nature portfolio

Peer Review File

Reservoir Direct Feedback Alignment: Deep Learning by Physical Dynamics

Corresponding Author: Dr Mitsumasa Nakajima

This file contains all reviewer reports in order by version, followed by all author rebuttals in order by version.

Version 0:

Reviewer comments:

Reviewer #1

(Remarks to the Author)
Nakajima et al report in their manuscript ‘Reservoir Direct Feedback Alignment: Deep Learning by Physical Dynamics’ the training of deep neural networks (DNN) using the nonlinear dynamics of a reservoir computer (RC) as the information used for updating the DNN’s weight. The RC is driven by information related to the error signal obtained at the DNN’s final layer, and the weights of the DNN are updated using the product between a constant random matrix and the RC’s states. In order to augment the dimensionality of the RC, the authors furthermore time-multiplex this update factors by using the temporal evolution of the RC. This allows to design a RC that is smaller than the dimensionality of the DNN’s parameter space.
The manuscript provides an extensive amount of data and analysis, comparing numerical models with an experimental hardware implementation using a photonic RC implemented in a delay optoelectronic system. The work is of very high quality and extremely timely, as the question on how to train DNN’s using concepts that are better suited for unconventional hardware than error back propagation (BP) is one of the main concerns in the field of unconventional computing. The manuscript is well written. As such, I recommend publication after minor modifications.

More detailed comments:

One of the points I find a bit lacking at the moment is the discussion of hardware affinity. Here, the authors strongly focus on the comparison to classical BP with they RC-DFA method. However, that has been discussed numerous times before in the manuscript. Actually, they authors repeat that point a bit much, and I would prefer if this is just mentioned at the beginning, and then all discussion related to that point are exclusively discussed at that section. However, I find that one of the most interesting questions is why one would use the RD-DFA, and not simply DFA in the first place. This point should be included specifically in the sections related to hardware affinity and energy efficiency / scaling.

‘the discussion of the figure is described in later’ -> ‘The discussion of the figure is described later’

Figure 3: label for the colorbar is missing

In general: for each plots and tables that report MSE or accuracies the particular test as well as the DNN’s or RC’s size should be specified in the caption.

Why is the test accuracy in Fig 4c higher than the training accuracy in Fig 4b?

‘As DFA does not requires’ -> ‘As DFA does not require’

Reviewer #3

(Remarks to the Author)
This manuscript describes a method to train deep neural networks based on the nonlinear dynamics of physical reservoirs. This so-called reservoir computing direct feedback alignment (RC-DFA) method is first used numerically to train a multi-layer perceptron using different models of reservoirs and extensively characterizing different parameter dependencies. Second, a PIC-based optoelectronic delay reservoir is used to experimentally demonstrate the feasibility of the present training approach. The authors demonstrate that performance of the RC-DFA method for training deep learning models is

comparable to classic back-propagation and to their previously published direct feedback alignment methods. A particularly interesting result of the present work is the possibility to use nonlinear dynamics at the edge of chaos (or fully chaotic, if certain time scales requirements are met) to train large deep neural networks.

I think the work is convincing and relevant for other experts in the topic, however the authors should further elaborate on the differences between the current approach and their previously presented approach for augmented DFA. Two are the main points I refer: the requirements for offline analysis of the errors minimization (both in DFA and RC-DFA) and their different potential for scalability. Particularly the fact that (following the authors discussion) their RC-DFA method seems to only be able to reduce the number of required nodes by one order of magnitude, and this is already leveraging the time multiplexing.

Version 1:

Reviewer comments:

Reviewer #1

(Remarks to the Author)
I would like to thank the authors for addressing my comments.

I recommend publication of the manuscript.

Reviewer #3

(Remarks to the Author)
The authors have addressed and considered in detail all my previous comments. I recommend the manuscript for publication.

