

Supplementary Information for

Scalable physical deep learning using optical dynamics with state-skipping training

Yongbo Zhang^{*1}, Mitsumasa Nakajima², Katsuma Inoue¹, Toshikazu Hashimoto², Yasuo Kuniyoshi^{1,3}, and Kohei Nakajima^{1,3}

¹Graduate School of Information Science and Technology, The University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Tokyo 113-8656, Japan

²NTT Device Technology Laboratories, 3-1 Morinosato-Wakamiya, Atsugi, Kanagawa, Japan

³Next Generation Artificial Intelligence Research Center, The University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Tokyo 113-8656, Japan

S1. Details of model architecture and training

In this section, we systematically list the training configurations of each model on the MNIST and CIFAR-10 datasets in each experiment, including hidden dimensions, optimizer type, learning rate, momentum coefficient, batch size, number of training epochs, label smoothing coefficient, and data augmentation strategies. These settings are presented in Tables S1 and S2.

Dataset	MNIST		
Architecture	most basic	scaled up in size	scaled up in depth
Hidden nodes	500	500 ~ 5000	500
Optimizer	SGD	SGD	SGD
Learning rate	0.005	0.005	0.005
Momentum	0.9	0.9	0.9
Batch size	64	64	64
Epochs	120	300	300
Label smoothing	-	0.1	0.3
Random resized crop and horizontal flip	-	-	-
Mixup	-	-	-

Table S1. Training settings for different experiments on MNIST task

The parameters in each experiment are adjusted based on the model architecture and dataset characteristics to ensure that the models are fully trained and to avoid overfitting. For example, more training epochs and stronger label smoothing are used in the depth-scaled-up experiments, and significantly more training epochs and data augmentation strategies (such as random cropping, horizontal flipping, and mixup) are employed in the CIFAR-10 task to accommodate the more complex image distribution. For the parameters of data augmentation strategies, we utilize the default settings for the random resized crop and horizontal flip. For mixup, we follow [1] to use a strong strength of 0.8. Each experiment in this work was conducted five times to ensure statistical reliability.

^{*}Corresponding author: zhang@isi.imi.i.u-tokyo.ac.jp

Dataset	CIFAR-10			
Architecture	most basic	scaled up in size	scaled up in depth	ResMLP exp.
Hidden nodes	1000	500 ~ 5000	1000	1024
Optimizer	SGD	SGD	SGD	LION
Learning rate	0.01	0.01	0.01	$1e-5$
Momentum	0.0	0.0	0.0	-
Batch size	128	128	128	4096
Epochs	120	500	500	1000
Label smoothing	-	0.1	0.3	0.3
Random resized crop and horizontal flip	-	Yes	Yes	Yes
Mixup	-	-	-	Yes

Table S2. Training settings for different experiments on CIFAR-10 task

S2. Results of training fully-coupled OASP-RC using BPTT and linear regression

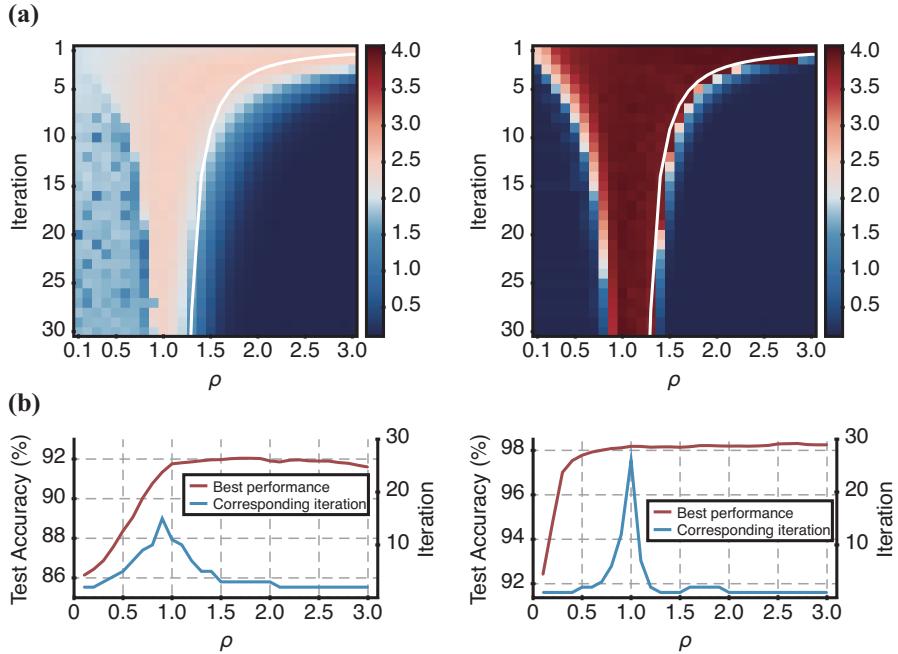


Figure S1. (a) The test accuracy of fully-coupled OASP-RC trained with BPTT as a function of T and ρ . $|\log(\epsilon)|$ is harnessed to represent the test accuracy, with higher values indicating better performance, and ϵ denotes the prediction error. The white line denotes the Lyapunov time T_L , which corresponds to the boundary. (b) The best performance achieved at different ρ , and the number of state evolution iterations T required to achieve this performance.

In this section, we report the performance of a fully-coupled OASP-RC trained with existing training strategies for comparison. We employ linear regression and the mainstream training method—BPTT—to train and evaluate the model on the MNIST task. The experimental design and parameter settings strictly follow the standards outlined in Section 2.2 of the main text; the only controlled parameters are the two key factors governing reservoir layer dynamics—the number of state-evolution iterations T and the expanding coefficient ρ . The results are presented in Figure S1.

As observed, the performance trends obtained with these two training algorithms align with

the results of the SS-DFA in the main text. In particular, when ρ is near 1.0 or exceeds 1.0—in other words at the edge of chaos or when entering the expanding/chaotic regime—the model achieves its optimal performance interval (specifically, for BPTT, optimal performance of 98.31% is achieved at $\rho = 2.7, T = 1$; for linear regression, 92.04% is achieved at $\rho = 1.7, T = 3$). Notably, the white curves in Figure S1a indicate the system’s Lyapunov time T_L . It can be seen that the T_L is highly consistent with the transition of model performance within the expanding/chaotic regime. This indicates that T_L can serve as a reliable indicator by approximating a reasonable upper bound on the state-evolution iterations T under different ρ values. This trend is also consistent with the SS-DFA results.

