

Product Jacobi-Theta Boltzmann machines with score matching

Andrea Pasquale^{1,2}, Daniel Krefl³, Stefano Carrazza^{1,2,4} and Frank Nielsen⁵

1 TIF Lab, Dipartimento di Fisica, Università degli Studi di Milano and INFN Sezione di Milano Via Celoria 16, 20133, Milan, Italy.

2 Quantum Research Center, Technology Innovation Institute, Abu Dhabi, UAE.

3 Department of Computational Biology, University of Lausanne, Switzerland

4 CERN, Theoretical Physics Department, CH-1211, Geneva 23, Switzerland.

5 Sony Computer Science Laboratories Inc, Tokyo, Japan

E-mail: andrea.pasquale@unimi.it, daniel.krefl@unil.ch, stefano.carrazza@unimi.it, frank.nielsen.x@gmail.com

Abstract. The estimation of probability density functions is a non trivial task that over the last years has been tackled with machine learning techniques. Successful applications can be obtained using models inspired by the Boltzmann machine (BM) architecture. In this manuscript, the product Jacobi-Theta Boltzmann machine (pJTBM) is introduced as a restricted version of the Riemann-Theta Boltzmann machine (RTBM) with diagonal hidden sector connection matrix. We show that score matching, based on the Fisher divergence, can be used to fit probability densities with the pJTBM more efficiently than with the original RTBM.

1. Introduction

The modelling of general probability density functions (PDFs) is a difficult task, even in the one-dimensional setting. Perhaps most commonly applied for this purpose is the method of kernel density estimation (KDE) [1]. Though often KDE yields a relatively good approximation to the underlying density in small dimensions, its general modelling capacity is limited as only the type of kernel and hyperparameter bandwidth are optimizable degrees of freedom.

More recently, significant success in modelling arbitrary probability densities, in particular in the high dimensional setting, has been achieved with new models like *variational autoencoders* (VAE) [2] and *normalizing flows* [3, 4]. Though these models shine in their modelling capacity and ability to easily draw samples from the modelled densities, obtaining derived quantities like moments, cumulative density functions, conditionals or marginalizations, is not straight-forward.

Unrelated to the above developments, another novel method has been proposed in recent years to model arbitrary probability densities [5]. The method is based on a generalization of the well-known Boltzmann machine [6], dubbed Riemann-Theta Boltzmann machine [5] (RTBM). The RTBM features a novel parametric density involving the Riemann-Theta (RT) function, which possesses a significant modelling capacity. The analytic expression for the density allows to explicitly calculate derived quantities. However, one drawback of the RTBM, which hindered so far a wider usage, is that it is computationally expensive, in particular for high dimensional

hidden state spaces. The reason being that the computational burden of the RT function scales exponentially with the dimension d .

As a step to overcome this computational challenge, we introduce in this work a simplified version of the RTBM by making use of a factorizing property of the RT function. This simplified version, which we will refer to as *product Jacobi-Theta Boltzmann machine*, is very suitable for a more efficient optimization (training) using the score matching method of [7].

As we will demonstrate experimentally on two and three-dimensional examples, this simplified model yields a significant improvement over the original RTBM based modelling.

1.1. Riemann-Theta Boltzmann machine

The Riemann-Theta Boltzmann machine (RTBM) is a novel variant of a Boltzmann machine, introduced in [5]. The RTBM is defined by the following energy model:

$$E(x) = \frac{1}{2}x^tAx + B^tx \quad \text{with} \quad x = \begin{pmatrix} h \\ v \end{pmatrix}, \quad B = \begin{pmatrix} B_h \\ B_v \end{pmatrix}, \quad A = \begin{pmatrix} Q & W^t \\ W & T \end{pmatrix}, \quad (1)$$

where A is a positive-definite $N \times N$ matrix and B a N dimensional bias vector. For reasons becoming more clear below, x , B and A are decomposed into a visible sector v of dimension N_v and a hidden sector h of dimension N_h so that $N = N_v + N_h$. In detail, Q corresponds to the connection matrix of the hidden sector, T to the connection matrix of the visible sector and W encodes the coupling between the two sectors. The Boltzmann distribution for this model is defined as $P(v, h) = \frac{e^{-E(v, h)}}{Z}$, where Z denotes the canonical partition function.