October 2nd, 2024
Communications Physics
The Campus
4 Crinan Street, London N1 9XW
United Kingdom

Dear Editor,

We deeply thank for the editors and reviewers for taking the time to check our manuscript. We are also sincerely grateful for your thoughtful comments on our manuscript. We have revised the manuscript according to your advice. We have read each comment carefully and have made the corresponding corrections, which we hope to address your concerns. Before responding each referee's comments point-by-point, we would first like to summarize the main changes that we have made to the manuscript and supplementary materials. All the corrections in the main and supplementary articles are highlighted in red.

Summary of changes:

(1) We added descriptions of scalability and energy efficiency for RC-DFA in the discussion part.
(2) We added the section "Hyperparameter optimization of RC-DFA" in discussion part.
(3) We added "Required number of active components for physical implementation" as Supplementary Information S2 for describing the physical scalability of RC-DFA.
(4) We added detailed additional explanations to clarify the reviewer's concerns.
(5) An English rewriting service corrected the words and expressions where needed.

To Reviewer #1

Comment 1

Reviewer #1 (Remarks to the Author):

Nakajima *et al* report in their manuscript 'Reservoir Direct Feedback Alignment: Deep Learning by Physical Dynamics' the training of deep neural networks (DNN) using the nonlinear dynamics of a reservoir computer (RC) as the information used for updating the DNN's weight. The RC is driven by information related to the error signal obtained at the DNN's final layer, and the weights of the DNN are updated using the product between a constant random matrix and the RC's states. In order to augment the dimensionality of the RC, the authors furthermore time-multiplex this update factors by using the temporal evolution of the RC. This allows to design a RC that is smaller than the dimensionality of the DNN's parameter space.
The manuscript provides an extensive amount of data and analysis, comparing numerical models with an experimental hardware implementation using a photonic RC implemented in a delay optoelectronic system. The work is of very high quality and extremely timely, as the question on how to train DNN's using concepts that are better suited for unconventional hardware than error back propagation (BP) is one of the main concerns in the field of unconventional computing. The manuscript is well written. As such, I recommend

publication after minor modifications.

We thank the referee for the positive feedback on our manuscripts. We sincerely appreciate the time and effort you have taken to review our work. We could re-recognize the main advantage of our paper to the relevant fields thanks to your comment. We have carefully checked your comments and revised the paper accordingly. We believe that the revisions address your concerns.

Comment 2:

One of the points I find a bit lacking at the moment is the discussion of hardware affinity. Here, the authors strongly focus on the comparison to classical BP with they RC-DFA method. However, that has been discussed numerous times before in the manuscript. Actually, they authors repeat that point a bit much, and I would prefer if this is just mentioned at the beginning, and then all discussion related to that point are exclusively discussed at that section. However, I find that one of the most interesting questions is why one would use the RD-DFA, and not simply DFA in the first place. This point should be included specifically in the sections related to hardware affinity and energy efficiency / scaling.

Thank you very much for your valuable feedback on our manuscript. In response to your comment regarding the comparison with BP, we agree that we were somewhat redundant in the results section. As you suggested, we have removed the description from the results section and merged it into the discussion section, ensuring that the point is addressed more cohesively. Additionally, we have expanded the discussion section to include the energy efficiency of RC-DFA, addressing the point raised regarding hardware affinity and why one would choose RC-DFA over traditional DFA. We hope these changes clarify the distinctions and provide a more robust discussion on the topic. As this comment is partially related to Reviewer 2's comment #2, we believe the answer to this question may support the further understanding of RC-DFA.

Changes:

- We deleted the description of comparison between the BP and RC-DFA from the resultx section and merged it into "Hardware affinity of training algorithm" in the discussion section.
- We added a description of energy efficiency in "Computational efficiency" in the discussion section.

Comment 3:

'the discussion of the figure is described in later' -> 'The discussion of the figure is described later'

Thank you for pointing out the writing issue. We checked the manuscript once again and rewrote some sentences again. In addition, the manuscript has also once again been reviewed by technical rewriting service.