Figure S1b further summarizes the optimal performance of fully-coupled OASP-RC trained with two strategies at various ρ values, along with the corresponding iteration count T required to achieve them. Compared to the results of SS-DFA, the overall trend remains consistent. In other words, within the contractive region, the number of iterations T required for the model to achieve optimal performance increases continuously as ρ increases. However, upon reaching the edge of chaos and entering the expanding/chaotic regime, the required T decreases continuously with increasing ρ . Moreover, it is noteworthy that relative to BPTT and linear regression, our SS-DFA typically reaches its optimal performance with fewer iterations, highlighting its advantage in training efficiency.

S3. Verification of proposed SS-DFA on the feed-forward echo-state network (FFESN)

In this work, we propose a novel training strategy, SS-DFA. To validate its effectiveness, we first evaluate it on the chaotic feed-forward echo state network (FFESN) introduced by [2]. The structure of FFESN is similar to that of our fully-coupled OASP-RC, with the only difference being that FFESN uses the hidden layer setting of a conventional ESN, that is, the tanh activation function. We conduct experiments on the MNIST task using the same settings as [2], in other words, 500 nodes and a connection density of 0.5 in the hidden layer, and other training parameters and architectures are consistent with the most basic architecture experiments in this work. The performance as a function of ρ and T is shown in Figure S2. The experimental results show that the FFESN trained with SS-DFA exhibits consistent trends in performance with respect to ρ and T compared to the FFESN

trained with BPTT in [2]. Furthermore, under the settings of $\rho = 1.2$ and $T = 1$, the model achieves a test accuracy of 98.07%, which is nearly identical to the accuracy of 98.23% reported in [2]. Additionally, the Lyapunov time T_L can still clearly delineate the boundary between superior and inferior model performance. These results indicate that SS-DFA, which eliminates access to intermediate states and thus uses imprecise gradient transfer, does not alter the model’s dependence on dynamic properties, nor does it lead to performance degradation. This validates the effectiveness of the proposed training strategy and lays the foundation for its application OASP-RC models.

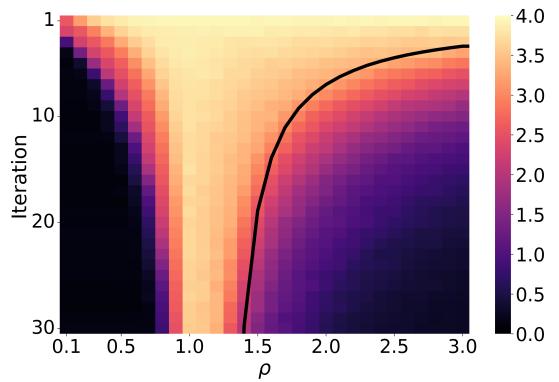


Figure S2. The test accuracy of the FFESN trained with SS-DFA as a function of T and ρ . $|\log(\epsilon)|$ is harnessed to represent the test accuracy, with higher values indicating better performance, and ϵ denotes the prediction error. The black line denotes the Lyapunov time T_L of the system, which corresponds to the boundary.

S4. Number of accumulations of system topology information in the backward process

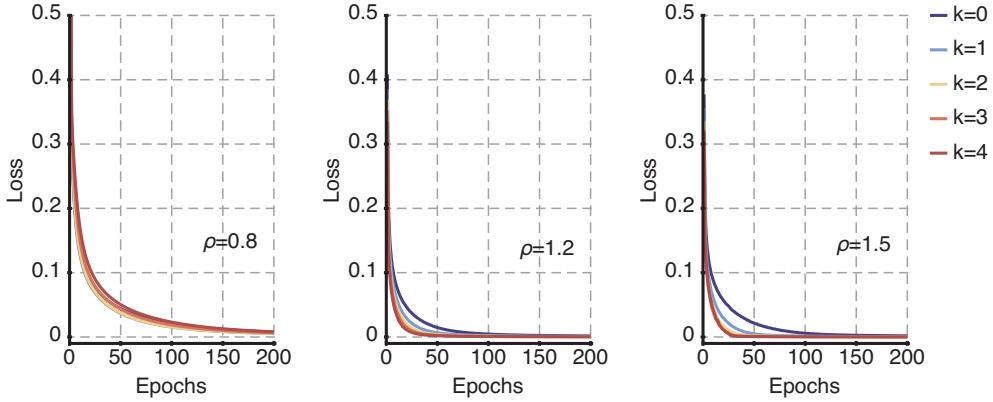


Figure S3. Typical training curves of the fully-coupled OASP-RC trained with SS-DFA on different k settings

In the Section 2.1, we discuss the necessity for the accumulations of system topology information in the reservoir layer when training fully-coupled OASP-RC with SS-DFA, although the accumulation count k need not correspond to the iteration count T . This section further explores the selection of the k parameter in the Equation 3 in the main text. Here, we conduct experiments on the MNIST task, setting the state-evolution iteration T to 5, resulting in a k parameter range from 0 to 4. Figure S3 displays typical training curves for the framework under different k and ρ values. It is evident that regardless of the dynamic state of the system, the value of k does not necessarily strictly equal $T - 1$. For the system in the expanding/chaotic regime (i.e., $\rho = 1.2$ and $\rho = 1.5$), increasing the value of k accelerates convergence speed. Conversely, for the system in the contractive regime (i.e., $\rho = 0.8$), increasing k slows convergence. We attribute this to the accumulation of spectral radius ρ of W acting as a scaling factor that amplifies or attenuates the magnitude of the error signal used for updates during the backward process. We also make a preliminary exploration of the selection of k values for larger state-evolution iteration T , discovering that training fails to work effectively when k is too small. This dependency will be further investigated in future work. It is noteworthy that regardless of the value of k , since W is fixed, the effect of $(W^\top)^k$ is essentially equivalent to a scaling term. Therefore, it can be precomputed offline, incurring no additional computational overhead during training.

