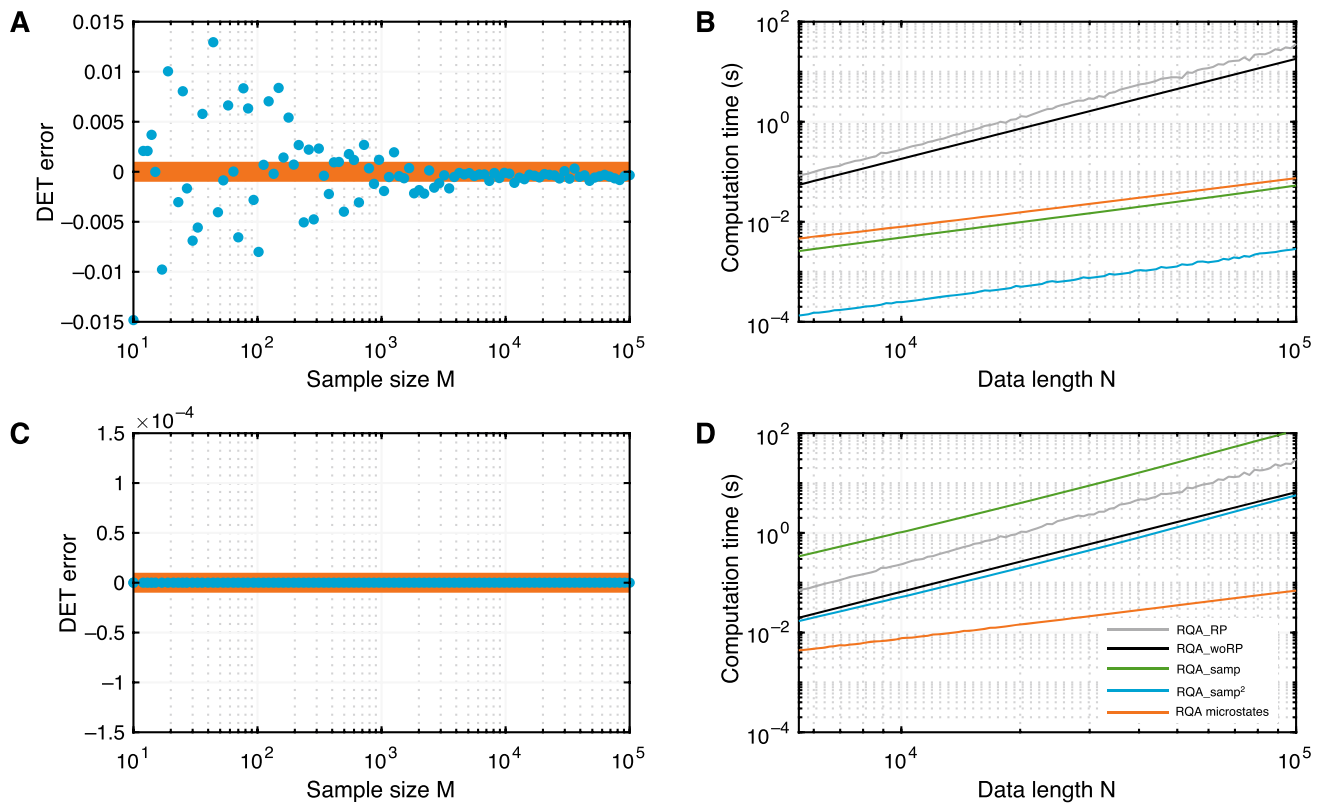


Fig. 6 (A, C) Estimation error $DET_{\text{sampling}} - DET_{\text{true}}$ and (B, D) computation times as in Figs. 2 and 3 for the (A, B) the logistic map and (C, D) a periodic signal. In (A, C) results for $N = 25,000$ and using RQA_Samp² are shown

