

Bayesian Transferability Assessment for Spiking Neural Networks

Anonymous authors

Paper under double-blind review

Abstract

Brain-inspired spiking neural networks (SNNs) attract broad interest in neuromorphic computing but suffer the problem of being difficult to optimize. Concurrently, pre-trained models (PTMs) have become a foundation for developing and applying artificial intelligence. Therefore, it is expected that pre-trained SNNs can alleviate the optimization difficulty of training from scratch. However, with a lot of PTMs available in the model hubs, effectively selecting the most appropriate PTM for a given task remains a significant challenge, often necessitating exhaustive fine-tuning and grid-searching. While several solutions to this challenge have been proposed for the mainstream artificial neural network (ANNs), aimed at developing efficient methods to assess the transferability of PTMs on target tasks, the realm of SNNs remains unexplored. The currently most used transferability assessment method for ANNs predicts transferability in a Bayesian perspective. Feature maps extracted by the PTM backbone on the target task are used to calculate the maximum model evidence as the indicator of transferability. However, ANNs and SNNs differ in architecture, rendering the existing Bayesian method incompatible with SNNs. To solve this problem, this paper introduces a novel approach to using the feature maps averaged over the time domain to calculate maximum evidence. Our proposed **Maximum Evidence method with Averaged Features (MEAF)** demonstrates effectiveness for SNNs. Additionally, the current algorithm calculates maximum evidence in an iterative way. To accelerate the selection of PTMs, an approximation method is proposed to avoid iteration in the calculation of maximum evidence, significantly reducing time consumption. It is shown through experiment that the proposed MEAF method is effective for the transferability assessment of SNNs and surpasses other information theory-based assessment methods such as LEEP and NCE, which can directly adapt to SNNs, underscoring its potential to streamline PTM selection and application in the realm of SNNs.

1 Introduction

Spiking neural networks (SNNs), known as the third generation of neural networks (Maass, 1997), have obtained broad attention in the field of neuromorphic computing (Roy et al., 2019). Inspired by the way biological neurons communicate with spikes in the brain, neurons in SNNs encode information into sequential binary spikes, offering a potential way towards energy-efficient machine intelligence (Dampfhofer et al., 2022). Based on its distinctive mechanism, SNNs have achieved outstanding performance in robotics (Bing et al., 2019; Debat et al., 2021), event-based vision (Zhu et al., 2022; Hagenaaars et al., 2021), and other fields (Pei et al., 2019). However, SNN models are much harder to optimize compared to ANNs (Wu et al., 2019; Wang et al., 2020) because of the non-differentiable nature of spiking signals.

Pre-trained models (PTMs) have fueled the development and applications of artificial intelligence (Jiang et al., 2022). In computer vision, neural networks pre-trained on ImageNet (Deng et al., 2009) exhibit improved performance (Yosinski et al., 2014; He et al., 2019; Kornblith et al., 2019), and recent advances in vision-language pre-training provide large-scale visual encoders capable of even zero-shot generalization on downstream tasks (Radford et al., 2021; Jia et al., 2021). Meanwhile, language models self-supervised pre-trained on a large-scale unlabeled corpus demonstrate in-context learning capabilities (Devlin et al.,

2018; Brown et al., 2020; Radford et al., 2019; Du et al., 2022). The pre-training and adaptation paradigm has become a foundational methodology for the implementation of large neural network models. With lots of PTMs available, model hubs such as TorchVision (maintainers & contributors, 2016) and Hugging Face Transformers (Wolf et al., 2020) have attracted widespread interest. It is expected that pre-trained SNNs can alleviate the difficulty of optimization when training SNNs from scratch.

However, when using PTMs in a model hub, selecting the most suitable PTM for a given task without exhaustive fine-tuning and grid-searching (You et al., 2021) remains a challenge. Several model hub scheduling methods have been proposed to address this challenge for the mainstream artificial neural network (ANN) (Tran et al., 2019; Nguyen et al., 2020; You et al., 2021). They focus on designing efficient methods to assess the transferability of PTMs on a target dataset and then select the highest-ranked candidate. To the best of our knowledge, however, there is no research focusing on the scheduling of SNN PTMs. Therefore, to bridge the gap, it is meaningful to explore model hub scheduling methods for SNN PTMs.

One problem is that not all existing transferability assessment methods for ANNs can be applied to SNNs directly. Current methods use the pseudo-output (Tran et al., 2019; Nguyen et al., 2020) or the feature map (You et al., 2021; 2022) generated by the PTMs to assess the transferability on the target task. The pseudo-output approaches, which are information theory-based, including LEEP (Nguyen et al., 2020) and NCE (Tran et al., 2019), can be applied directly to SNN. However, the feature-based approach logME (You et al., 2021), which is more effective on ANNs, cannot be adopted to SNN directly. In this logME approach, feature maps extracted by the PTM backbone are used to calculate the maximum model evidence as an indicator of transferability from a Bayesian perspective. Differences between ANNs and SNNs in architecture render the feature-based methods incompatible with SNNs from two aspects. i) The feature map extracted by the SNN backbone has an additional temporal dimension, which cannot adapt to the current model evidence calculation framework. ii) SNNs generate predictions based on the feature with a nonlinear transform, while the maximum model evidence method for ANNs supposes a linear hypothesis.

Taking these two factors into consideration, this paper introduces a Bayesian method of PTM transferability for SNNs inspired by the counterpart for ANNs. To match the feature map of SNN with the current calculation framework of maximum model evidence, we propose to use feature maps averaged through the time domain for the assessment of transferability. Theoretical analysis shows that the nonlinear transform of the SNN layer can be estimated by a linear transform with the proposed averaged features, thus satisfying the linear assumption of the maximum model evidence method. Moreover, the calculation of maximum model evidence requires MacKay’s algorithm (MacKay, 1999) to maximize the model evidence iteratively. To further accelerate the selection of PTMs, we propose an approximated maximum model evidence method to avoid iteration while giving similar results as logME. Experiment results show that the proposed **Maximum Evidence** method with **Averaged Features** (MEAF) is effective for SNNs and outperforms information theory-based NCE and LEEP. The effectiveness of the approximated maximum model evidence method is also validated. Our contributions are as follows:

1. We identify the gaps of current Bayesian transferability assessment methods between ANNs and SNNs, and propose the **Maximum Evidence** method with **Averaged Features** (MEAF) for the assessment of transferability of pre-trained SNNs.
2. To further accelerate the selection of PTMs, we provide an approximated maximum model evidence method, which avoids iteration while giving effective results.
3. We validate our proposed method through experiments, and confirm that our Bayesian-based method outperforms information theory-based methods like NCE and LEEP for SNNs.

2 Related works

2.1 Bayesian model selection

Bayesian model selection methods use model evidence (or marginal likelihood) among different hypotheses as the criterion for selection (Wasserman, 2000). It is believed that model evidence implicitly incorporates

the Occam Razor Principle, which balances the complexity and the representation capability of a given model (Knuth et al., 2015). Gull uses model evidence in image reconstruction to determine a critical prior hyperparameter (Gull, 1989). Inspired by Gull’s work, model evidence is used in the field of machine learning by MacKay to determine the degree of polynomials in interpolation problems (MacKay, 1992a). It is also used to get the optimal regularization coefficient in the training of regression (MacKay, 1992c) and classification (MacKay, 1992b) problems, even in multi-layer backpropagation neural network models, for better generalization performance (Wolpert, 1992). An algorithm to automatically determine the regularization coefficient in linear regression problems is proposed by MacKay by maximizing the model evidence iteratively (MacKay, 1999).

