
NP-Match: When Neural Processes meet Semi-Supervised Learning

Jianfeng Wang¹ Thomas Lukasiewicz^{2,1} Daniela Massiceti³ Xiaolin Hu⁴ Vladimir Pavlovic⁵
Alexandros Neophytou⁶

Abstract

Semi-supervised learning (SSL) has been widely explored in recent years, and it is an effective way of leveraging unlabeled data to reduce the reliance on labeled data. In this work, we adjust neural processes (NPs) to the semi-supervised image classification task, resulting in a new method named NP-Match. NP-Match is suited to this task for two reasons. Firstly, NP-Match implicitly compares data points when making predictions, and as a result, the prediction of each unlabeled data point is affected by the labeled data points that are similar to it, which improves the quality of pseudo-labels. Secondly, NP-Match is able to estimate uncertainty that can be used as a tool for selecting unlabeled samples with reliable pseudo-labels. Compared with uncertainty-based SSL methods implemented with Monte Carlo (MC) dropout, NP-Match estimates uncertainty with much less computational overhead, which can save time at both the training and the testing phases. We conducted extensive experiments on four public datasets, and NP-Match outperforms state-of-the-art (SOTA) results or achieves competitive results on them, which shows the effectiveness of NP-Match and its potential for SSL.

1. Introduction

Deep neural networks have been widely used in computer vision tasks (Krizhevsky et al., 2012; Simonyan & Zisserman, 2014; Szegedy et al., 2015; 2016; He et al., 2016) due to their strong performance. Training deep neural net-

works relies on large-scale labeled datasets, but annotating large-scale datasets is time-consuming, which encourages researchers to explore semi-supervised learning (SSL). SSL aims to learn from few labeled data and a large amount of unlabeled data, and it has been a long-standing problem in computer vision and machine learning (Sohn et al., 2020; Zhang et al., 2021; Rizve et al., 2021; Pham et al., 2021; Li & Zhou, 2014; Liu et al., 2010; Berthelot et al., 2019; 2020). In this work, we focus on SSL for image classification.

Most recent approaches to SSL for image classification are based on the combination of consistency regularization and pseudo-labeling (Sohn et al., 2020; Li et al., 2021; Rizve et al., 2021; Zhang et al., 2021; Nassar et al., 2021; Pham et al., 2021; Hu et al., 2021). They can be further classified into two categories, namely, deterministic (Sohn et al., 2020; Li et al., 2021; Zhang et al., 2021; Nassar et al., 2021; Pham et al., 2021; Hu et al., 2021) and probabilistic ones (Rizve et al., 2021). A deterministic approach aims at directly making predictions, while a probabilistic approach tries to additionally model the predictive distribution, such as using Bayesian neural networks (BNNs), which are implemented by Monte Carlo (MC) dropout (Gal & Ghahramani, 2016). As a result, the former cannot estimate the uncertainty of the model’s prediction, and unlabeled samples are selected only based on high-confidence predictions. In contrast, the latter can give uncertainties for unlabeled samples, and the uncertainties can be combined with high-confidence predictions for picking or refining pseudo-labels.

Current SOTA methods for the semi-supervised image classification task are deterministic, including FixMatch (Sohn et al., 2020), CoMatch (Li et al., 2021), and FlexMatch (Zhang et al., 2021), which have achieved promising results on public benchmarks. In contrast, progress on probabilistic approaches lags behind, which is mainly shown by the fact that there are only few studies on this task and MC dropout becomes the only option for implementing the probabilistic model (Rizve et al., 2021). In addition, MC dropout also dominates the uncertainty-based approaches to other SSL tasks (Sedai et al., 2019; Shi et al., 2021; Wang et al., 2021; Yu et al., 2019; Zhu et al., 2020). MC dropout, however, is time-consuming, requiring several feedforward passes to get uncertainty at both the training and the testing stages, especially when some large models are used.

¹Department of Computer Science, University of Oxford, UK. ²Institute of Logic and Computation, TU Wien, Austria. ³Microsoft Research, Cambridge, UK. ⁴Department of Computer Science and Technology, Tsinghua University, Beijing, China. ⁵Department of Computer Science, Rutgers University, New Jersey, USA. ⁶Microsoft, Applied Science Group, Reading, UK. Correspondence to: Jianfeng Wang <jianfeng.wang@cs.ox.ac.uk>.