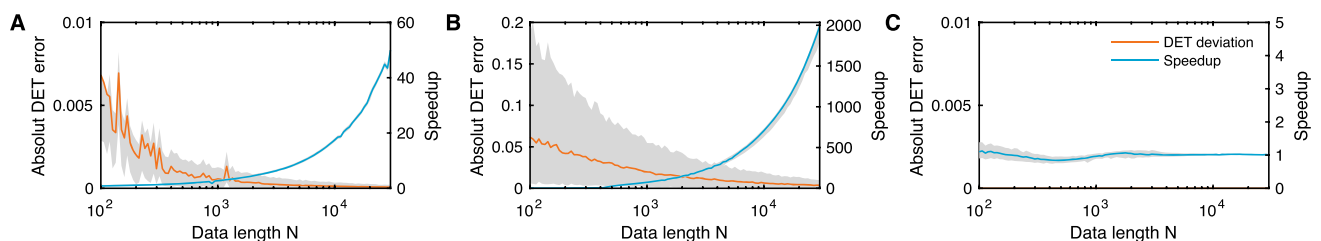


Fig. 7 Speedup of computation time ($t_{\text{RQA_samp}^2}/t_{\text{RQA_woRP}}$) and absolute estimation error $|DET_{\text{sampling}} - DET_{\text{true}}|$ as function of data length using the RQA_samp² approach for (A) the Rössler system, (B) the logistic map and (C) a periodic signal. Shaded areas indicate the 5% and 95% quantiles obtained from 500 repetitions; solid lines denote the median

References

1. S.G. Paul, A. Saha, M.S. Arefin, T. Bhuiyan, A.A. Biswas, A.W. Reza, N.M. Alotaibi, S.A. Alyami, M.A. Moni, A comprehensive review of green computing: past, present, and future research. *IEEE Access* **11**, 87445–87494 (2023). <https://doi.org/10.1109/access.2023.3304332>
2. ...S. Zhang, F. Haohuan, W. Lixin, Y. Li, H. Wang, Y. Zeng, X. Duan, W. Wan, L. Wang, Y. Zhuang, H. Meng, X. Kai, X. Ping, L. Gan, Z. Liu, W. Sihai, Y. Chen, Yu. Haining, S. Shi, L. Wang, X. Shiming, W. Xue, W. Liu, Q. Guo, J. Zhang, G. Zhu, T. Yang, J. Edwards, A. Baker, J. Yong, M. Yuan, Yu. Yangyang, Q. Zhang, Z. Liu, M. Li, D. Jia, G. Yang, Z. Wei, J. Pan, P. Chang, G. Danabasoglu, S. Yeager, N. Rosenbloom, Y. Guo, Optimizing high-resolution community earth system model on a heterogeneous many-core supercomputing platform. *Geosci. Model Dev.* **13**(10), 4809–4829 (2020). <https://doi.org/10.5194/gmd-13-4809-2020>
3. B. Bruers, M. Cruces, M. Demleitner, G. Duckeck, M. Düren, N. Eich, T. Enßlin, J. Erdmann, M. Erdmann, P. Fackeldey, C. Felder, B. Fischer, S. Fröse, S. Funk, M. Gasthuber, A. Grimshaw, D. Hadasch, M. Hannemann, A. Kappes, R. Kleinemühl, O.M. Kozlov, T. Kuhr, M. Lupberger, S. Neuhaus, P. Niknejadi, J. Reindl, D. Schindler, A. Schneidewind, F. Schreiber, M. Schumacher, K. Schwarz, A. Streit, R.F.v. Cube, R. Walker, C. Walther, S. Wozniowski,