Evidence-based Bayesian method provides a practical framework to automatically tune the hyperparameters of deep neural networks without the use of validation data. However, a main obstacle lies in the difficulty of estimating the model evidence when the model is deep, as conventional methods proposed by MacKay are only suitable for basic linear regression problems. The Hessian matrix of loss with respect to the trainable parameters is required to calculate the model evidence, which is impossible to efficiently calculate for large-scale neural networks. To solve this problem, scalable Laplace approximation is used to calculate model evidence for deep neural networks (Daxberger et al., 2021; Immer et al., 2021). However, this process is still time-consuming and suffers the problem of inaccurate estimation (Kunstner et al., 2019).

2.2 Transferability assessment

Assessment of the transferability of PTMs starts from measuring the similarity between the source task where the models are pre-trained and the target task to transfer. Several distance-based metrics are proposed to measure the difference between source and target tasks, which gives error bounds of the transferred model in theory (Ben-David & Schuller, 2003; Ben-David et al., 2006; Mansour et al., 2009). While it is obvious that models trained on more challenging datasets have better performance when transferred to simple tasks, the opposite is not true. Therefore, distance, which is a symmetry metric, is not an ideal measure of transferability (Nguyen et al., 2020).

As an improvement, some new metrics other than distance have been proposed. Tran et al. (2019) uses conditional entropy (NCE) to estimate the test set accuracy after transfer by calculating the conditional entropy of the target task relative to the source task. Conditional entropy is calculated based on the probability distributions to generate source and target tasks. It is asymmetrical and can reflect the difficult relationship between tasks, achieving good results in experiments.

Contrary to the above methods, Nguyen et al. (2020) no longer uses the relation between tasks as the basis for evaluation, but instead directly inputs the target tasks into the pre-trained model, and quantitatively evaluates the transferability through the output of the pre-trained model. They calculate the expected probability of the real result obtained from the pre-trained model based on the joint distribution between the results obtained from the pre-trained model and the real result of the target task, as a basis for evaluation.

Model evidence-based Bayesian methods are also used in transferability assessment. Kim et al. (2016) uses model evidence to select pre-trained convolutional neural network (CNN) models for best feature extraction. You et al. (2021; 2022) use model evidence to predict the transferability of different PTMs on a given target dataset. In this approach, features extracted by the PTMs instead of predicted results are used as the basis for transferability evaluation. Compared to prediction results, the features extracted from the middle layer are more universal, so this method can be applied not only to classification problems but also to other problems such as regression.

3 Methodology

3.1 Background

We consider a supervised learning task on observed data $\mathcal{D} = \{(x_i, t_i)\}_{i=1}^N$ with inputs $x_i \in \mathbb{R}^N$ and targets $t_i \in \mathbb{R}$. SNN model \mathcal{M}_θ predicts y_i for each input x_i , where $\theta \in \mathbb{R}^D$ stands for the trainable weights of the model. The model evidence of \mathcal{M} on task \mathcal{D} is defined as the probability the model generates the observed

data, which can be calculated by

$$p(\mathcal{D}|\mathcal{M}) = \int_{\mathbb{R}^D} p(\mathcal{D}, \theta|\mathcal{M}) d\theta. \quad (1)$$

For different candidate models, the best one can be selected by maximizing the model evidence in the Bayesian perspective.

Model evidence is used to automatically determine the regularization coefficients in ridge regression tasks. Given representations $f_i, i = 1, \dots, N$, the model predicts $\hat{t}_i, i = 1, \dots, N$ through a linear transform. It is known that L-2 regularization is equivalent to subjecting parameters to a Gaussian prior. We denote the Gaussian prior distribution of the parameters as $\theta \sim \mathcal{N}(0, \alpha^{-1}\mathbf{I})$ and the Gaussian prior distribution of observed data as $t_i \sim \mathcal{N}(\theta^T f_i, \beta^{-1})$.

For determination of regularization coefficients, linear regression models with different regularization coefficients α, β are viewed as different candidate models, and the evidence is calculated as a function of α, β like

$$\begin{aligned} Evidence(\alpha, \beta) &= p(\mathcal{D}|\alpha, \beta) \\ &= \int_{\mathbb{R}^D} p(\mathcal{D}|\theta, \beta) p(\theta|\alpha) d\theta \\ &= \left(\frac{\beta}{2\pi}\right)^{\frac{N}{2}} \alpha^{\frac{D}{2}} \exp\left(-\frac{\beta}{2} \|\mathbf{F}m - \mathbf{t}\|^2 - \frac{\alpha}{2} m^T m\right) (\det \mathbf{A})^{-\frac{1}{2}}, \end{aligned} \quad (2)$$

where $\mathbf{F} \in \mathbb{R}^{N \times D}$ denotes inputs and $\mathbf{t} \in \mathbb{R}^N$ denotes targets, $\mathbf{A} = \alpha\mathbf{I} + \beta\mathbf{F}^T\mathbf{F}$, $m = \beta\mathbf{A}^{-1}\mathbf{F}^T\mathbf{t}$.

For convenience of calculation, the logarithm of equation 2 is more often used as

$$\mathcal{L}(\alpha, \beta) = \frac{N}{2} \log \beta - \frac{N}{2} \log 2\pi + \frac{D}{2} \log \alpha - \frac{\beta}{2} \|\mathbf{F}m - \mathbf{t}\|^2 - \frac{\alpha}{2} m^T m - \frac{1}{2} \log \det \mathbf{A}. \quad (3)$$

We donate $\gamma = \sum_{i=1}^D \frac{\beta\sigma_i}{\alpha + \beta\sigma_i}$, where σ_i is the i_{th} singular value of $\mathbf{F}^T\mathbf{F}$.

By iteratively updating α, β with MacKay's algorithm

$$\alpha \leftarrow \frac{\gamma}{m^T m}, \beta \leftarrow \frac{N - \gamma}{\|\mathbf{F}m - \mathbf{t}\|^2}, \quad (4)$$

model evidence can be maximized with α, β optimized simultaneously.

Linear probe is usually used as a method of transfer learning, especially when the pre-trained backbone is large (Radford et al., 2021) as the retraining of large models is sometimes prohibitive. In such a setting, pre-trained backbones are used to extract representations, and a fully connected layer is used to predict results based on the representations. This process is a linear transform, and the model evidence can be calculated analytically.

Therefore, to assess and compare the transferability of different models on target task \mathcal{D} , for each pre-trained backbone \mathcal{M}_i , evidence $\mathcal{L}_i(\alpha, \beta)$ can be calculated and maximized with MacKays's algorithm, resulting in the logarithm of maximum evidence $\mathcal{L}_i(\alpha^*, \beta^*)$. Eventually, the transferability can be compared in terms of $\mathcal{L}(\alpha^*, \beta^*)$, and the candidate models can be ranked.

3.2 Adapting maximum evidence method to SNN with averaged features

Architectural differences between SNNs and ANNs impede the direct application of existing transferability assessment methods for ANNs to SNNs. SNN uses multiple time steps to generate results. Therefore, the feature map extracted by an SNN backbone is a two-dimensional matrix instead of a single-dimensional vector. Additionally, SNN models use nonlinear transform in classification tasks.

The most used leaky-integral and fire (LIF) spiking neuron model in SNNs is given by

$$\begin{cases} \tau \frac{d\mathbf{u}^n(t)}{dt} = -\mathbf{u}^n(t) + \theta^n \mathbf{o}^{n-1}(t), \\ o_i^n(t) = 1, \text{ if } u_i^n(t) \geq u_{th}, \\ o_i^n(t) = 0, \text{ if } u_i^n(t) < u_{th}. \end{cases} \quad (5)$$

Each neuron in the last layer is depicted by its membrane potential $u \in \mathbb{R}$ and output spike $o \in \{0, 1\}$. The membrane potential $u_i^n(t)$ will be reset to u_0 (usually $u_0 = 0$) if $o_i^n(t) = 1$. The spiking frequency is usually used as the logit of the target category. This nonlinear transform makes existing ANN transferability assessment methods incompatible with SNNs.