S5. Operational principles of the SS-DFA and its version that ignores the scaling factor during the backward process

This section aims to demonstrate the operational principles underlying the basic characteristics and performance trends of the SS-DFA-trained fully-coupled OASP-RC presented in the main text. Using the most basic model (comprising only one input, one reservoir and one output layers), we analyze the largest singular values (LSVs) of the input matrix W^{in} and output matrix W^{out} before and after training on the MNIST dataset. We compare scenarios where ρ is retained versus omitted, in other words, we normalize W^\top , during the backward process. The results are shown in the Figure S4. It can be observed that regardless of ρ 's inclusion in the backward process, the boundaries where the LSVs of W^{in} and W^{out} become excessively large or small align with the trends observed in the test accuracy results. To be specific, where ρ is retained, when $\rho < 1$, as the iteration T increases, the error signal used to update W^{in} significantly attenuates due to the cumulative effect of the contractive scaling factor, leading to the LSVs of W^{in} showing almost no change in magnitude compared to its initial value, which

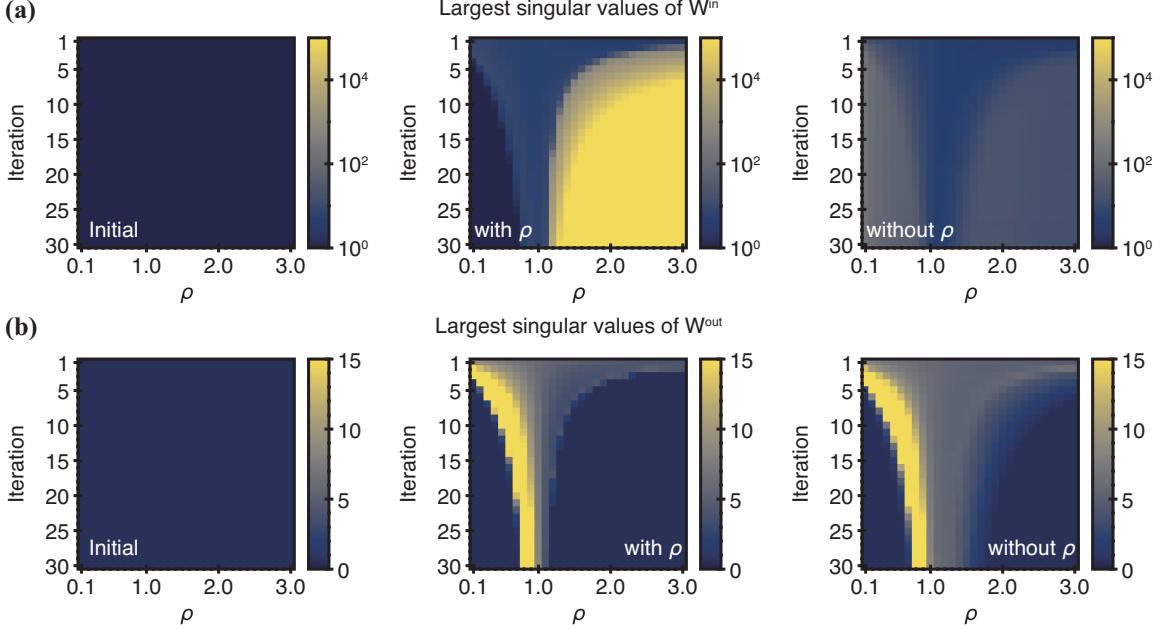


Figure S4. (a) The largest singular values (LSVs) of the input matrix W^{in} as a function of ρ and iteration count before and after SS-DFA training. (b) The LSVs of the readout matrix W^{out} as a function of ρ and iteration count before and after SS-DFA training. Initial: represents the LSVs of initialized W^{in} ; with ρ : refers to preserving the effect of ρ during the backward process, i.e., not normalizing W^{\top} ; without ρ : refers to omitting the ρ during the backward process.

indicates that poor performance stems both from the system's contractive dynamics and from insufficient update magnitude failing to enable effective optimizing for W^{in} . When $\rho > 1$, the update signal exponentially amplifies when T increases, causing W^{in} 's LSVs to expand exponentially (e.g., at $\rho \simeq 3.0$ and $T \simeq 30$, W^{in} 's LSVs reach 10^{16} magnitude). Performance collapse in this region stems not only from the system's inherent excessive chaos but also from overly expanded update signals altering W^{in} 's scale, making it difficult for the input mapping to produce meaningful representations during the forward process. In contrast, within the optimal performance regions (edge of chaos and the transient expanding/chaotic regime), W^{in} 's LSVs are moderately amplified (to approximately the order of 10^1). This indicates that training selectively amplifies input information in meaningful directions and enhances separability of features, while avoiding pushing the system into excessively contractive and chaotic regimes; consequently, superior generalization and accuracy are achieved.

Wherever we omit ρ , we can effectively eliminate the cumulative impact of the scaling factor during the backward process: On one hand, the attenuation of update signals caused by contractive scaling factor when $\rho < 1$ is removed; on the other hand, the exponential amplification of update signals induced by the expanding factor when $\rho > 1$ is also suppressed. Consequently, the cumulative effect of ρ on W^{in} is eliminated. As shown in Figure S4, the regions with excessively small or large LSVs of W^{in} no longer appear. It is crucial to emphasize that the relatively higher LSVs in the lower-left and lower-right regions of the figure, exceeding the LSVs in the optimal range, do not stem from cumulative scaling factors. Instead, they result from the system's inherent contractive or excessive chaotic characteristics, which degrade prediction quality and increase errors, thereby inducing larger magnitude of the update signals. In summary, by ignoring ρ during the backward pass of SS-DFA, we eliminate the additional amplification/attenuation effects introduced by training algorithms. This allows prediction to be driven solely by the system's intrinsic dynamics, yielding a broader effective operating range. Notably, BPTT struggles to avoid cumulative scaling factor effects due to its reliance on precise

information in the error backpropagation, whereas our simplified SS-DFA approach completely eliminates such impacts. We believe that this strategy—minimizing backward dependencies while emphasizing forward dynamics—merits further systematic investigation.