The peculiarity of the RTBM is that the hidden state space is taken to be \mathbb{Z}^{N_h} , while the visible state space is \mathbb{R}^{N_v} . This quantization of the hidden state space makes the computation of the canonical partition function tractable. In particular, by marginalizing out the hidden sector, a closed form analytic expression for the visible sector probability density function, denoted as $P(v)$, can be derived as [5]

$$P(v) = \sqrt{\frac{\det T}{(2\pi)^{N_v}}} e^{-\frac{1}{2}((v+T^{-1}B_v)^tT(v+T^{-1}B_v))} \frac{\tilde{\theta}(B_h^t + v^tW|Q)}{\tilde{\theta}(B_h^t - B_v^tT^{-1}W|Q - W^tT^{-1}W)}, \quad (2)$$

where $\tilde{\theta}$ is given by the RT function [8] with rescaled arguments, *i.e.*,

$$\tilde{\theta}(z|\Omega) := \theta\left(\frac{z}{2\pi i} \middle| \frac{i\Omega}{2\pi}\right) = \sum_{n \in \mathbb{Z}^k} e^{-\frac{1}{2}n^t\Omega n + n^tz}. \quad (3)$$

1.2. Product Jacobi-Theta Boltzmann machine

The general θ -function factorizes for a diagonal second argument matrix Ω into one-dimensional θ -functions, which are known as Jacobi-Theta functions: $\theta(z|\Omega) = \prod_i \theta(z_i|\Omega)_{ii}$. Consequently, $\tilde{\theta}$ also factorizes similarly. This factorization has significant computational advantages, as the computational complexity of the non-factorized θ grows exponentially with the dimension of Ω . For illustration, we plotted in Fig. 1 a run-time comparison between the factorized and non-factorized calculation of θ with diagonal Ω : While for small d the factorized computation run-time grows experimentally as $d^{1.2}$, the non-factorized shows clearly exponential growth under increasing k .

A diagonal Q in $P(v)$ corresponds to a hidden sector which is not inter-connected, and therefore is similar in spirit to the well-known restricted Boltzmann machine, *cf.*, [10]. However, for a RTBM with diagonal Q , which we will refer to as *product Jacobi-Theta Boltzmann Machine* (pJTBM), the θ -function occurring in the normalization of $P(v)$ does in general not factorize, as that would require that W diagonalizes T^{-1} . Hence, at first sight the computational advantage of a pJTBM is limited, as the normalization would still be computationally expensive.

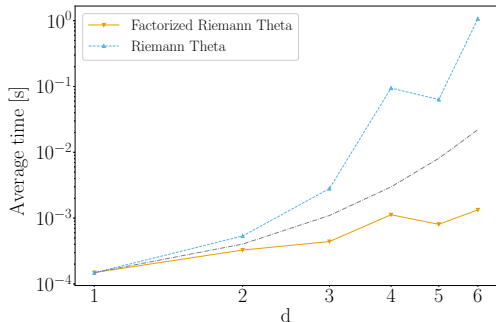


Figure 1. Average time to evaluate the RT function [9] for different dimensions d of a diagonal matrix Ω using the factorized form (yellow curve) and the standard form (light-blue curve). Note that both axes are plotted in logarithmic scale. The matrix elements of Ω have been sampled uniformly from the imaginary unit interval, and we averaged over 10 independent runs. Exponential growth is marked with a gray dashed-dotted line. A linear regression in log-log space implies that the average time for the factorized RT grows as $\propto d^{1.2}$ with $R^2 = 0.93$.

2. The Fisher cost function

2.1. Introduction

In the original work introducing the RTBM [5] it has been shown that the probability density of the visible sector $P(v)$ can be used to approximate the underlying density of a given dataset via maximum likelihood estimation (MLE).

The novel contribution of this work consists in adopting a different learning method for the RTBM known as *score matching*, first introduced in [7]. Score matching is particularly efficient for non-normalized models, *i.e.*, models where the partition function Z cannot be computed either analytically or numerically. For instance, this can be the case for high-dimensional densities where the computation of the partition function requires the calculation of a non-trivial multi-dimensional integral. This parameter learning method can also be exploited for models where the partition function can be computed, such as the RTBM, but is computationally expensive: Avoidance of the computation of the partition function during optimization may strongly accelerate the training process. In particular, this applies to the pJTBM defined above, for which the normalization is the main remaining computational obstacle.