We propose to use features averaged through time steps for the evaluation of the transferability of SNN PTMs. Consider a pre-trained SNN model with T time steps, for input data x , the feature extracted by the model is $\mathbf{f}^{(t)} \in \mathbb{R}^D, t = 1, \dots, T$. We denote averaged feature as $\bar{\mathbf{f}} = \frac{1}{T} \sum_{t=1}^T \mathbf{f}^{(t)} \in \mathbb{R}^D$. To get outputs \mathbf{y} , the feature $\mathbf{f}^{(t)}$ is processed by equation 5 as

$$\mathbf{y} = \frac{1}{T} \sum_{t=1}^T \mathbf{o}^{(t)}, \quad (6)$$

$$\mathbf{o}^{(t)} = s(\mathbf{u}^{(t)} - \mathbf{u}_{th}), \quad (7)$$

$$\mathbf{u}^{(t)} = e^{-\frac{\Delta t}{\tau}} \mathbf{u}^{(t-1)} (\mathbf{1} - \mathbf{o}^{(t-1)}) + \theta^T \mathbf{f}^{(t)}. \quad (8)$$

where $e^{-\frac{\Delta t}{\tau}}$ denotes the decay of membrane potential in a time step Δt . $s(\cdot)$ is the function to describe the spiking behavior of a spiking neuron, giving 1 when input is non-negative and 0 otherwise. u_{th} is the threshold to spike and τ is a constant to describe the speed of decay.

In the classification problems where targets are 1 or 0, predicted results y_i should tend towards their targets. Suppose that y_i is equal to 1, then $o_i^{(t)} = 1$ for all $t = 1, \dots, T$. Therefore, the membrane potential in each time step is a linear transform of the feature of this time step as

$$u_i^{(t)} = \theta_i^T \mathbf{f}^{(t)}. \quad (9)$$

From equation 7, it is clear that $\theta_i^T \mathbf{f}^{(t)} = u_i^{(t)} \geq u_{th}, \forall t = 1, \dots, T$. Therefore, we can get that

$$\theta_i^T \bar{\mathbf{f}} = \theta_i^T \left(\frac{1}{T} \sum_{t=1}^T \mathbf{f}^{(t)} \right) = \frac{1}{T} \sum_{t=1}^T \theta_i^T \mathbf{f}^{(t)} \geq \frac{1}{T} \sum_{t=1}^T u_{th} = u_{th}. \quad (10)$$

Conversely, for $t_i = 0$, we have $\theta_i^T \mathbf{f} < u_{th}$. This shows that the SNN transform of a single layer can be approximated by a linear classification problem, and then transformed into a linear problem on average features.

3.3 Accelerating convergence with approximated maximum evidence

MacKay’s algorithm iteratively updates the regularization coefficients and maximizes the model evidence. However, the convergence of MacKay’s algorithm is not guaranteed (You et al., 2022), and the iteration number might be large in some cases. To reduce the time consumption and thus further improve the efficiency of model hub scheduling, here we propose an alternative method that avoids the iteration while giving an approximated result of maximum model evidence based on MacKay’s algorithm (hereinafter referred to as MacKay’s method). We name our proposed method *approximated maximum model evidence* method (hereinafter referred to as *the approximated method*).

The starting point of designing the approximated version is to make sure that our approximated results share the same properties as MacKay’s method. To achieve this, we state three properties of MacKay’s method. These properties describe the invariant of MacKay’s method with some operation on the input feature map. We design our approximated method adhering to these properties.

Properties 1 (invariant with stacking). Consider input data $\mathbf{F} \in \mathbb{R}^{N \times D}$ and labels $\mathbf{t} \in \mathbb{R}^N$. Suppose that MacKay’s algorithm converges to α^* and β^* and the maximized logarithm of model evidence is \mathcal{L}_1^* , then when input data is stacked q times as $[\mathbf{F}, \dots, \mathbf{F}] \in \mathbb{R}^{N \times qD}$, MacKay’s method converges to $q\alpha^*$ and β^* , and the maximized logarithm of model evidence \mathcal{L}_2^* equals exactly to \mathcal{L}_1^* .

Properties 2 (invariant with padding zero). Consider input data $\mathbf{F} \in \mathbb{R}^{N \times D}$ and labels $\mathbf{t} \in \mathbb{R}^N$. Suppose that MacKay’s algorithm converges to α^* and β^* and the maximized logarithm of model evidence is \mathcal{L}_1^* , then when input data is padded with zeros as $[\mathbf{F}, \mathbf{0}] \in \mathbb{R}^{N \times qD}$, MacKay’s method converges to $q\alpha^*$ and β^* , and the maximized logarithm of model evidence \mathcal{L}_2^* equals exactly to \mathcal{L}_1^* .

Properties 3 (invariant with scalar multiplication). Consider input data $\mathbf{F} \in \mathbb{R}^{N \times D}$ and labels $\mathbf{t} \in \mathbb{R}^N$. Suppose that MacKay’s algorithm converges to α^* and β^* and the maximized logarithm of model evidence is \mathcal{L}_1^* , then when input data is multiplied by a scalar $q > 0$ as $q\mathbf{F} \in \mathbb{R}^{N \times D}$, MacKay’s method converges to $q^2\alpha^*$ and β^* , and the maximized logarithm of model evidence \mathcal{L}_2^* equals exactly to \mathcal{L}_1^* .

The proof of properties 1, 2 can be found in reference You et al. (2022). The proof of properties 3 is given in appendix B.

The approximated maximum model evidence method is given below.

Consider input feature $\mathbf{F} \in \mathbb{R}^{N \times D}$ and labels $\mathbf{t} \in \mathbb{R}^N$. Let

$$\lambda \leftarrow \frac{\|\mathbf{F}\|_2^2}{N}, \quad (11)$$

$$m \leftarrow (\lambda \mathbf{I} + \mathbf{F}^T \mathbf{F})^{-1} \mathbf{F}^T \mathbf{t}, \quad (12)$$

$$\beta_0 \leftarrow \frac{N}{\|\mathbf{F}m - \mathbf{t}\|^2 + \lambda \|m\|^2}, \quad (13)$$

$$\alpha_0 \leftarrow \lambda \beta_0. \quad (14)$$

Then the approximate model evidence can be calculated by equation equation 3 with α_0 and β_0 .

Theorem 1. The model evidence calculated by our method satisfies properties 1, 2, and 3.

Proof. See appendix B. □

The experimental results in section 4.4 show that our method gives similar results to MacKay’s method when the number of samples N is larger than the dimension of features D . Therefore, the approximated result can serve as an approximated maximum model evidence under such circumstances. Moreover, with more reasonable prior values, MacKay’s method can be effectively accelerated.

4 Results

4.1 Experimental setup

Datasets The experiment is conducted on both neuromorphic datasets and static datasets. Neuromorphic datasets are vision datasets recorded with event cameras (Gallego et al., 2020), which are the commonly used benchmark for SNNs (He et al., 2020). As SNNs can also apply to conventional frame-based images (static datasets), we also validate our method on these datasets. For experiments on neuromorphic datasets, we use ES-ImageNet dataset (Lin et al., 2021) for pre-training and DVS128 Gesture, CIFAR10-DVS, N-Caltech101 and N-MNIST for fine-tuning, and for that on static datasets, the SNN models are pre-trained on ImageNet (Deng et al., 2009) and fine-tuned on CIFAR10, CIFAR100, Caltech101 and MNIST.

Pre-train SNN models We pre-train SNN models as the foundation of our experiment. We use spike-element-wise (SEW) ResNet as well as the spiking version of vanilla ResNet (Fang et al., 2021) as candidate PTMs. Eight different kinds of SNN models are used as different kinds of PTMs, the details of which are given in appendix A. For experiments on neuromorphic datasets, models are pre-trained with Adam

optimizer (Kingma & Ba, 2014) with momentum 0.9 and 0.999 for 10 epochs with 8 RTX3090 GPUs. We use a step learning rate scheduler with an initial learning rate of 0.01, step size 3, and $\gamma = 0.3$. For experiments on static datasets, we use PTMs from (Fang et al., 2021).