Changes:

- We revised the description that you pointed out.
- We used English rewriting services to check for grammatical errors and typos.

Comment 4:

Figure 3: label for the colorbar is missing

Thank you for pointing out the missing color bar.

Change:

- We added the colorbar in the Fig. 3.

Comment 5:

In general: for each plots and tables that report MSE or accuracies the particular test as well as the DNN's or RC's size should be specified in the caption.

We appreciate your pointing out the lack of detailed description in the figures.

Change:

- We added the RC size in the figure captions.

Comment 6:

Why is the test accuracy in Fig 4c higher than the training accuracy in Fig 4b?

Figure 4b shows the plot within a single epoch (i.e. 600 minibatch). Fig. 4c shows the accuracies after five-epoch training. Thus, the results in 4c are superior to the ones in 4b. To prevent the misunderstanding, we added the description.

Change:

- We added an explanation of why the results in Fig. 4(c) are superior to the ones in Fig. 4(b).

Comment 7:

'As DFA does not requires' -> 'As DFA does not require'

Thank you for pointing out the writing issue. We checked the manuscript once again and rewrote some sentences again. In addition, the manuscript has also once again been reviewed by a technical rewriting service.

Changes:

- We revised the description that you pointed out.
- We used an English rewriting service to check for grammatical errors and typos.

Reviewer #2 (Remarks to the Author):

Comment 1

This manuscript describes a method to train deep neural networks based on the nonlinear dynamics of physical reservoirs. This so-called reservoir computing direct feedback alignment (RC-DFA) method is first used numerically to train a multi-layer perceptron using different models of reservoirs and extensively characterizing different parameter dependencies. Second, a PIC-based optoelectronic delay reservoir is used to experimentally demonstrate the feasibility of the present training approach. The authors demonstrate that performance of the RC-DFA method for training deep learning models is comparable to classic back-propagation and to their previously published direct feedback alignment methods. A particularly interesting result of the present work is the possibility to use nonlinear dynamics at the edge of chaos (or fully chaotic, if certain time scales requirements are met) to train large deep neural networks.

We thank the referee for the positive feedback on our manuscript, and we deeply thank the reviewer for taking the time to check our manuscript and providing insightful comments. We have carefully checked your comments and revised the paper accordingly. We believe that the revisions address your concerns.

Comment 2:

I think the work is convincing and relevant for other experts in the topic, however the authors should further elaborate on the differences between the current approach and their previously presented approach for augmented DFA. Two are the main points I refer: the requirements for offline analysis of the errors minimization (both in DFA and RC-DFA) and their different potential for scalability. Particularly the fact that (following the authors discussion) their RC-DFA method seems to only be able to reduce the number of required nodes by one order of magnitude, and this is already leveraging the time multiplexing.

Thank you for your valuable feedback on our paper. We have addressed your points as follows:

On the first point regarding "the requirements for offline analysis of the errors minimization":

We appreciate your insightful comment. We understand your comment is referring to the requirements to minimize the test error for RC-DFA and standard DFA, especially how to select the hyperparameters. Regarding the standard DFA, it just uses the linear random projection of the final error to compute the quasi-errors in each layer. Thus, there are few hyperparameters for improving the performance, excluding the learning rate and hyperparameters in the target DNN model itself. In general, the learning rate should be small enough to suppress the divergence of the training, but too small a learning rate requires many training iterations. Thus, we should select the learning rate by considering the performance and experimental resources. The learning rate for RC-DFA can be determined with the same strategy. For our RC-DFA experiment, we used the learning rate of 0.0005 for the 300-epoch training and 0.01 for the 20-epoch training.