- K. Zhou, Resource-aware research on universe and matter: call-to-action in digital transformation. *Eur. Phys. J. Spec. Top.* (2024). <https://doi.org/10.1140/epjs/s11734-024-01436-4>
4. B. Kocot, P. Czarnul, J. Proficz, Energy-aware scheduling for high-performance computing systems: a survey. *Energies* **16**(2), 890 (2023). <https://doi.org/10.3390/en16020890>
 5. J. Muraña, S. Nesmachnow, F. Armenta, A. Tcherykh, Characterization, modeling and scheduling of power consumption of scientific computing applications in multicores. *Cluster Comput.* **22**(3), 839–859 (2019). <https://doi.org/10.1007/s10586-018-2882-8>
 6. P. Augier, C.F. Bolz-Tereick, S. Guelton, A.V. Mohanan, Reducing the ecological impact of computing through education and python compilers. *Nat. Astron.* **5**(4), 334–335 (2021). <https://doi.org/10.1038/s41550-021-01342-y>
 7. J.-P. Eckmann, S.O. Kamphorst, D. Ruelle, Recurrence plots of dynamical systems. *Europhys. Lett.* **4**(9), 973–977 (1987). <https://doi.org/10.1209/0295-5075/4/9/004>
 8. N. Marwan, M.C. Romano, M. Thiel, J. Kurths, Recurrence plots for the analysis of complex systems. *Phys. Rep.* **438**(5–6), 237–329 (2007). <https://doi.org/10.1016/j.physrep.2006.11.001>
 9. C.L. Webber Jr., N. Marwan, *Recurrence Quantification Analysis—Theory and Best Practices* (Springer, Cham, 2015). <https://doi.org/10.1007/978-3-319-07155-8>
 10. J.P. Zbilut, C.L. Webber Jr., Embeddings and delays as derived from quantification of recurrence plots. *Phys. Lett. A* **171**(3–4), 199–203 (1992). [https://doi.org/10.1016/0375-9601\(92\)90426-M](https://doi.org/10.1016/0375-9601(92)90426-M)
 11. N. Marwan, A historical review of recurrence plots. *Eur. Phys. J. Spec. Top.* **164**(1), 3–12 (2008). <https://doi.org/10.1140/epjst/e2008-00829-1>
 12. N. Marwan, A bibliographic view on recurrence plots and recurrence quantification analyses, in *Recurrence Plots and Their Quantifications: Methodological Breakthroughs and Interdisciplinary Discoveries*, ed. by Y. Hirata, M. Shiro, M. Fukino, C.L. Webber Jr., K. Aihara, N. Marwan (Springer, Cham, 2025), pp. 1–27. https://doi.org/10.1007/978-3-031-91062-3_1
 13. L. Poenaru-Olaru, J. Sallou, L. Cruz, J. Rellermeyer, A. van Deursen, Sustainable machine learning retraining: optimizing energy efficiency without compromising accuracy (2025). <https://doi.org/10.48550/arxiv.2506.13838>
 14. R. Rózycki, D.A. Solarska, G. Waligóra, Energy-aware machine learning models—a review of recent techniques and perspectives. *Energies* **18**(11), 2810 (2025). <https://doi.org/10.3390/en18112810>
 15. T. Rawald, M. Sips, N. Marwan, U. Leser, Massively parallel analysis of similarity matrices on heterogeneous hardware, in *Proceedings of the EDBT/ICDT Joint Conference 2015*, pp. 56–62 (2015). URL <http://ceur-ws.org/Vol-1330/#paper-11>
 16. T. Rawald, M. Sips, N. Marwan, PyRQA—conducting recurrence quantification analysis on very long time series efficiently. *Comput. Geosci.* **104**, 101–108 (2017). <https://doi.org/10.1016/j.cageo.2016.11.016>
 17. T. Rawald, M. Sips, N. Marwan, D. Dransch, Fast computation of recurrences in long time series, in *Translational Recurrences—From Mathematical Theory to Real-World Applications*, ed. by N. Marwan, M.A. Riley, A. Giuliani, C.L. Webber, Jr., vol. 103 (Springer, Cham, 2014), pp. 17–29. https://doi.org/10.1007/978-3-319-09531-8_2
 18. D. Schultz, S. Spiegel, N. Marwan, S. Albayrak, Approximation of diagonal line based measures in recurrence quantification analysis. *Phys. Lett. A* **379**(14–15), 997–1011 (2015). <https://doi.org/10.1016/j.physleta.2015.01.033>
 19. S. Spiegel, D. Schultz, N. Marwan, Approximate recurrence quantification analysis (arQA) in code of best practice, in *Recurrence Plots and Their Quantifications: Expanding Horizons*, ed. by C.L. Webber, Jr., C. Ioana, N. Marwan (Springer, Cham, 2016), pp. 113–136. https://doi.org/10.1007/978-3-319-29922-8_6
 20. G. Corso, T.D.L. Prado, G.Z. dos Santos Lima, J. Kurths, S.R. Lopes, Quantifying entropy using recurrence matrix microstates. *Chaos* **28**(8), 083108 (2018). <https://doi.org/10.1063/1.5042026>
 21. F.E.L. da Cruz, T.D.L. Prado, S.R. Lopes, N. Marwan, J. Kurths, Density-based recurrence measures from microstates. *Phys. Rev. E* **111**, 044212 (2025). <https://doi.org/10.1103/PhysRevE.111.044212>
 22. G.V. Ferreira, F.E.L. da Cruz, G. Marghoti, T.D.L. Prado, S.R. Lopes, N. Marwan, J. Kurths, RecurrenceMicrostateAnalysis.jl: a Julia library for analyzing dynamical systems with recurrence microstates. *Chaos* **35**(11), 113123 (2025). <https://doi.org/10.1063/5.0293708>
 23. R.V. Donner, Y. Zou, J.F. Donges, N. Marwan, J. Kurths, Recurrence networks—a novel paradigm for nonlinear time series analysis. *New J. Phys.* **12**(3), 033025 (2010). <https://doi.org/10.1088/1367-2630/12/3/033025>
 24. N. Marwan, J.F. Donges, Y. Zou, R.V. Donner, J. Kurths, Complex network approach for recurrence analysis of time series. *Phys. Lett. A* **373**(46), 4246–4254 (2009). <https://doi.org/10.1016/j.physleta.2009.09.042>
 25. M. Colafranceschi, M. Papi, A. Giuliani, G. Amiconi, A. Colosimo, Simulated point mutations in the A α -chain of human fibrinogen support a role of the α C domain in the stabilization of fibrin gel. *Pathophysiol. Haemost. Thromb.* **35**(6) (2006)
 26. E. García-Ochoa, J. González-Sánchez, N. Acuña, J. Euan, Analysis of the dynamics of Intergranular corrosion process of sensitised 304 stainless steel using recurrence plots. *J. Appl. Electrochem.* **39**(5), 637–645 (2009). <https://doi.org/10.1007/s10800-008-9702-4>
 27. N. Marwan, S. Foerster, J. Kurths, Analysing spatially extended high-dimensional dynamics by recurrence plots. *Phys. Lett. A* **379**(10–11), 894–900 (2015). <https://doi.org/10.1016/j.physleta.2015.01.013>
 28. M. Mosdorf, Method for detecting software anomalies based on recurrence plot analysis. *J. Theor. Appl. Comput. Sci.* **6**(1), 3–12 (2012) <http://www.jtacs.org/archive/2012/1/1>
 29. N. Marwan, K.H. Kraemer, Trends in recurrence analysis of dynamical systems. *Eur. Phys. J. Spec. Top.* **232**, 5–27 (2023). <https://doi.org/10.1140/epjs/s11734-022-00739-8>