Fine-tune pre-trained models The fine-tuning of pre-trained SNN models are as the settings in logME (You et al., 2021). The parameters of the pre-trained neural network backbone are frozen during fine-tuning, and an FC layer is re-trained as the classifier of the PTMs. On neuromorphic datasets, each model is re-trained for 200 epochs. The hyper-parameters, learning rate, and weight decay, are determined by grid-searching from $1e-1$ to $1e-3$ and $1e-5$ to $1e-8$ respectively. On static datasets, each model is re-trained for 100 epochs. The same hyper-parameters are determined by grid-searching from $1e-1$ to $1e-3$ and $1e-5$ to $1e-7$ respectively. The test set accuracy of the best model on the validation set is used as the ground truth of transferability.

Assessment of transferability For each PTM, the transferability is assessed with MEAF, LEEP, and NCE. We calculate the logarithm of maximum model evidence on the train set. Note that ranking models with the maximum model evidence does not require the involvement of the validation or test set. This is also an advantage of the assessment method over the fine-tuning methods. Same as You et al. (2021), the Kendall correlation coefficient is used as a metric to evaluate the validity of our assessed score of transferability. Given N candidate models, $\{T_i, 1 \leq i \leq N\}$ is the list of their ground truth transferability and $\{S_i, 1 \leq i \leq N\}$ is the list of predicted scores. Kendall’s coefficient τ gives a quantitative illustration of how the ranking of S_i can stand for the ranking of T_i that when $S_i > S_j$, then the probability that $T_i > T_j$ is $\frac{\tau+1}{2}$ (Fagin et al., 2003).

4.2 Transferability assessment on neuromorphic datasets

We demonstrate the effectiveness of the maximum evidence method with averaged features on neuromorphic datasets. The results are shown in figure 1. It is shown that our MEAF method gives the best prediction result on all datasets, compared to NCE and LEEP. The Kendall coefficients are above 0.8 on all datasets, meaning that the predicted scores can represent the ranking of real transferability with a probability of more than 90%. It is also shown that LEEP also gives effective predictions on SNNs but with a lower Kendall coefficient. Note that on N-Caltech101 and N-MNIST, NCE gives predictions with correlation coefficients nearly equal to zero. The reason why information theory-based methods do not work well with SNNs is left for further research.

4.3 Transferability assessment on static datasets

We demonstrate the effectiveness of the MEAF method on static datasets, and the results are shown in figure 2. Similar to the results on neuromorphic datasets, MEAF gives the best prediction result on all datasets compared with NCE and LEEP. Although the Kendall coefficients are relatively low in this experiment, they are still above 0.6 on all datasets, meaning that the predicted scores can represent the ranking of real transferability with a probability of more than 80%. We suspect that the transferabilities (represented by the test set accuracy) of the PTMs are much closer in this experiment, thus the Kendall coefficients are relatively lower than on the neuromorphic datasets.

4.4 Validation of approximated maximum evidence

We first show that the approximated maximum evidence method gives consistent results with MacKay’s method when the number of samples is large. In this experiment, we generate toy data for linear regression and add Gaussian noise to the data with different SNRs from -10 dB to 10 dB. The dimension of the features (D) is fixed at 100 and the number of samples (N) increases from 100 to 3,000. The results are illustrated in figure 3, showing that the approximated maximum evidence converges to MacKay’s results as N tends to infinity.

We then validate that the approximated maximum evidence is effective in model hub scheduling. Figure 4 shows that with the approximated maximum evidence, the Kendall coefficients are above 0.85 on CIFAR10-

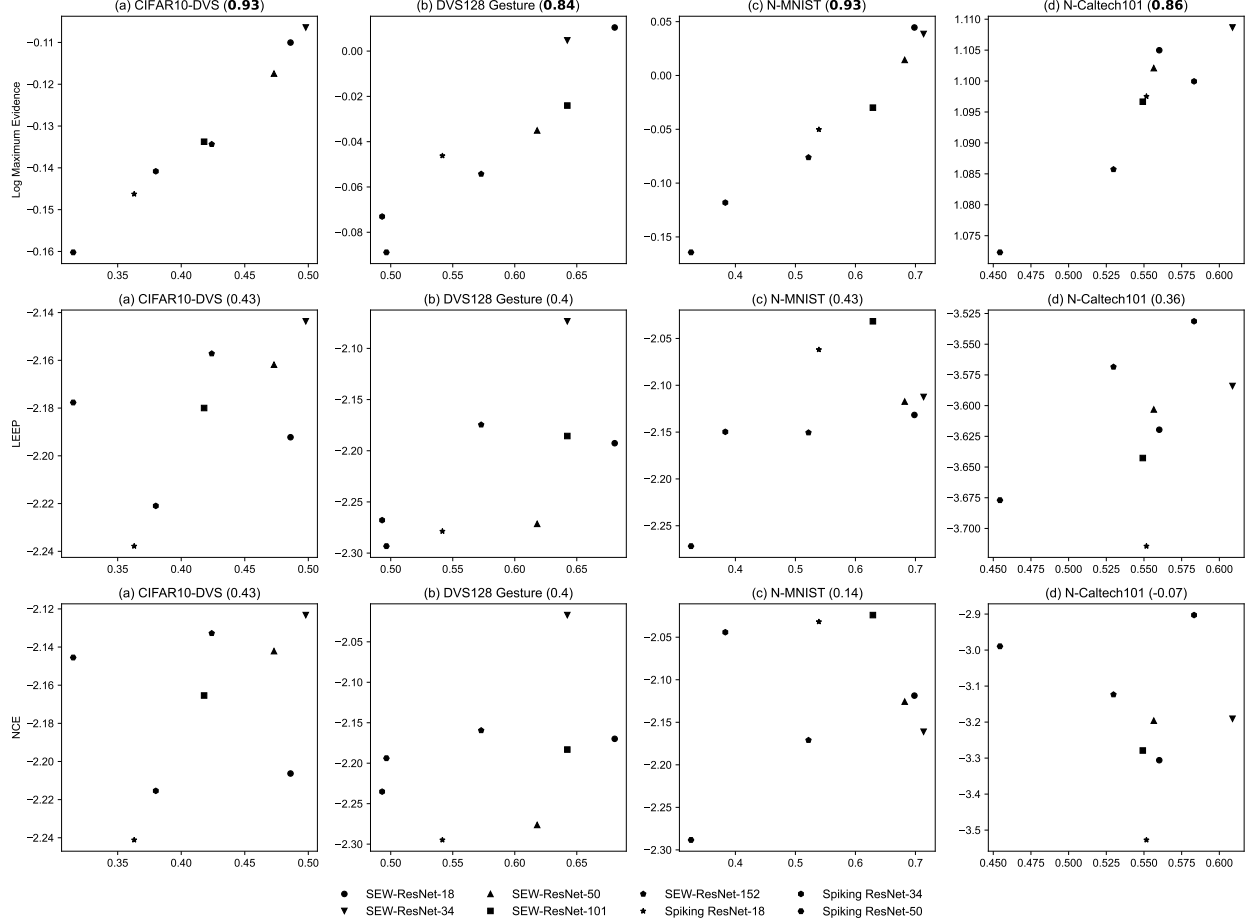


Figure 1: Result of assessed transferability on neuromorphic datasets. The abscissa represents the test set accuracy on the corresponding dataset, and the ordinate represents the scores of transferability given by different assessment methods. The first row gives the result of the MEAF. The second and third row shows the results of LEEP and NCE. Kendall correlation coefficients are labeled in the title of each graph after the dataset name.