S6. Detailed characteristics analysis of the single-map framework

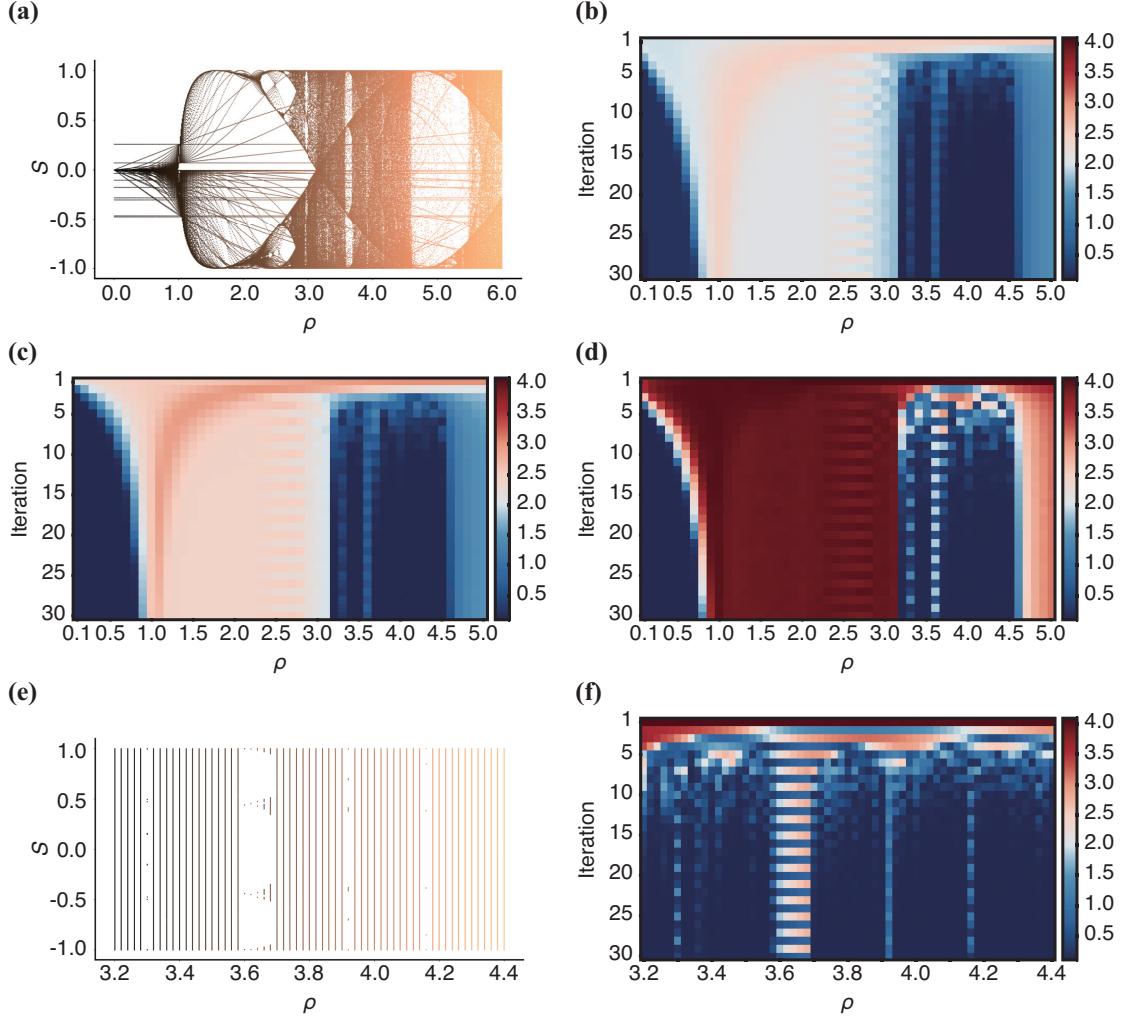


Figure S5. (a) The bifurcation diagram of the proposed single-map architecture. (b) The test accuracy of single-map OASP-RC trained with ridge regression as a function of T and ρ . $|\log(\epsilon)|$ is harnessed to represent the test accuracy, with higher values indicating better performance, and ϵ denotes the prediction error. (c) The results of single-map OASP-RC trained with readout-only learning using stochastic gradient descent (SGD). (d) The results of single-map OASP-RC trained with proposed SS-DFA. (e) The bifurcation diagram at finer ρ resolution in the chaotic region. (f) The performance of single-map OASP-RC trained with SS-DFA at finer ρ resolutions within the chaotic region.

In this section, we supplement the main text by discussing in greater detail the fundamental characteristics of the proposed single-map OASP-RC. Figure S5a and S5e present the bifurcation diagram of the single-map architecture and the higher-resolution bifurcation diagram obtained through fine sampling of ρ in the chaotic region, respectively; Figure S5b, S5c, and S5d show the results of training the single-map OASP-RC using readout-only learning with ridge regression and stochastic gradient descent (SGD), and using SS-DFA, respectively; Figure S5f further details the SS-DFA training results obtained through high-resolution scanning of ρ in the chaotic region.

Across broader scans of ρ , the performance trends produced by these training strategies closely align with the bifurcation diagram's general structure. Furthermore, SS-DFA exhibits consistent overall trends with the other two baseline strategies: Optimal performance regions are achieved at the bifurcation point corresponding to the transition from single fixed point to bistable states ($\rho \sim 1$) and in the transient expanding/chaotic region, with SS-DFA demonstrating considerably superior performance. This indicates that applying SS-DFA to single-map OASP-RC leaves the system's intrinsic properties unchanged while enabling more effective training, meaning joint optimization of input matrix W^{in} offers advantages over training W^{out} alone.

Moreover, as shown in Figure S5e, multiple non-chaotic windows emerge within the chaotic region as the ρ sampling resolution increases. Correspondingly, Figure S5f reveals effective performance within these windows. This demonstrates strong consistency between bifurcation analysis and framework performance, providing a theoretical basis and guidance for framework design and parameter selection. It should be noted that when T is small, model performance no longer strictly follows the bifurcation diagram's indications. As discussed in the main text, transient dynamics dominate at this scale, making bifurcation diagrams that reflect global property ineffective for interpretation. Analysis should instead leverage the finite-time maximum Lyapunov exponent (FTMLE), which reflects local time-domain characteristics. Relevant results and discussions are presented later in the supplementary material.