On the other hand, in the case of RC-DFA, we have additional freedom for the training because the quasi-errors for each layer are computed by the nonlinear dynamics. From the numerical analysis, we clarified that two unique hyperparameters are the critical factors for the RC-DFA training. One is the spectral radius (SR), and the other is the time step for recording the temporal evolution in a reservoir (*T*). Regarding the SR, the performance was maximized near SR=1, where there is a boundary region between chaotic and ordered dynamics (so-called edge of chaos) as shown in Fig. 2(a)–(c). However, this region also tends to require longer training iteration steps to converge [Fig. 2(d)]. Therefore, SR values slightly below 1 (e.g., 0.4–0.8) should be selected to achieve a balance between convergence speed and accuracy improvement. From the alignment angle analysis in the Figs. 3(f)–(h), the region near SR=1 tends to result in the smallest average angles. This means that the average direction of weight updates in the edge chaos region is close to the direction of the BP updates. However, an increase in the outliers for the alignment angles near SR=1 was also observed. The angles of outliers were typically over 90 degrees. This suggests that a nonnegligible amount of the update direction did not agree with the direction of the BP update, which contributes to the degradation in accuracy. Similar behavior was observed in the chaotic regime where SR>1, with an increasing number of cases where alignment angles approached or exceeded 90 degrees. Thus, the alignment angle analysis also suggests that SR values should be slightly below 1. Regarding the optimum *T* value, we observed that too short a time led to accuracy degradation, while too long a time resulted in accuracy deterioration, particularly in the small-SR region. However, when the SR was in the preferred range (SR = 0.4–0.8), stable performance was obtained across a relatively wide range of time steps (T = 5–50). Therefore, selecting time steps within this range seems appropriate from the viewpoint of accuracy test optimization. In physical implementations, if *T* is too long, the main signal decays, leading to a reduced signal-to-noise ratio. Thus, choosing a shorter *T* within this range is practical. The impact of noise is analyzed in Supplementary Information S2.

As the reviewer pointed out, the above-described points for optimizing the hyper parameter were not clear in the original manuscript. Therefore, we have added a “Hyperparameter optimization” section in the discussion section. Additionally, if your comment was pointing to the unclear experimental setup for the offline analysis in Fig. 3(a)−(e), we have added some details of the experimental configuration to address this as well.

<u>Change:</u>

- We added "Hyperparameter optimization of RC-DFA" in the discussion section.
- We added some details of the experimental configuration for the offline analysis in Fig. 3(a)−(e).