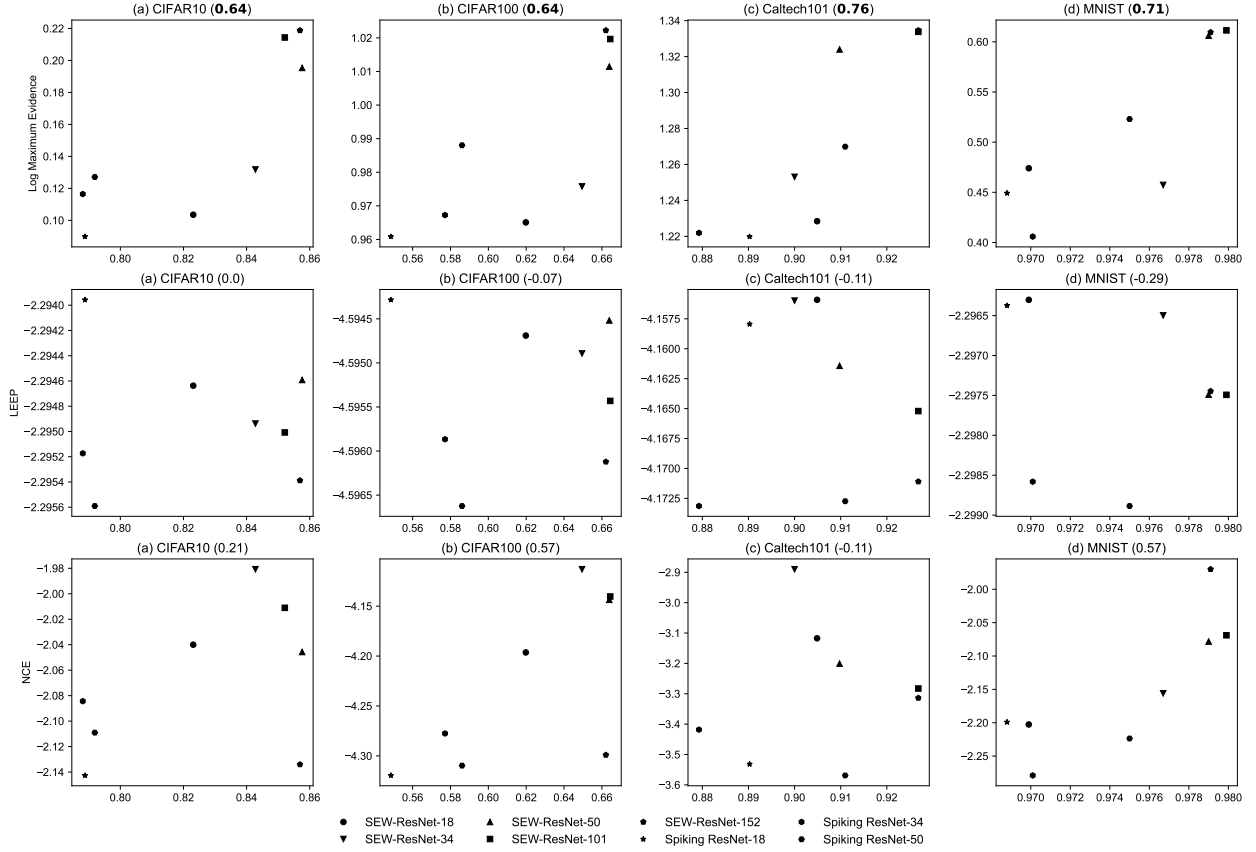


Figure 2: Result of assessed transferability on static datasets. The abscissa represents the test set accuracy on the corresponding dataset, and the ordinate represents the scores of transferability given by different assessment methods. Kendall correlation coefficients are labeled in the title of each graph after the dataset name.

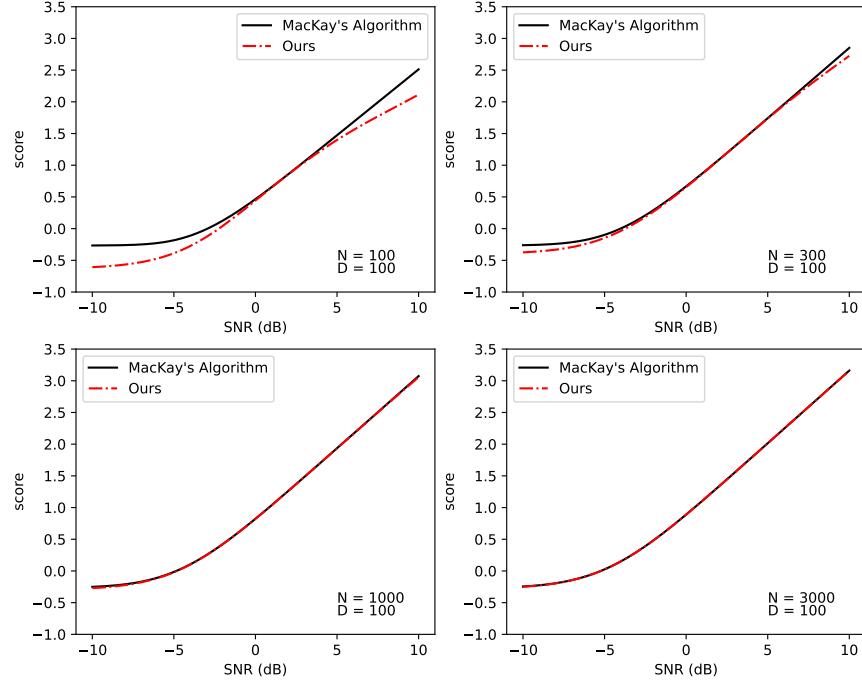


Figure 3: Consistency between the results of MacKay’s algorithm and our approximated result on toy data. Dimension of feature (D) is fixed as 100 and the number of samples (N) is 100, 3,000, 1,000, and 3,000 from (a) to (d). The abscissa represents the SNR of the toy data, and the ordinate represents the scores of transferability given by MacKay’s method and our approximated method.

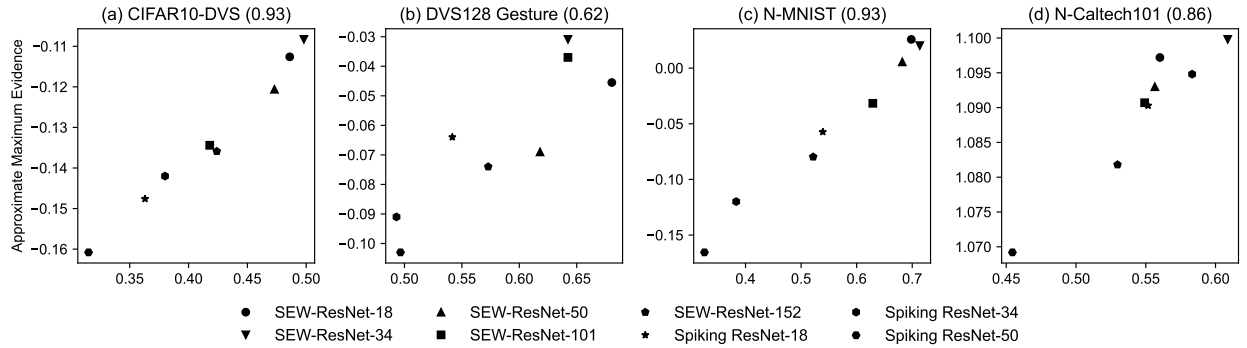


Figure 4: Result of assessed transferability on neuromorphic datasets with approximated maximum evidence. The abscissa represents the test set accuracy on the corresponding dataset, and the ordinate represents the scores of transferability given by the approximated maximum evidence method. Kendall correlation coefficients are labeled in the title of each graph after the dataset name.