Suppose that we are trying to model the probability density $p(x)$ using a parametric density $q(x, \theta)$, where θ are the parameters to be optimized. The divergence metric associated with score matching [7] reads

$$D_F(p||q_\theta) = \int p(x) \left| \frac{\nabla_x p(x)}{p(x)} - \frac{\nabla_x q(x, \theta)}{q(x, \theta)} \right|^2 dx, \quad (4)$$

which is often referred to as Fisher divergence or relative Fisher information [11]. It has been shown [7, 12] that if $q(x, \theta)$ is sufficiently regular and differentiable, then we can simplify the above expression via integration by parts. For a finite sample v of $p(x)$ we obtain (up to an irrelevant constant)

$$D_F(p||q_\theta) \approx \mathcal{C}^F = \sum_{i=1}^N \left| \nabla_{v_i} \log q(v_i, \theta) \right|^2 + 2\Delta_{v_i} \log q(v_i, \theta), \quad (5)$$

where ∇ and Δ denotes the gradient and Laplacian operators, respectively. We refer to \mathcal{C}^F as the Fisher cost function.

2.2. Derivation of the Fisher cost function for the RTBM and pJTBM

If we model the probability density $p(x)$ using a RTBM, we can exploit the fact that we have an analytical expression for $q(v, \theta)$, given by $P(v)$ in Eq. 2, to compute explicitly the Fisher cost function. The two terms in Eq. 5 for \mathcal{C}^F are given for the RTBM by

$$\partial_{v_i} \log P(v) = -(Tv)_i - (Bv)_i + (WD)_i, \quad \partial_{v_i}^2 \log P(v) = -T_{ii} + (WHW^t)_{ii} + (WD)_i^2, \quad (6)$$

with D the normalized gradient and H the normalized Hessian:

$$(D)_i = \frac{\nabla_i \tilde{\theta}(B_h^t + v^t W | Q)}{\tilde{\theta}(B_h^t + v^t W | Q)}, \quad (H)_{ij} = \frac{\nabla_i \nabla_j \tilde{\theta}(B_h^t + v^t W | Q)}{\tilde{\theta}(B_h^t + v^t W | Q)}. \quad (7)$$

The calculation of D and H can be computationally expensive, especially in the case of $N_h > 1$. However, for the pJTBM, *i.e.*, a *restricted* RTBM with diagonal Q matrix, we can simplify the above expressions using the fact that the θ -function (and consequently also $\tilde{\theta}$) factorizes. This yields

$$\begin{aligned} \partial_{v_i} \log P(v) &= -(Tv)_i - (Bv)_i + \sum_{j=1}^{N_h} \frac{\partial_{v_i} \tilde{\theta}((B_h^t + v^t W)_j | Q_{jj})}{\tilde{\theta}((B_h^t + v^t W)_j | Q_{jj})} W_{ji}, \\ \partial_{v_i}^2 \log P(v) &= -T_{ii} + \sum_{j=1}^{N_h} \frac{\partial_{v_i}^2 \tilde{\theta}((B_h^t + v^t W)_j | Q_{jj})}{\tilde{\theta}((B_h^t + v^t W)_j | Q_{jj})} W_{ji}^2 - \sum_{j=1}^{N_h} \frac{(\partial_{v_i} \tilde{\theta}((B_h^t + v^t W)_j | Q_{jj}))^2}{\tilde{\theta}((B_h^t + v^t W)_j | Q_{jj})} W_{ji}. \end{aligned} \quad (8)$$

Note that, the above simplification is in a computational sense. The complexity of the Fisher cost function is now growing linearly with N_h , because we have only to compute derivatives of Jacobi-Theta functions. This motivates us to adopt a dedicated implementation for the Jacobi-Theta function, which is more efficient, instead of using the RT function implementation of [9]. For more details about the implementation of the Jacobi-Theta function we refer to the Appendix.

3. Example

As an example we employ RTBMs to model the empirical bivariate distribution given by the `uranium` dataset [13], which is included in the R package `copula` [14]. We also fit the models to the empirical joint density of daily log returns of three different equities: `AAPL`, `MSFT` and `GOOG`.

We compare the performance of the standard RTBM, *i.e.*, a generic RTBM optimized via maximum likelihood estimation of parameters, with a pJTBM trained using the Fisher cost function, as described in section 2. Note that we make use of the transformation property of the RTBM under affine transformations [15] to preprocess the data using a z-score normalization and PCA [16]. In particular, we perform a pre-training step by fitting first the marginals of the underlying distribution using 1-dimensional RTBM or pJTBM.