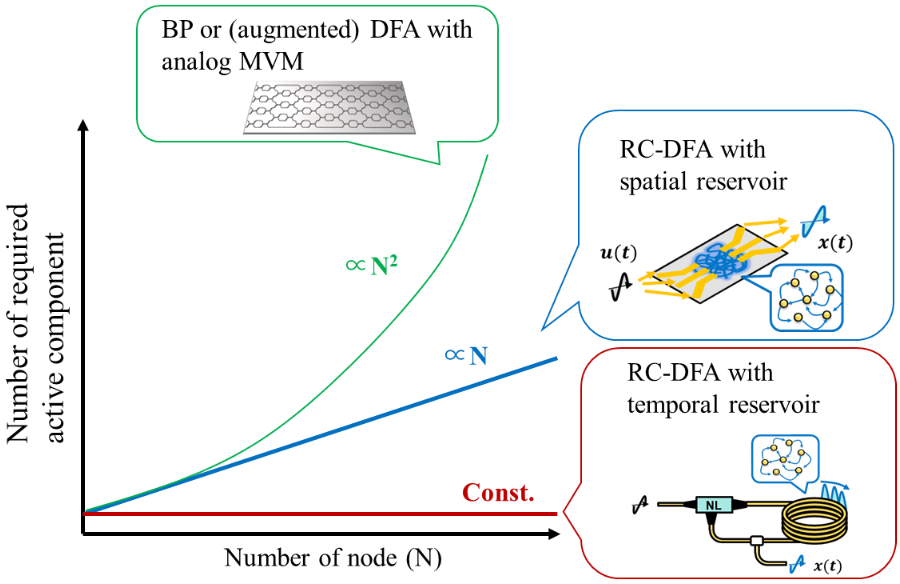

- ***Fig. S2. Required physical component and node count.*** *Number of required active elements to implement the training procedure on a specific physical computing device. BP, DFA, and augmented DFA require programmable matrix vector multiplication (MVM), which requires O(N^2) active components (e.g., MZI, MRR for optical implementation) and O(N) of input/output components (e.g., laser, modulator, and detectors for optical implementation). On the other hand, RC-DFA does not require programmable devices for MVM. Instead, passive random or fixed connections like a ring topology can be used, which drastically reduce the number of required active component to O(N) in the case of spatial reservoir computer. By selecting temporal reservoir implementation, the number of required components becomes independent of the number of nodes, which has significant advantage for hardware affinity.*

On the second point regarding the scalability differences between RC-DFA and DFA:

We thank you for pointing out the importance of further discussion about the scalability differences between RC-DFA and DFA (or BP) in the original manuscript. Our method aims not only to reduce the number of required computational nodes as you pointed out, but also to provide easier and more energy efficient physical implementation by leveraging the concept of physical RC. To explain the later advantage directly, we have added the above figure and related description as supplementary material S2 (please see Fig. S2). The figure illustrates the requirement of physical active components (e.g., MZI, MRR for optical implementation) for implementing of BP, DFA, and RC-DFA as a function of node count. The BP and DFA require $O(N^2)$ and $O(N_f N)$ active elements for solving the matrix vector multiplication (MVM), where N *and* N_f are the node numbers of target layer and final layer (e.g. N=1000 and N_f=10 in the fully connected network for the 10-class classification task). By assuming N=1000 and N_f=10, they require reconfigurable active components on the order of 10^6 (for BP) and 10^4 (for DFA) for executing the MVM. These values are still highly challenging for the physical implementation; e.g., the recent maximum number of active components in a photonic integrated circuit is around 10^4 [R1].

In contrast, our proposed RC-DFA can be operated on the passive nonlinear circuit (i.e. physical reservoir) with the number of reservoir node N', which does not require the active component for the each MVM operation. Thus, in general, only N' modulation elements for signal input (e.g. optical modulators), N' receiving elements for signal output (e.g. photodetectors), and N' nonlinear components (e.g. optical amplifiers) are required. Thus, we can reduce the required number of active components from $O(N^2)$ and $O(MN)$ to $O(N')$ even using the standard spatially multiplexed physical RC circuit [R2]. In addition, we can select an N' value that is much smaller than N (e.g. $N'=N/10$ in our experiment). Thus, we can reduce the number of active components from 10^6 (for BP) and 10^4 (for DFA) to 10^2 (for RC-DFA) in the case of the above-described example. Interestingly, we can further reduce the required component by using the time domain multiplexing scheme [R3]. In this scheme, the reservoir node is multiplexed in a single physical lane (e.g. optical waveguide or fiber) by dividing each time slot in the physical lane as reservoir nodes. Thus, this scheme requires only a single modulator, detector, and nonlinear component, which drastically reduces the required number of active components (in principle, it does not depend on node number N). It is also worth to note that the number of active components is highly related to energy consumption for the physical analog computation. Since the RC-DFA can reduce the number of active components drastically, it is more effective than BP and standard DFA even from the viewpoint of energy efficiency. We have added a discussion about energy efficiency in “Computational efficiency” in the discussion section in the revised manuscript.

The discussions about the scalability of RC-DFA were not clear in the original manuscript as the reviewer pointed out. Therefore, we have added “Required number of active components for physical implementation” in the supplementary material S2. In addition, we have added a description about the scalability and energy efficiency in “Hardware affinity of training algorithm” and “Computational efficiency” in the discussion section in the revised manuscript.

[R1] Shekhar, S. *et al.* "Roadmapping the next generation of silicon photonics." *Nature Commun.* 15, 751 (2024).
[R2] Vandoorne, K. *et al*. "Experimental demonstration of reservoir computing on a silicon photonics chip." *Nature Commun.* 5, 3541(2014).
[R3] Larger, L. *et al.* "Photonic information processing beyond Turing: an optoelectronic implementation of reservoir computing." *Opt. Express* 20, 3241-3249 (2012).