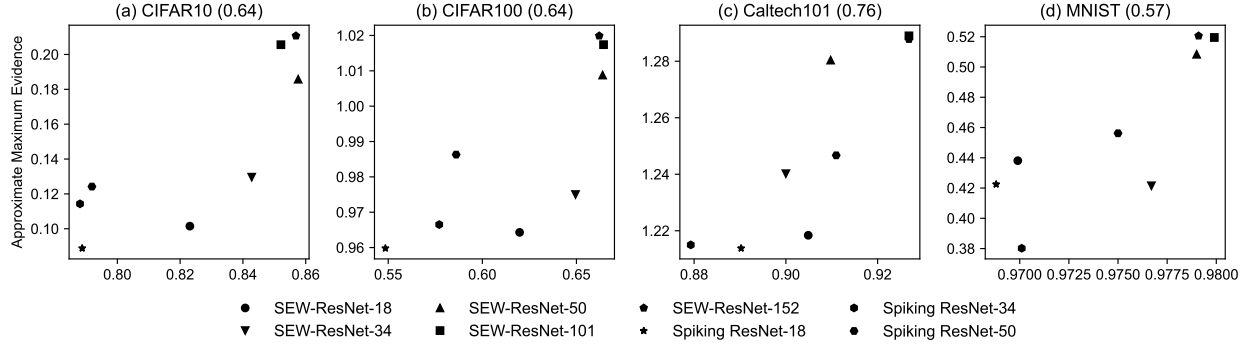


Figure 5: Result of assessed transferability on static datasets with approximated maximum evidence. The abscissa represents the test set accuracy on the corresponding dataset, and the ordinate represents the scores of transferability given by the approximated maximum evidence method. Kendall correlation coefficients are labeled in the title of each graph after the dataset name.

DVS, N-MNIST, and N-Caltech101 datasets, which is not lower than the results of MacKay’s method. On the DVS128 Gesture dataset, the Kendall coefficient is 0.62, which is still higher than LEEP (0.4) and NCE (0.4). The results on static datasets (Figure 5) give the Kendall coefficient not lower than the results by MacKay’s method in CIFAR10, CIFAR100 and Caltech101 datasets. For MNIST dataset, the Kendall coefficient is 0.57, lower than the result of MacKay’s method (0.71), but still comparable to the better results given by LEEP (-0.29) and NCE (0.57).

5 Discussion

This paper introduces a novel approach to assessing the transferability of SNN PTMs. We highlight the challenges of directly applying the existing Bayesian method to SNN transferability assessment and propose the MEAF method to address these issues. Additionally, an approximated maximum evidence method is introduced to reduce the computational time of model hub scheduling. We hope this work can help SNN PTMs be more widely and efficiently used.

Our work is limited in the diversity of SNN PTMs. There are currently no large-scale SNN model hubs, and our computational resources for pre-training are restricted. A more extensive collection of SNN PTMs would enhance the accuracy and generalizability of our conclusions.

Acknowledgments

Use unnumbered third level headings for the acknowledgments. All acknowledgments, including those to funding agencies, go at the end of the paper. Only add this information once your submission is accepted and deanonymized.

References

- Arnon Amir, Brian Taba, David Berg, Timothy Melano, Jeffrey McKinstry, Carmelo Di Nolfo, Tapan Nayak, Alexander Andreopoulos, Guillaume Garreau, Marcela Mendoza, et al. A low power, fully event-based gesture recognition system. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 7243–7252, 2017.
- Shai Ben-David and Reba Schuller. Exploiting task relatedness for multiple task learning. In *Learning Theory and Kernel Machines: 16th Annual Conference on Learning Theory and 7th Kernel Workshop, COLT/Kernel 2003, Washington, DC, USA, August 24-27, 2003. Proceedings*, pp. 567–580. Springer, 2003.

- Shai Ben-David, John Blitzer, Koby Crammer, and Fernando Pereira. Analysis of representations for domain adaptation. *Advances in neural information processing systems*, 19, 2006.
- Zhenshan Bing, Zhuangyi Jiang, Long Cheng, Caixia Cai, Kai Huang, and Alois Knoll. End to end learning of a multi-layered snn based on r-stdp for a target tracking snake-like robot. In *2019 International Conference on Robotics and Automation (ICRA)*, pp. 9645–9651. IEEE, 2019.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- Manon Dampfhoffer, Thomas Mesquida, Alexandre Valentian, and Lorena Anghel. Are snns really more energy-efficient than anns? an in-depth hardware-aware study. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 2022.
- Erik Daxberger, Agustinus Kristiadi, Alexander Immer, Runa Eschenhagen, Matthias Bauer, and Philipp Hennig. Laplace redux-effortless bayesian deep learning. *Advances in Neural Information Processing Systems*, 34:20089–20103, 2021.
- Guillaume Debat, Tushar Chauhan, Benoit R Cottureau, Timothée Masquelier, Michel Paindavoine, and Robin Baures. Event-based trajectory prediction using spiking neural networks. *Frontiers in computational neuroscience*, 15:658764, 2021.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255. Ieee, 2009.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Zhengxiao Du, Yujie Qian, Xiao Liu, Ming Ding, Jiezhong Qiu, Zhilin Yang, and Jie Tang. Glm: General language model pretraining with autoregressive blank infilling. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 320–335, 2022.
- Ronald Fagin, Ravi Kumar, and Dakshinamurthi Sivakumar. Comparing top k lists. *SIAM Journal on discrete mathematics*, 17(1):134–160, 2003.
- Wei Fang, Zhaofei Yu, Yanqi Chen, Tiejun Huang, Timothée Masquelier, and Yonghong Tian. Deep residual learning in spiking neural networks. *Advances in Neural Information Processing Systems*, 34:21056–21069, 2021.
- Guillermo Gallego, Tobi Delbrück, Garrick Orchard, Chiara Bartolozzi, Brian Taba, Andrea Censi, Stefan Leutenegger, Andrew J Davison, Jörg Conradt, Kostas Daniilidis, et al. Event-based vision: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 44(1):154–180, 2020.
- Stephen F Gull. Developments in maximum entropy data analysis. In *Maximum Entropy and Bayesian Methods: Cambridge, England, 1988*, pp. 53–71. Springer, 1989.
- Jesse Hagenaars, Federico Paredes-Vallés, and Guido De Croon. Self-supervised learning of event-based optical flow with spiking neural networks. *Advances in Neural Information Processing Systems*, 34:7167–7179, 2021.
- Kaiming He, Ross Girshick, and Piotr Dollár. Rethinking imagenet pre-training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 4918–4927, 2019.
- Weihua He, YuJie Wu, Lei Deng, Guoqi Li, Haoyu Wang, Yang Tian, Wei Ding, Wenhui Wang, and Yuan Xie. Comparing snns and rnns on neuromorphic vision datasets: Similarities and differences. *Neural Networks*, 132:108–120, 2020.