S7. Analysis of finite-time maximum Lyapunov exponents (FTMLE) for the single-map framework

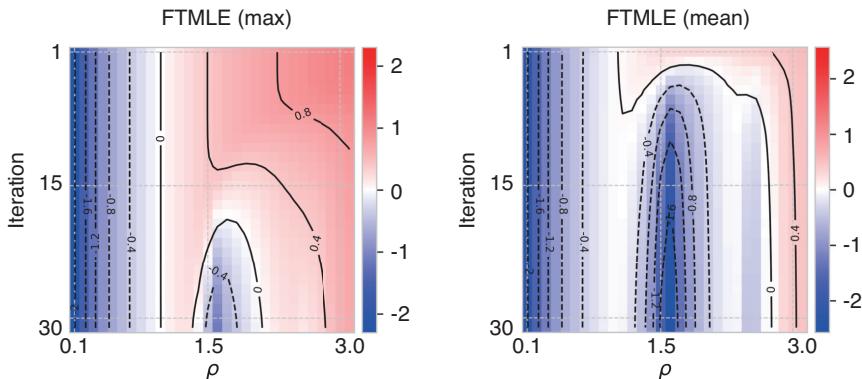


Figure S6. The dependence of the finite-time maximum Lyapunov exponent (FTMLE) of the single-map framework on ρ and the number of iterations T . The left panel shows the system's maximum FTMLE, while the right panel shows the system's mean FTMLE.

In this section, we employ finite-time maximum Lyapunov exponent (FTMLE) to investigate the transient dynamics of the single-map framework. As stated in the main text, bifurcation analysis reflects the long-term global properties of the system; consequently, it cannot accurately represent the system's behavior for relatively small iteration settings. To characterize short-term behavior, we utilize the FTMLE, which quantitatively captures local instability over finite time. Similar to the maximum Lyapunov exponent (MLE)[3], the metric that captures global features, an FTMLE greater than 0 indicates that the system is sensitive to initial conditions, demonstrating expanding and chaotic behavior or transient instability. When the FTMLE is less than 0, the system tends toward an attractor or stable periodic orbit, indicating that it has good local stability and is controllable on a finite timescale.

Since there are no connections between nodes/maps within the single-map framework, focusing solely on the system's maximum FTMLE is insufficient in high-dimensional scenarios (e.g., for a dimension of 1000, there are 1000 independent nodes/maps that do not interact during

state iteration). It is therefore necessary to explore the FTMLE of each node/map individually to summarize the aggregate behavior. Accordingly, to demonstrate the overall dynamics of the single-map layer, we report the system’s mean FTMLE as a representative metric. The results for maximum and mean FTMLE with respect to ρ and iteration T are shown in the Figure S6. It can be observed that when $\rho > 1$ and T is small, the mean FTMLE remains positive, indicating that the system as a whole resides in a transient expanding/chaotic regime. This aligns with the performance trends observed in the single-map framework, further confirming that the great performance indeed stems from the transient expanding/chaotic property. Additionally, when T is large, the mean FTMLE results show high consistency with bifurcation analysis, thereby validating the effectiveness of these analyses. It should be noted that in certain regions, the maximum FTMLE is positive while the mean FTMLE is negative (e.g., the double fixed point region). This indicates that some nodes/maps within the single-map layer exhibit the expanding/chaotic property, yet the overall average remains in contractive dynamics.

S8. Results on the more complex dataset CIFAR-10

After outlining the fundamental characteristics of the single-map framework and its overall performance trends under SS-DFA training, this section further evaluates the framework’s performance on the more challenging CIFAR-10 dataset. We still employ the most basic and general three-layer architecture: an input layer, a single reservoir layer, and an output layer. The results are shown in Figure S7: On more challenging tasks, the framework remains stable and effective, with its behavior as a function of T and ρ consistent with the previous MNIST experiments, that is, the optimal regions are concentrated in the bifurcation point corresponding to the transition from single fixed point to bistable states and the transient expanding/chaotic regime. Specifically,

the model achieves optimal performance at $\rho = 0.9$ and $T = 13$, with a test accuracy of 54.85%. This result outperforms the BP-trained MLP with the same architecture, clearly demonstrating the comprehensive performance advantages of our method.

S9. Gradient-free training: SS-aDFA

This section provides a systematic analysis of the SS-aDFA, a gradient-free extension of the proposed SS-DFA. Based on the experimental design and parameter settings of Section 2.2 in the main text, we conduct experiments for both fully-coupled and single-map OASP-RCs on the MNIST task to verify the effectiveness of the SS-aDFA. We employ a phase-shift cosine function as the surrogate nonlinearity g to replace the precise derivative function f' during the backward process; the updated rule of fully-coupled OASP-RC is then derived as:

$$e^{(l)} = B^{(l)} e^{(L)} \cdot (W^\top)^k \odot \cos(\omega z_0 + \theta) W^\top \quad (1)$$