At hand of these examples, we want to address two main points: firstly, we aim to quantify the speed-up achieved using this new learning methodology by benchmarking the time necessary to run 500 iterations using CMA-ES as optimizer [17]. Secondly, we want to compare the quality of the fits by using as *Goodness-of-Fit* test a generalized version of the Kolmogorov-Smirnov (KS) test called Fasano-Franceschini test (FF) [18]. In particular, we are interested to check if the pJTBM can achieve a comparable level of accuracy, despite the fewer parameters due to diagonal Q .

Results are given in Table 1. We observe that the time required to train a pJTBM is exponentially suppressed compared to the RTBM. This enabled us to push the number of hidden nodes up to 8 with acceptable computational times under our used optimization scheme. Note that for the RTBM we could not go beyond $N_h = 4$ in reasonable time.

The FF values are in general comparable between RTBM and pJTBM under the error bounds. However, we observe that the variance for the pJTBM is in general larger compared to the RTBM, despite the lower number of parameters. A possible explanation could be that a model with diagonal Q is more rigid and therefore has more difficulties to reach a good solution from a given initial condition.

	N_h	parameters	Execution Time	FF
pJTBM	2	13	9.2 s	2.32 ± 0.23
RTBM	2	14	59.6 s	1.72 ± 0.10
pJTBM	4	21	20.6 s	1.81 ± 0.32
RTBM	4	27	5 min	1.50 ± 0.08
pJTBM	6	29	51.6 s	1.63 ± 0.22
RTBM	6	44	-	
pJTBM	8	37	82.9 s	1.76 ± 0.32
RTBM	8	65	-	
pJTBM	3	24	32.2 s	2.91 ± 0.24
RTBM	3	27	6 min	2.41 ± 0.43

Table 1. Execution time and FF values for RTBM and pJTBM with different number of hidden nodes N_h . Each result is averaged over 10 independent runs. The results with N_h even correspond to fitting the models to the `uranium` dataset, while for $N_h = 3$ the models are fitted to the joint `AAPL`, `MSFT` and `GOOG` stock daily log return distribution. The error on the FF corresponds to the standard deviation between 10 runs of the training process. A – indicates that the computation did not finish in reasonable time.

We believe that under a different optimization scheme (optimization on manifolds), the pJTBM could become less dependent on initial condition and also more suitable for far higher-dimensional problems, leading to interesting applications. We hope to investigate this aspect in future work.

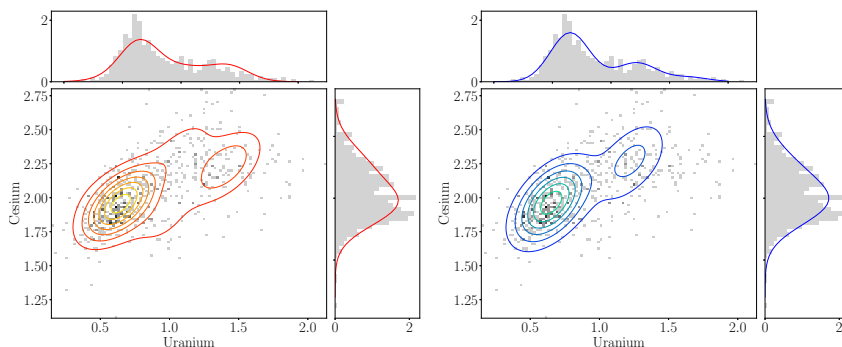


Figure 2. Contour plots for the pJTBM (left, red) and for the RTBM (right, blue) for the `uranium` dataset. The $N_h = 4$ model with the lowest FF value over 10 independent runs is plotted.

Appendix

The computation of the RT function and its derivatives is computationally challenging, as the function is defined as an infinite sum over a N -dimensional integer lattice \mathbb{Z}^N , *cf.*, Eq. 3. For a general dimension N , the RT function and its derivatives can be calculated up to any desired precision ϵ via summing over a finite subset of lattice points falling inside an ellipsoid of radius $R(\epsilon)$ [19, 20].

For the purpose of this work we mainly need an efficient implementation of the RT function and its derivatives for $N = 1$. A highly efficient implementation has been proposed in [21], albeit not including the derivatives. For simplicity, we only consider here the naive algorithm given in [21], because it can be easily generalized to compute as well the derivatives, and gives us already a significant performance gain compared to [19], see Fig. A1.