- Alexander Immer, Matthias Bauer, Vincent Fortuin, Gunnar Rätsch, and Khan Mohammad Emtiyaz. Scalable marginal likelihood estimation for model selection in deep learning. In *International Conference on Machine Learning*, pp. 4563–4573. PMLR, 2021.
- Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International conference on machine learning*, pp. 4904–4916. PMLR, 2021.
- Junguang Jiang, Yang Shu, Jianmin Wang, and Mingsheng Long. Transferability in deep learning: A survey. *arXiv preprint arXiv:2201.05867*, 2022.
- Yong-Deok Kim, Taewoong Jang, Bohyung Han, and Seungjin Choi. Learning to select pre-trained deep representations with bayesian evidence framework. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5318–5326, 2016.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kevin H Knuth, Michael Habeck, Nabin K Malakar, Asim M Mubeen, and Ben Placek. Bayesian evidence and model selection. *Digital Signal Processing*, 47:50–67, 2015.
- Simon Kornblith, Jonathon Shlens, and Quoc V Le. Do better imagenet models transfer better? In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 2661–2671, 2019.
- Frederik Kunstner, Philipp Hennig, and Lukas Balles. Limitations of the empirical fisher approximation for natural gradient descent. *Advances in neural information processing systems*, 32, 2019.
- Yihan Lin, Wei Ding, Shaohua Qiang, Lei Deng, and Guoqi Li. Es-imagenet: A million event-stream classification dataset for spiking neural networks. *Frontiers in neuroscience*, 15:1546, 2021.
- Wolfgang Maass. Networks of spiking neurons: the third generation of neural network models. *Neural networks*, 10(9):1659–1671, 1997.
- David JC MacKay. Bayesian interpolation. *Neural computation*, 4(3):415–447, 1992a.
- David JC MacKay. The evidence framework applied to classification networks. *Neural computation*, 4(5):720–736, 1992b.
- David JC MacKay. A practical bayesian framework for backpropagation networks. *Neural computation*, 4(3):448–472, 1992c.
- David JC MacKay. Comparison of approximate methods for handling hyperparameters. *Neural computation*, 11(5):1035–1068, 1999.
- TorchVision maintainers and contributors. Torchvision: Pytorch’s computer vision library. <https://github.com/pytorch/vision>, 2016.
- Yishay Mansour, Mehryar Mohri, and Afshin Rostamizadeh. Domain adaptation: Learning bounds and algorithms. *arXiv preprint arXiv:0902.3430*, 2009.
- Cuong Nguyen, Tal Hassner, Matthias Seeger, and Cedric Archambeau. Leep: A new measure to evaluate transferability of learned representations. In *International Conference on Machine Learning*, pp. 7294–7305. PMLR, 2020.
- Jing Pei, Lei Deng, Sen Song, Mingguo Zhao, Youhui Zhang, Shuang Wu, Guanrui Wang, Zhe Zou, Zhenzhi Wu, Wei He, et al. Towards artificial general intelligence with hybrid tianjic chip architecture. *Nature*, 572(7767):106–111, 2019.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.

- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pp. 8748–8763. PMLR, 2021.
- Kaushik Roy, Akhilesh Jaiswal, and Priyadarshini Panda. Towards spike-based machine intelligence with neuromorphic computing. *Nature*, 575(7784):607–617, 2019.
- Anh T Tran, Cuong V Nguyen, and Tal Hassner. Transferability and hardness of supervised classification tasks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 1395–1405, 2019.
- Xiangwen Wang, Xianghong Lin, and Xiaochao Dang. Supervised learning in spiking neural networks: A review of algorithms and evaluations. *Neural Networks*, 125:258–280, 2020.
- Larry Wasserman. Bayesian model selection and model averaging. *Journal of mathematical psychology*, 44(1):92–107, 2000.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pp. 38–45, Online, October 2020. Association for Computational Linguistics. URL <https://www.aclweb.org/anthology/2020.emnlp-demos.6>.
- David Wolpert. On the use of evidence in neural networks. *Advances in neural information processing systems*, 5, 1992.
- Yujie Wu, Lei Deng, Guoqi Li, Jun Zhu, Yuan Xie, and Luping Shi. Direct training for spiking neural networks: Faster, larger, better. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pp. 1311–1318, 2019.
- Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. How transferable are features in deep neural networks? *Advances in neural information processing systems*, 27, 2014.
- Kaichao You, Yong Liu, Jianmin Wang, and Mingsheng Long. Logme: Practical assessment of pre-trained models for transfer learning. In *International Conference on Machine Learning*, pp. 12133–12143. PMLR, 2021.
- Kaichao You, Yong Liu, Ziyang Zhang, Jianmin Wang, Michael I Jordan, and Mingsheng Long. Ranking and tuning pre-trained models: a new paradigm for exploiting model hubs. *The Journal of Machine Learning Research*, 23(1):9400–9446, 2022.
- Lin Zhu, Xiao Wang, Yi Chang, Jianing Li, Tiejun Huang, and Yonghong Tian. Event-based video reconstruction via potential-assisted spiking neural network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3594–3604, 2022.

A Appendix

A.1 Details of pre-training SNNs

We use SEW ResNet Fang et al. (2021) and a spiking version of vanilla ResNet as candidate PTMs. On neuromorphic datasets, the accuracy on valid datasets is shown in Table 1. We use the same model architectures as in Fang et al. (2021). To adapt the neuromorphic dataset, we modified the input layer of the model, as shown in Table 2. The unchanged hyper-parameters are not notified in the table. The number of time steps on static datasets is 4 to be consistent with Fang et al. (2021) and 8 on neuromorphic datasets for better leveraging the temporal dimension of SNNs.

Table 1: Pre-training SNN models on ES-ImageNet dataset

MODEL	TOP 1 ACCURACY	TOP 5 ACCURACY
SEW ResNet-18	0.3443	0.5840
SEW ResNet-34	0.3507	0.5922
SEW ResNet-50	0.3140	0.5470
SEW ResNet-101	0.3031	0.5347
SEW ResNet-152	0.2615	0.4956
Spiking ResNet-18	0.3156	0.5486
Spiking ResNet-34	0.2376	0.4475
Spiking ResNet-50	0.1661	0.3568

Table 2: Modification of the input layer of the model on neuromorphic datasets

ORIGINAL (STATIC)	NEUROMORPHIC
Conv2d(in_channels=3)	Conv2d(in_channels=2)
Batch normalization	Batch normalization
LIF	AvgPool2d(kernel_size=3)
MaxPool2d(kernel_size=3)	LIF

A.2 Proofs

A.2.1 Proof of Properties 3

Proof. The logarithm of maximum evidence is calculated as equation 3,

$$\mathcal{L}_1^* = \frac{N}{2} \ln \beta^* + \frac{D}{2} \ln \alpha^* - \frac{N}{2} \ln 2\pi - \frac{\beta^*}{2} \|\mathbf{F}m - \mathbf{t}\|^2 - \frac{\beta^*}{2} m^T m - \frac{1}{2} \ln \det \mathbf{A}, \quad (15)$$

where $\mathbf{A} = \alpha^* \mathbf{I} + \beta^* \mathbf{F}^T \mathbf{F}$, $m = \beta^* \mathbf{A}^{-1} \mathbf{F}^T \mathbf{t}$.

When the input data \mathbf{F} is multiplied by q as $\tilde{\mathbf{F}} = q\mathbf{F}$, let $\tilde{\alpha}$ be $q^2\alpha^*$ and $\tilde{\beta}$ be β^* , we have

$$\tilde{\mathbf{A}} = q^2\alpha^*\mathbf{I} + \beta^*q\mathbf{F}^Tq\mathbf{F} = q^2\mathbf{A}, \quad (16)$$

$$\tilde{m} = \beta^*\tilde{\mathbf{A}}^{-1}\tilde{\mathbf{F}}^T\mathbf{t} = \frac{1}{q}m, \quad (17)$$

$$\det \tilde{\mathbf{A}} = q^{2D} \det \mathbf{A}. \quad (18)$$

\mathcal{L}_2^* is the maximum of logarithm of model evidence over α and β . Therefore we have

$$\begin{aligned} \mathcal{L}_2^* &\geq \mathcal{L}_2(q^2\alpha^*, \beta^*) \\ &= \frac{N}{2} \ln \beta^* + \frac{D}{2} \ln q^2\alpha^* - \frac{N}{2} \ln 2\pi - \frac{\beta^*}{2} \|\tilde{\mathbf{F}}\tilde{m} - \mathbf{t}\|^2 - \frac{\beta^*}{2} \tilde{m}^T \tilde{m} - \frac{1}{2} \ln \det \tilde{\mathbf{A}} \\ &= \frac{N}{2} \ln \beta^* + \frac{D}{2} \ln \alpha^* - \frac{N}{2} \ln 2\pi - \frac{\beta^*}{2} \|\mathbf{F}m - \mathbf{t}\|^2 - \frac{\beta^*}{2} m^T m - \frac{1}{2} \ln \det \mathbf{A} \\ &= \mathcal{L}_1^*. \end{aligned} \quad (19)$$

Similarly we have $\mathcal{L}_1^* \geq \mathcal{L}_2^*$. Therefore $\mathcal{L}_1^* = \mathcal{L}_2^*$. □

A.2.2 Proof of Theorem 1

Proof. Consider input data \mathbf{F} and label \mathbf{t} . Let approximated maximum model evidence as \mathcal{L}_1^A .