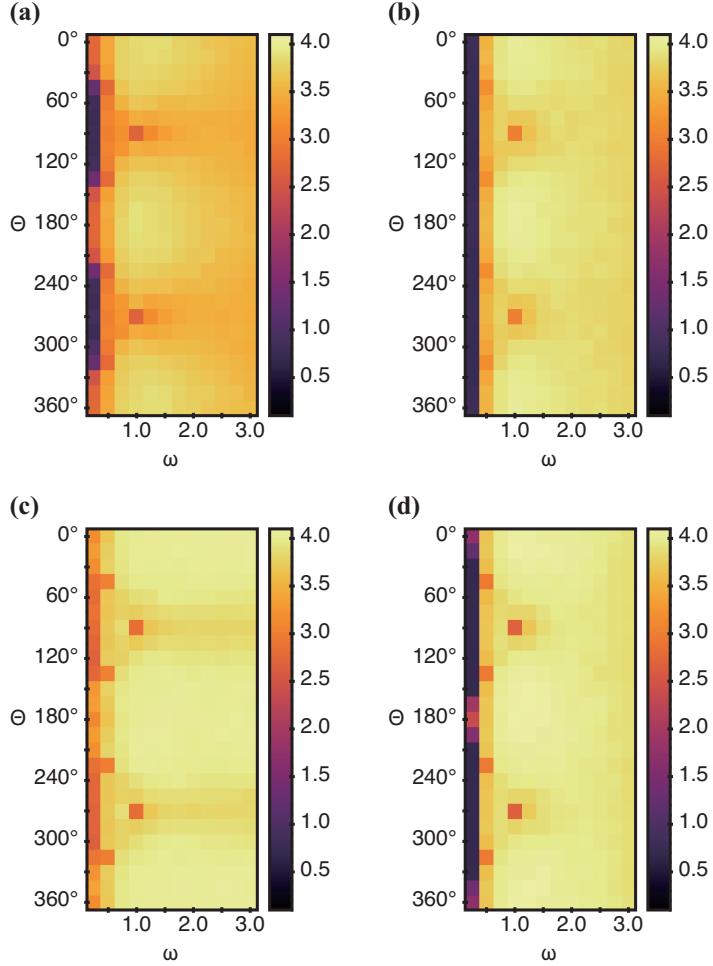


Figure S8. (a), (b) The test accuracy of the fully-coupled OASP-RC trained with SS-aDFA as a function of angular frequency ω and phase θ under different ρ and T conditions. (a: general setting, $\rho = 1.2$, $T = 5$; b: optimal parameter combination identified in basic characteristic experiments, $\rho = 2.8$, $T = 1$.) $|\log(\epsilon)|$ is harnessed to represent the test accuracy, with higher values indicating better performance, and ϵ denotes the prediction error. (c), (d) The results of the single-map OASP-RC trained with SS-aDFA. (c: general setting, $\rho = 1.2$, $T = 5$; d: optimal parameter combination identified in basic characteristic experiments, $\rho = 2.7$, $T = 1$.)

and for the single-map case, the update rule is:

$$e^{(l)} = B^{(l)} e^{(L)} \odot \cos(\omega z_0 + \theta) \quad (2)$$

where the angular frequency ω ranges from 0.25 to 3.0, and the phase θ changes from 0° to 360° .

Figure S8a and S8b show the results of the test accuracy of the fully-coupled OASP-RC as a function of ω and θ under different ρ and T conditions. Figure S8c and S8d report the results of single-map OASP-RC. We evaluate SS-aDFA under two sets of ρ and T parameters: (i) a general configuration for comparison, $\rho = 1.2, T = 5$ (see Fig. S8a and S8c). (ii) the optimal settings determined based on the SS-DFA experiments in the main text, fully-coupled with $\rho = 2.8, T = 1$ (Figure S8b) and single-map with $\rho = 2.7, T = 1$ (Figure S8d). It can be seen that the extension SS-aDFA can work effectively for both architectures under almost all ω, θ settings, which proves its effectiveness. Specifically, under the general ρ and T setting, the fully-coupled OASP-RC achieves 97.98% test accuracy at $\omega = 1.0, \theta = 180^\circ$, while the single-map OASP-RC achieves 98.27% at $\omega = 2.5, \theta = 345^\circ$; both slightly exceed their SS-DFA trained baselines (97.86% and 98.18%, respectively). Under the optimal settings, the fully-coupled OASP-RC achieves 98.29%

at $\omega = 1.0, \theta = 165^\circ$, and the single-map OASP-RC reaches 98.36% at $\omega = 1.25, \theta = 165^\circ$, again marginally surpassing the SS-DFA-trained counterparts (98.26% and 98.33%, respectively). The above results not only indicate that the proposed SS-DFA can be effectively extended to a gradient-free version but also demonstrate that performance can be further improved by optimizing the parameters in the alternative nonlinear function. In this work, we primarily employ grid search to measure the selection of nonlinear parameters. Exploring more efficient methods to optimize parameters in the alternative nonlinear function will be one of our future research interests.

S10. Optimization of the single-map framework

This section discusses our preliminary exploration of optimization research on single-map frameworks. The main configuration in this paper employs a single mapping within the reservoir layer, that is, all nodes/maps share the same ρ for state updates, and grid search is used to identify optimal parameter combinations in the (ρ, T) space. In fact, the single-map framework need not be confined to a single identical mapping; we can introduce multiple mappings to enrich intra-layer dynamics, or even perform adaptive updates of ρ to optimize performance. In this experiment, we first construct a multi-map framework: On a mapping layer of dimension 1000, we uniformly scan ρ from 0.1 to 5.1 with a step size of 0.005, assigning distinct ρ values to each node/map. We use the most basic architecture with one input layer, one reservoir layer and one output layer, and evaluate its performance on MNIST and Fashion-MNIST [4] datasets. Subsequently, we optimize ρ using SS-DFA, with results summarized in Table S3. It can be observed that the multiple-map framework continues to function effectively, achieving performance superior to the baseline MLP. Optimizing ρ using SS-DFA yields a marginal performance improvement. It should be noted that these results are actually inferior to the single-map performance reported in the main text. We attribute this discrepancy to the inclusion of extremely small ρ settings, such as 0.1, within the multiple-map framework used in the experiments. Furthermore, we observe that the initial setting of ρ is crucial for its optimization, as it determines whether the trained ρ converges toward bifurcation points or transient expanding/chaotic regions. The uniform distribution ranging from 0.1 to 5.1 used in our experiments may not constitute an optimal initial value, potentially explaining why the optimized performance dose not significantly improve. Optimizing ρ to enhance performance and extending this mechanism to deep architectures are possible future research directions.