The naive algorithm for evaluation of the 1d θ function starts from the partial summation

$$S_B(z, \Omega) = 1 + \sum_{0 < n < B} q^{n^2} (e^{2\pi i n z} + e^{-2\pi i n z}) =: 1 + \sum_{0 < n < B} v_n, \quad (\text{A.1})$$

where $q = e^{i\pi\Omega}$ and $v_n = q^{n^2} (e^{2\pi i n z} + e^{-2\pi i n z})$. One can prove that for specific ranges of Ω and z there exist a $B(\epsilon)$ such that an approximation to the θ function with absolute precision ϵ is

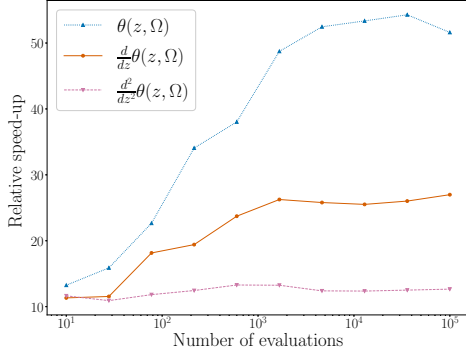


Figure A1. Relative speed-up factor (t_1/t_2) between the execution time of [19] (t_1), and the naive algorithm from [21] extended to the calculation of the derivatives (t_2), as described in the Appendix. The parameters of the RT function have been sampled uniformly from the imaginary unit interval and we averaged over 10 independent runs.

obtained. Once bounded, S_B can be computed via exploiting the fact that the summands v_n can be computed recursively, *i.e.*, $v_{n+1} = q^{2n}v_1v_n - q^{4n}v_{n-1}$. In the following, we will generalize this recursive algorithm to the computation of the derivatives. We made an implementation of this algorithm available inside the *Theta* (v0.0.2) Python package [22].

Appendix A.1. Partial derivatives of the 1d Riemann-Theta function

The recursive algorithm recalled above can be easily generalized as follows: First, note that it is not hard to see that the partial sums for the first and second derivative read

$$\begin{aligned} U_B(z, \Omega) &= \sum_{1 < n < B} -4\pi n q^{n^2} \sin(2\pi n z) =: \sum_{1 < n < B} w_n, \\ V_B(z, \Omega) &= \sum_{1 < n < B} -8\pi^2 n^2 q^{n^2} \cos(2\pi n z) =: \sum_{1 < n < B} \xi_n. \end{aligned} \quad (\text{A.2})$$

The recurrence relations for w_n and ξ_n can be derived quite easily by making use of Chebyshev's polynomials. This yields

$$w_{n+1} = (n+1) \left[\frac{2 \cos(2\pi z)}{n} q^{2n+1} w_n - \frac{q^{4n}}{n-1} w_{n-1} \right], \quad (\text{A.3})$$

and

$$\xi_{n+1} = (n+1)^2 \left[\frac{2 \cos(2\pi z)}{n^2} q^{2n+1} \xi_n - \frac{1}{(n-1)^2} q^{4n} \xi_{n-1} \right]. \quad (\text{A.4})$$

It remains to derive the ranges of validity for which the partial sums for U_B and V_B give an accurate approximation for the gradients of the RT function:

Proposition 1 *Suppose that $\Im(\tau) \geq 0.742$ and $0 \leq \Im(z) \leq \Im(\tau)/2$.*

Then, for $B \geq 1$, $|\frac{d}{dz}\theta(z, \tau) - U_B(z, \tau)| \leq 3|q|^{(B-1)^2}$, where

$$U_B(z, \tau) = \sum_{0 < n < B} -4\pi n q^{n^2} \sin(2\pi n z). \quad (\text{A.5})$$

Proof. In order to demonstrate the proposition we bound the remainder of the series:

$$\begin{aligned} \left| \frac{d}{dz}\theta(z, \tau) - U_B(z, \tau) \right| &\leq 4\pi \sum_{n \geq B} |q|^{n^2} \frac{n}{2} |e^{2\pi i n z} - e^{-2\pi i n z}| \leq 4\pi \sum_{n \geq B} n |q|^{n^2 - n} \\ &\leq 4\pi \sum_{n \geq B} n |q|^{(n-1)^2} \leq 8\pi |q|^{(B-1)^2} \sum_{n \geq 0} n |q|^n = 8\pi \frac{|q|^{(B-1)^2 + 1}}{(1 - |q|)^2}. \end{aligned} \quad (\text{A.6})$$

Numerically it can be shown that for $\Im(\tau) \geq 0.742$, we have $\frac{8\pi|q|}{(1-|q|)^2} \leq 3$, which proves the proposition.