30. O.E. Rössler, An equation for continuous chaos. *Phys. Lett. A* **57**(5), 397–398 (1976). [https://doi.org/10.1016/0375-9601\(76\)90101-8](https://doi.org/10.1016/0375-9601(76)90101-8)
31. N. Marwan, RQA_openmp. Zenodo (2025). <https://doi.org/10.5281/zenodo.16786351>
32. R. Martinović, G. Zitzlsberger, Highly scalable algorithm for computation of recurrence quantitative analysis. *J. Supercomput.* **75**, 1175–1186 (2019). <https://doi.org/10.1007/s11227-018-2350-5>
33. M. Sherrington, *Mastering Julia: Enhance your analytical and programming skills for data modeling and processing with Julia* (Packt Publishing Ltd, 2024)
34. Dynamics UFPR. RecurrenceMicrostatesAnalysis.jl. Github (2025). <https://github.com/DynamicsUFPR/RecurrenceMicrostatesAnalysis.jl>
35. J. von Kistowski, H. Block, J. Beckett, C. Spradling, K.-D. Lange, S. Kounev, Variations in CPU power consumption, in *Proceedings of the 7th ACM/SPEC on International Conference on Performance Engineering*, pp. 147–158 (2016). <https://doi.org/10.1145/2851553.2851567>
36. P. Icha, T. Lauf, Entwicklung der spezifischen Treibhausgas-Emissionen des deutschen Strommix in den Jahren 1990–2024. Technical report, Umweltbundesamt (2025). <https://www.umweltbundesamt.de/publikationen/entwicklung-der-spezifischen-treibhausgas-11>
37. J.B. Gao, Recurrence time statistics for chaotic systems and their applications. *Phys. Rev. Lett.* **83**(16), 3178–3181 (1999). <https://doi.org/10.1103/PhysRevLett.83.3178>
38. K.H. Kraemer, N. Marwan, Border effect corrections for diagonal line based recurrence quantification analysis measures. *Phys. Lett. A* **383**(34), 125977 (2019). <https://doi.org/10.1016/j.physleta.2019.125977>
39. F. Censi, G. Calcagnini, S. Cerutti, Proposed corrections for the quantification of coupling patterns by recurrence plots. *IEEE Trans. Biomed. Eng.* **51**(5), 856–859 (2004). <https://doi.org/10.1109/TBME.2004.826594>
40. N. Marwan, N. Wessel, U. Meyerfeldt, A. Schirdewan, J. Kurths, Recurrence plot based measures of complexity and its application to heart rate variability data. *Phys. Rev. E* **66**(2), 026702 (2002). <https://doi.org/10.1103/PhysRevE.66.026702>