	MNIST	Fashion-MNIST
MLP	98.47%	89.28%
Fixed	98.63%	90.02%
Tuned	98.65%	90.04%

Table S3. Test accuracy of the fixed and optimized multiple-map frameworks under $T = 1$ and baseline MLP with identical architecture on the MNIST and Fashion-MNIST datasets

S11. Noise robustness

In this section, we demonstrate the noise robustness of our frameworks, providing a theoretical foundation for practical implementation. In physical implementation, physical signals (such as light intensity, electric fields, etc.) are often affected by noise within physical devices. To simulate this process, we add noise to each iteration of the reservoir layer as shown below:

$$s(t+1) = f(z_t) + \varepsilon \sigma(t) \quad (3)$$

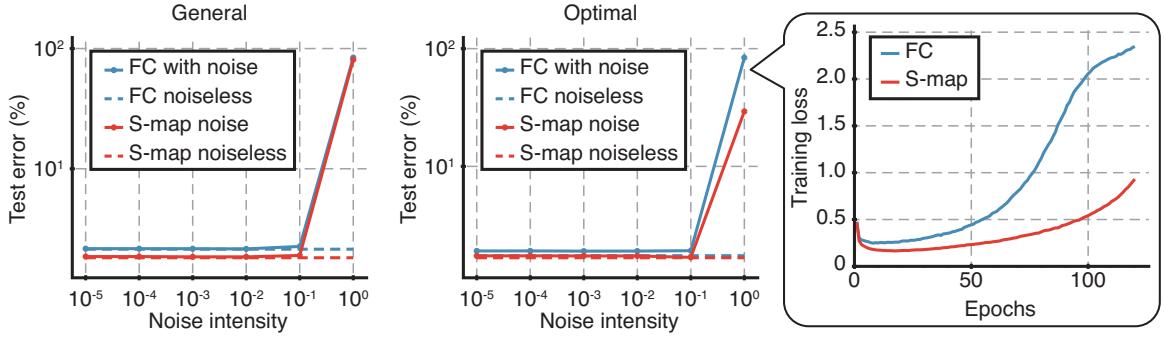


Figure S9. The test error of models trained with SS-DFA as a function of noise intensity under different (ρ, T) settings. The solid lines indicate performance with noise injection, and the dashed line indicates noiseless baseline performances. The rightmost panel displays training curves of models under optimal ρ and T settings when the noise intensity is equal to 10^0 .

where $z_t = Ws(t)$ for the fully-coupled OASP-RC, and $z_t = \rho s(t)$ for the single-map OASP-RC. $\sigma(t)$ denotes randomly generated white Gaussian noise at time t , and ε is the noise intensity. In this numerical simulation, we adopt the same architecture and experimental settings as those used in Section 2.2 in the main text on MNIST dataset. The value of ε varies from 10^{-5} to 10^0 . The test error results of the model at the final training epoch as a function of ε are shown in Figure S9. We still evaluate the noise robustness under two sets of ρ and T parameters: (i) a general configuration for comparison, $\rho = 1.2, T = 5$ (see the left panel in Figure S9) and (ii) the optimal settings determined based on the SS-DFA experiments in the main text, fully-coupled with $\rho = 2.8, T = 1$ and single-map with $\rho = 2.7, T = 1$ (see the middle panel in Figure S9). It can be seen that our frameworks perform very stably and is almost consistent with the noiseless case when the noise intensity ε is less than 10^{-1} . This indicates that our framework has very good noise robustness. The signal-to-noise ratio (SNR) of commonly used optoelectronic platforms is approximately 26 dB [5, 6], which corresponds to $\varepsilon = 4 \times 10^{-2}$, which is much lower than 10^0 . Therefore, our paradigm can theoretically meet the noise robustness requirements of most practical implementations. Notably, the rightmost panel shows the training curve of models under optimal settings when $\varepsilon = 10^0$. It can be observed that despite large noise intensity, the model initially makes learning progress. However, with continuous noise injection, training ultimately fails. This indicates that under carefully designed control of the training process and duration, an effective working window persists even under high-noise conditions. The single-map OASP-RC exhibits lower loss, demonstrating superior noise robustness compared to the fully-coupled OASP-RC.

S12. Physical implementation of readout-only training for the single-map framework

This section discusses preliminary experimental results from the physical implementation. To validate the effectiveness of the single-map OASP-RC constructed by the optoelectronic reservoir computer platform, we first employ the stochastic gradient descent (SGD)-based readout-only training strategy to examine its fundamental characteristics. To simplify the workflow and accelerate experiments, we utilize a reduced MNIST dataset (12,000 training images, 4,000 test images). The training set-

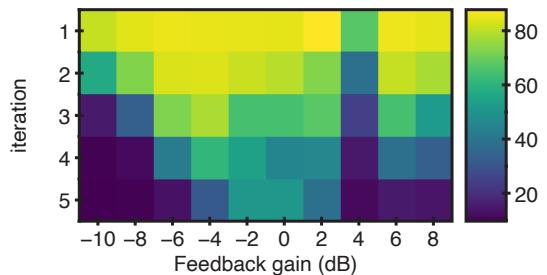


Figure S10. The readout-only training results of the physically implemented single-map framework as a function of T and ρ on MNIST dataset

tings are: an SGD optimizer with a learning rate of 5×10^{-3} , a reservoir layer with 1600 nodes, and training for 50 epochs. In the physical implementation, we adjust the system's feedback gain α by using a variable optical attenuator to control ρ of the single-map OASP-RC, thereby altering the system's dynamic properties. The correspondence between feedback gain α and ρ is given by the following equation:

$$\rho = \kappa 10^{(\alpha - \alpha_0)/20} \quad (4)$$

where κ is constant and α_0 denotes the reference baseline, set to 1 and 0, respectively. Here, $\alpha_0 = 0$ means that when $\alpha = 0$, ρ equals to 1.

The test accuracy as a function of feedback gain α and iteration count is shown in Figure S10. Although these results do not cover the entire range of ρ and iteration values in the main text experiments, the observed overall trend aligns with numerical simulations: When α is negative (corresponding to $\rho < 1$), performance gradually deteriorates as iterations increase; when α is positive, corresponding to $\rho > 1$, performance improves within the transient expanding/chaotic region—i.e., small iteration settings. When iterations and α become relatively large, the system enters the excessively chaotic region, causing performance to decline again. This consistency demonstrates that we have successfully reproduced the key dynamical behavior of the single-map OASP-RC on the optoelectronic platform, providing a crucial foundation for subsequent physical implementation of single-map OASP-RC trained with SS-DFA.