Changes:

- We added "Required number of active components for physical implementation" in Supplemental Information S2.
- We added a description about the scalability and energy efficiency of our proposed RC-DFA in "Hardware affinity of training algorithm" and "Computational efficiency" in the discussion section in the revised manuscript.

On the reference to augmented DFA (a-DFA):

You also mentioned about the relevance of our previous work named augmented-DFA (a-DFA) [R4]. This method was proposed for a physical neural network whose nonlinear function *f* is uncertain. We used alternative nonlinear functions *g* instead of the actual gradient of physical nonlinearity *f'*. As described in "Reservoir computing based direct feedback alignment" in the result section of the present manuscript [please see eq. (6)], the RC-DFA can also be augmented for such a case. An experimental demonstration can be found in Supplemental Information S3. As this point might not have been not clear to readers, we added a description in the revised manuscript.

[R4] Nakajima, M. *et al*. "Physical deep learning with biologically inspired training method: gradient-free approach for physical hardware", *Nat. Commun*. 13, 7847 (2022).

Change:

- We added the description for aRC-DFA in "Reservoir computing based direct feedback alignment" in result section.

Thank you again for your thoughtful feedback. We hope these revisions address your concerns and improve the clarity of our work.

Yours sincerely,
Mitsumasa Nakajima

Mitsumasa Nakajima
NTT Device Technology Laboratories
Nippon Telegraph and Telephone Corporation
3-1 Morinosato Wakamiya, Atsugi-shi, Kanagawa 243-0198, Japan
Tel: +81-46-240-2576
mitsumasa.nakajima @ntt.com

November 11th, 2024
Communications Physics
The Campus
4 Crinan Street, London N1 9XW
United Kingdom

Dear editors and reviewers,

We are grateful for your considering the enclosed manuscript “*Reservoir Direct Feedback Alignment: Deep Learning by Physical Dynamics*” by M. Nakajima, Y. Zhang, K. Inoue, Y. Kuniyoshi, T. Hashimoto, and K. Nakajima for publication in Communications Physics. We deeply thank the editors and reviewers for taking the time to check our manuscript and supporting our submission. We are also sincerely grateful for your and reviewers' final comments on our manuscript. We believe that this revision makes this paper eminently suitable for publication. We look forward to hearing from you and thank you for your attention to this submission.

To Reviewer #1

Comment:

I would like to thank the authors for addressing my comments.

I recommend publication of the manuscript.

We thank the referee for the positive feedback on our manuscripts. We sincerely appreciate the time and effort you have taken to review our work. Our manuscript has been highly improved thanks to your thoughtful and constructive comments at the first revision.

To Reviewer #3

Comment:

The authors have addressed and considered in detail all my previous comments. I recommend the manuscript for publication.

.

Thank you very much for your valuable feedback on our manuscript. We thank you for your constructive feedback. We believe that our manuscript has been improved highly thanks to your comments.

Additional changes:

We have slightly revised the manuscript according to the suggestion from editor team.

Changes:

(1) We added some explanations of figures in each figure caption

(2) We added the error range in Fig. 1c
(3) We realized that the term "limit cycle" is not appropriate in Fig. 3 because the term is usually for continuous-time dynamical systems. Thus, we changed the word “limit cycle” to “periodic”. We also added the explanations for the order, periodic, diverged, and chaotic region in the figure caption.
(4) We added the acknowledgement section.

Yours sincerely,
Mitsumasa Nakajima

Mitsumasa Nakajima
NTT Device Technology Laboratories
Nippon Telegraph and Telephone Corporation
3-1 Morinosato Wakamiya, Atsugi-shi, Kanagawa 243-0198, Japan
Tel: +81-46-240-2576
mitsumasa.nakajima @ntt.com