We first prove the case of properties 1. When the input data \mathbf{F} is stacked q times as $\tilde{\mathbf{F}} = [\mathbf{F}, \dots, \mathbf{F}]$. $\|\tilde{\mathbf{F}}\|_2 = \tilde{\sigma}_1$ increases \sqrt{q} times, where $\tilde{\sigma}_1$ denotes the first singular value of $\tilde{\mathbf{F}}$. N is unchanged in this case. Therefore, $\tilde{\lambda}$ is q times larger for $\tilde{\mathbf{F}}$.

From the proof of properties 1 in You et al. (2022), we have

$$\tilde{m} = \begin{bmatrix} \frac{1}{q}m \\ \vdots \\ \frac{1}{q}m \end{bmatrix}. \quad (20)$$

Therefore we have $\tilde{\beta}_0 = \beta_0, \tilde{\alpha} = \tilde{\lambda}\tilde{\beta} = q\alpha$. From the proof of properties 1 in You et al. (2022), we have $\mathcal{L}_2^A = \mathcal{L}_1^A$.

Then we prove the case of properties 2. When the input data \mathbf{F} is stacked q times as $\tilde{\mathbf{F}} = [\mathbf{F}, \mathbf{0}]$. $\|\tilde{\mathbf{F}}\|_2 = \tilde{\sigma}_1$ keeps the same and N is unchanged as well. Therefore, $\tilde{\lambda}$ unchanged for $\tilde{\mathbf{F}}$.

From the proof of properties 1 in You et al. (2022), we have

$$\tilde{m} = \begin{bmatrix} m \\ \mathbf{0}_{(q-1)D \times 1} \end{bmatrix}. \quad (21)$$

Therefore we have $\tilde{\beta}_0 = \beta_0, \tilde{\alpha} = \tilde{\lambda}\tilde{\beta} = \alpha$. From the proof of properties 1 in You et al. (2022), we have $\mathcal{L}_2^A = \mathcal{L}_1^A$.

We prove the case if properties 3. When the input data \mathbf{F} is multiplied by q as $\tilde{\mathbf{F}} = q\mathbf{F}$. $\|\tilde{\mathbf{F}}\|_2 = \tilde{\sigma}_1$ increases q times and N is unchanged. Therefore, $\tilde{\lambda}$ increases q^2 times for $\tilde{\mathbf{F}}$.

From the proof of properties 3, we have

$$\tilde{m} = \frac{1}{q}m. \quad (22)$$

Therefore we have $\tilde{\beta}_0 = \beta_0, \tilde{\alpha} = \tilde{\lambda}\tilde{\beta} = q^2\alpha$. From the proof of properties 3, we have $\mathcal{L}_2^A = \mathcal{L}_1^A$. \square

A.3 Validation of approximated maximum evidence

Table 3: Iteration number used for MacKay’s method

MODEL	CIFAR10-DVS	DVS128Gesture	N-MNIST	N-Caltech101	CIFAR10	CIFAR100	Caltech101	MNIST
SEW ResNet-18	6.0	5.2	2.3	7.9	2.0	3.2	4.3	1.1
SEW ResNet-34	5.6	5.2	3.3	7.6	2.0	3.0	4.0	1.4
SEW ResNet-50	4.5	6.3	4.0	7.1	3.1	4.4	6.5	2.0
SEW ResNet-101	3.6	6.6	3.5	6.3	3.2	4.3	6.3	2.1
SEW ResNet-152	4.1	6.8	3.7	5.3	3.1	4.2	6.3	2.0
Spiking ResNet-18	6.0	5.3	3.1	8.0	2.1	3.4	4.6	1.1
Spiking ResNet-34	5.0	5.3	2.9	7.0	2.1	3.4	4.8	0.7
Spiking ResNet-50	3.5	5.9	2.9	6.2	3.2	5.0	7.2	2.3
Average	4.8	5.8	3.2	6.9	2.6	3.9	5.5	1.6

Here we give the average number of iterations required for MacKay’s method in Table 3. The stopping criterion for the convergence of MacKay’s algorithm is set such that the relative error between two successive α/β is less than 0.01. This is a relatively lenient stopping criterion; employing a stricter criterion would result in a higher number of iterations.

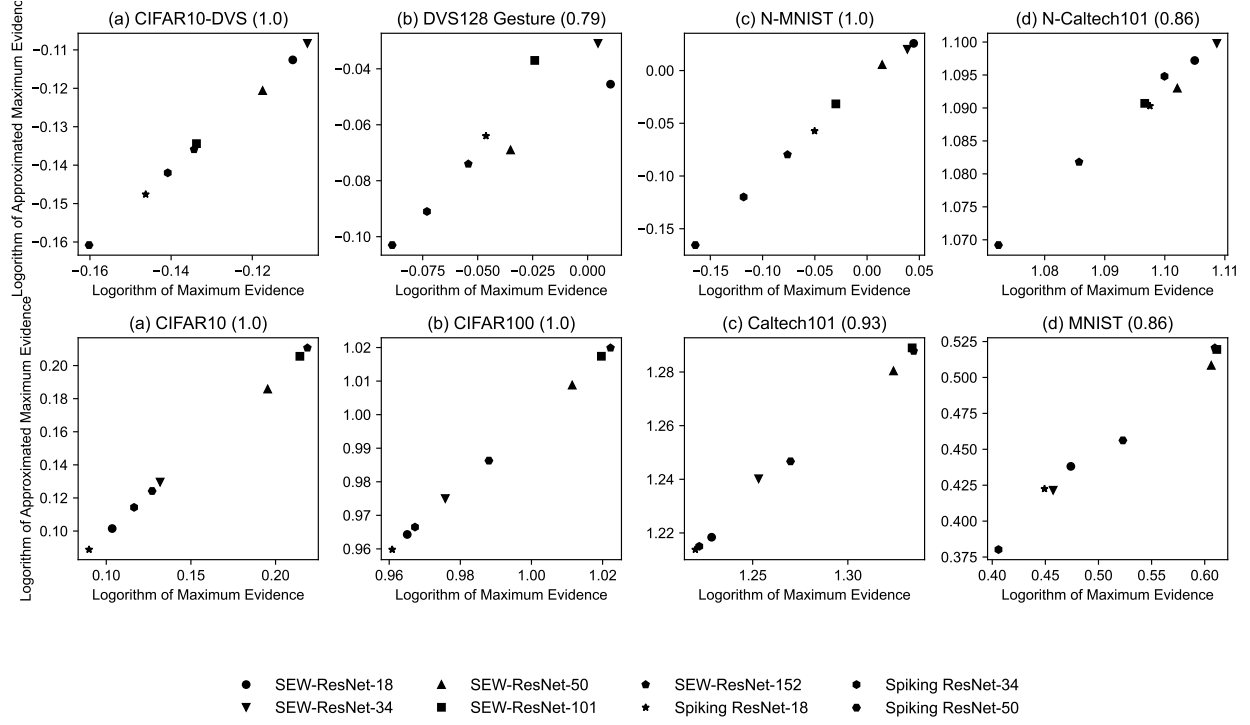


Figure 6: Compare scores given by the approximated method and MacKay's method. Kendall correlation coefficients are labeled in the title of each graph after the dataset name.

We also demonstrated a strong correlation between the approximated maximum evidence and those obtained using MacKay's algorithm⁶. It is worth noting that the correlation is relatively low on the DVS128Gesture dataset. We believe this is because the DVS128 Gesture is a very small dataset, with only 1176 samples in the training set (Amir et al., 2017), which contradicts our assumption of the number of samples being much greater than the dimension of features.