S13. Extending to other deep architectures

In this section, we investigate the scalability of the proposed model to other deep architectures, specifically including the standard block MLP architecture without residual connections and an architecture combined with convolutional neural networks (CNNs). We adopt the same experimental design and parameter settings from Section 2.5 in the main text. Figure S11a and S11c show the schematics of the utilized standard block architecture and the CNN-combined architecture, respectively. In the standard block architecture, we embed the single-map module into the original hidden layer and activation function positions. In the CNN-combined architecture, we replace the classifier part of VGG11 [7] with the single-map module. The experimental results for the two architectures are shown in Figure S11b and S11d, respectively. It can be seen that in the standard block architecture, the model with an embedded single-map module achieves significant improvements in test accuracy and convergence speed. In the CNN-combined architecture, we observe that the model performance is superior to that of the conventional classifier under all ρ settings, and when $\rho > 1$ (i.e., in the transient expanding region), the model performs better and converges slightly faster. Specifically, in the standard block architecture, the test accuracy improved from 63.74% for the conventional block to 66.29% for the block with the embedded single-map module; in the CNN-combined architecture, replacing the classifier improved the accuracy from 87.88% to 88.52%. These results demonstrate the scalability and effectiveness of the proposed approach.

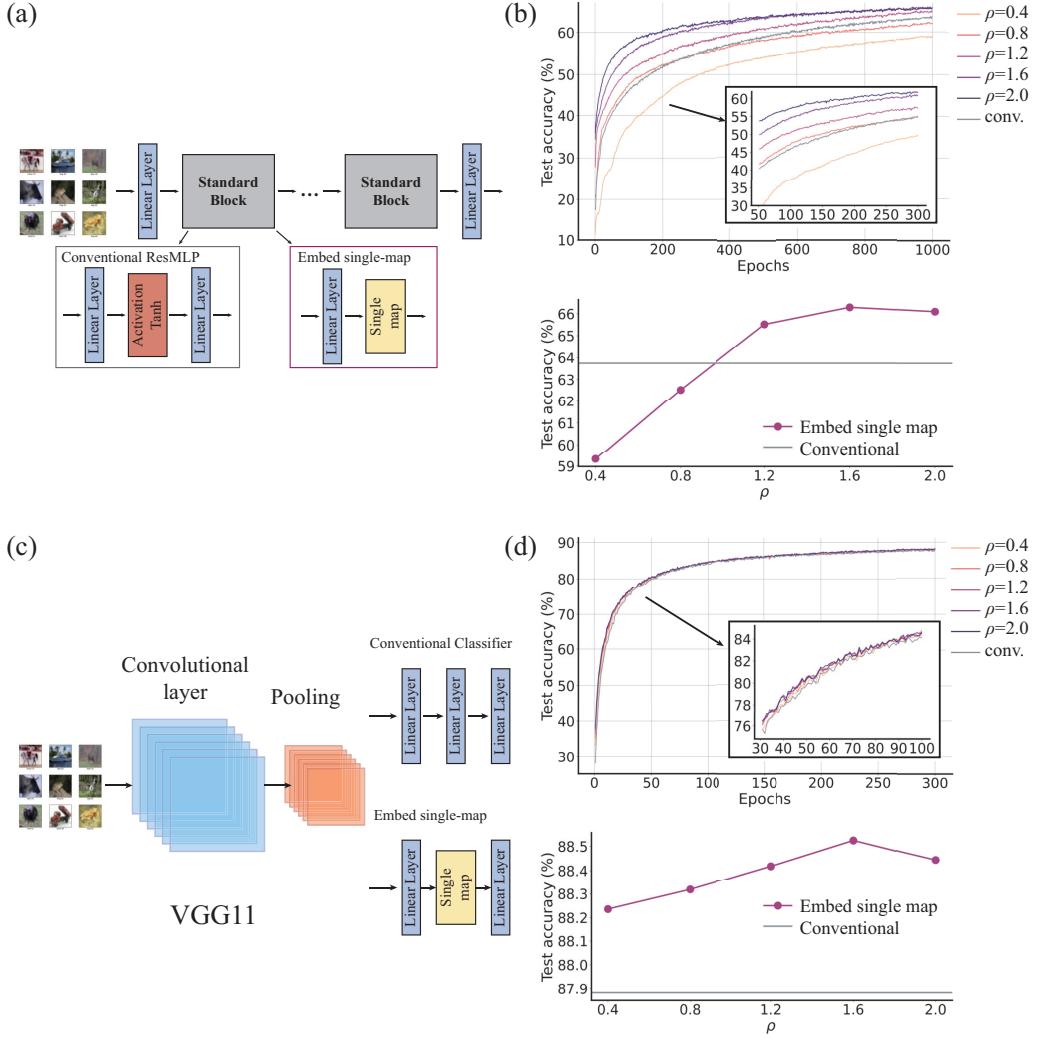


Figure S11. (a) Schematic of standard block architecture. (b) Performances of conventional standard block MLP and the proposed standard block structure with embedded single-map module. (c) Schematic of VGG11 with conventional and proposed embedded single-map module classifiers. (d) Performances of conventional VGG11 and VGG11 combined with the embedded single-map classifier.

References

- [1] Gregor Bachmann, Sotiris Anagnostidis, and Thomas Hofmann. Scaling mlps: A tale of inductive bias. *Advances in Neural Information Processing Systems*, 36:60821–60840, 2023.
- [2] Shuhong Liu, Nozomi Akashi, Qingyao Huang, Yasuo Kuniyoshi, and Kohei Nakajima. Exploiting chaotic dynamics as deep neural networks. *Physical Review Research*, 7(3):033031, 2025.
- [3] Angelo Vulpiani. *Chaos: from simple models to complex systems*, volume 17. World Scientific, 2010.
- [4] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.
- [5] Mitsumasa Nakajima, Katsuma Inoue, Kenji Tanaka, Yasuo Kuniyoshi, Toshikazu Hashimoto, and Kohei Nakajima. Physical deep learning with biologically inspired training method: gradient-free approach for physical hardware. *Nature communications*, 13(1):7847, 2022.

- [6] Mitsumasa Nakajima, Yongbo Zhang, Katsuma Inoue, Yasuo Kuniyoshi, Toshikazu Hashimoto, and Kohei Nakajima. Reservoir direct feedback alignment: deep learning by physical dynamics. *Communications Physics*, 7(1):411, 2024.
- [7] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.