Proposition 2 *Suppose that $\Im(\tau) \geq 0.882$ and $0 \leq \Im(z) \leq \Im(\tau)/2$. Then, for $B \geq 1$, $|\frac{d^2}{dz^2}\theta(z, \tau) - V_B(z, \tau)| \leq 3|q|^{(B-1)^2}$, where*

$$V_B(z, \tau) = \sum_{0 < n < B} -8\pi^2 n^2 q^{n^2} \cos(2\pi n z). \quad (\text{A.7})$$

Proof. We bound the remainder of the series as follow:

$$\begin{aligned} \left| \frac{d^2}{dz^2} \theta(z, \tau) - V_B(z, \tau) \right| &\leq 8\pi^2 \sum_{n \geq B} |q|^{n^2} \frac{n^2}{2} |e^{2\pi i n z}| + |e^{-2\pi i n z}| \leq 8\pi^2 \sum_{n \geq B} n^2 |q|^{n^2 - n} \\ &\leq 8\pi^2 \sum_{n \geq B} n^2 |q|^{(n-1)^2} \leq 32\pi^2 |q|^{(B-1)^2} \sum_{n \geq 0} n^2 |q|^n = \frac{64\pi^2 |q|^{(B-1)^2 + 2}}{(1-|q|)^3}. \end{aligned} \quad (\text{A.8})$$

Numerically it can be shown that for $\Im(\tau) \geq 0.882$, we have $\frac{8\pi|q|^2}{(1-|q|)^3} \leq 3$, which proves the proposition.

References

- [1] Rosenblatt M 1956 *The Annals of Mathematical Statistics* **27** 832 – 837
- [2] Kingma D and Welling M 2013 *ICLR*
- [3] Rezende D J and Mohamed S 2015 *Proceedings of the 32nd International Conference on International Conference on Machine Learning - Volume 37 ICML'15 (JMLR.org)* p 1530–1538
- [4] Kobyzev I, Prince S J and Brubaker M A 2020 *IEEE transactions on pattern analysis and machine intelligence* **43** 3964–3979
- [5] Krefl D, Carrazza S, Haghghat B and Kahlen J 2020 *Neurocomputing* **388** 334–345
- [6] Hinton G E and Sejnowski T J 1983 *Proceedings of the Fifth Annual Conference of the Cognitive Science Society, Rochester NY*
- [7] Hyvärinen A 2005 *Journal of Machine Learning Research* **6** 695–709
- [8] Mumford D 2007 *Tata Lectures on Theta I* (Springer) ISBN 978-0-8176-4572-4
- [9] Carrazza S and Krefl D openRT, an open-source implementation of the Riemann-Theta function
- [10] Ackley D H, Hinton G E and Sejnowski T J 1985 *Cognitive Science* **9** 147–169 ISSN 0364-0213
- [11] Otto F and Villani C 2000 *Journal of Functional Analysis* **173** 361–400
- [12] Lyu S 2009 *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence UAI '09* (Arlington, Virginia, USA: AUAI Press) p 359–366 ISBN 9780974903958
- [13] Cook R D and Johnson M E 1986 *Technometrics* **28** 123–131
- [14] Hofert M, Kojadinovic I, Maechler M and Yan J 2020 **copula**: Multivariate Dependence with Copulas
- [15] Carrazza S and Krefl D 2020 *Comput. Phys. Commun.* **256** 107464
- [16] Pearson K 1901 *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* **2** 559–572
- [17] Hansen N and Ostermeier A 2001 *Evolutionary Computation* **9** 159–195
- [18] Fasano G and Franceschini A 1987 *Monthly Notices of the Royal Astronomical Society* **225** 155–170
- [19] Deconinck B, Heil M, Bobenko A, van Hoeij M and Schmies M 2002 Computing Riemann Theta Functions
- [20] Swierczewski C and Deconinck B 2016 *Mathematics and Computers in Simulation* **127** 263–272 ISSN 0378-4754 special Issue: Nonlinear Waves: Computation and Theory-IX
- [21] Labrande H 2018 *Mathematics of Computation* **87** 1479–1508
- [22] Carrazza S and Krefl D 2017 RiemannAI/theta: theta v0.0.1