To solve this drawback and to further promote the related research, we need to find better probabilistic approaches for SSL. Considering that MC dropout is an approximation to the Gaussian process (GP) model (Gal & Ghahramani, 2016), we turn to another approximation model called neural processes (NPs) (Garnelo et al., 2018b), which can be regarded as an NN-based formulation that approximates GPs. Similarly to a GP, a neural process is also a probabilistic model that defines distributions over functions. Thus, an NP is able to rapidly adapt to new observations, with the advantage of estimating the uncertainty of each observation. There are two main aspects that motivate us to investigate NPs in SSL. Firstly, GPs have been preliminarily explored for some SSL tasks (Sindhwani et al., 2007; Jean et al., 2018; Yasarla et al., 2020), because of the property that their kernels are able to compare labeled data with unlabeled data when making predictions. NPs share this property, since it has been proved that NPs can learn non-trivial implicit kernels from data (Garnelo et al., 2018b). As a result, NPs are able to make predictions for target points conditioned on context points. This feature is highly relevant to SSL, which must learn from limited labeled samples in order to make predictions for unlabeled data, similarly to how NPs are able to impute unknown pixel values (i.e., target points) when given only a small number of known pixels (namely, context points) (Garnelo et al., 2018b). Due to the learned implicit kernels in NPs (Garnelo et al., 2018b) and the successful application of GPs to different SSL tasks (Sindhwani et al., 2007; Jean et al., 2018; Yasarla et al., 2020), NPs could be a suitable probabilistic model for SSL, as the kernels can compare labeled data with unlabeled data in order to improve the quality of pseudo-labels for the unlabeled data at the training stage. Secondly, previous GP-based works for SSL do not explore the semi-supervised large-scale image classification task, since GPs are computationally expensive, which usually incur a $\mathcal{O}(n^3)$ runtime for n training points. But, unlike GPs, NPs are more efficient than GPs, providing the possibility of applying NPs to this task. NPs are also computationally significantly more efficient than current MC-dropout-based approaches to SSL, since, given an input image, they only need to perform one feedforward pass to obtain the prediction with an uncertainty estimate.

In this work, we take the first step to explore NPs in large-scale semi-supervised image classification, and propose a new probabilistic method called NP-Match. NP-Match still rests on the combination of consistency regularization and pseudo-labeling, but it incorporates NPs to the top of deep neural networks, and therefore it is a probabilistic approach. Compared to the previous probabilistic method for semi-supervised image classification (Rizve et al., 2021), NP-Match not only can make predictions and estimate uncertainty more efficiently, inheriting the advantages of NPs, but also can achieve a better performance on public benchmarks.

Summarizing, the main contributions of this paper are:

- We propose NP-Match, which adjusts NPs to SSL, and explore its use in semi-supervised large-scale image classification. To our knowledge, this is the first such work. In addition, NP-Match has the potential to break the monopoly of MC dropout as the probabilistic model in SSL.
- We experimentally show that the Kullback-Leibler (KL) divergence in the evidence lower bound (ELBO) of NPs (Garnelo et al., 2018b) is not a good choice in the context of SSL, which may negatively impact the learning of global latent variables. To tackle this problem, we propose a new uncertainty-guided skew-geometric Jensen-Shannon (JS) divergence ($JS^{G_{\alpha_u}}$) for NP-Match.
- We show that NP-Match outperforms SOTA results or achieves competitive results on four public benchmarks, demonstrating its effectiveness for SSL. We also show that NP-Match estimates uncertainty faster than the MC-dropout-based probabilistic model, which can improve the training and the test efficiency.

The rest of this paper is organized as follows. In Section 2, we review related methods. Section 3 presents NP-Match and the uncertainty-guided skew-geometric JS divergence ($JS^{G_{\alpha_u}}$), followed by the experimental settings and results in Section 4. In Section 5, we give a summary and an outlook on future research. The source code is available at: <https://github.com/Jianf-Wang/NP-Match>.

2. Related Work

We now briefly review related works, including *semi-supervised learning (SSL) for image classification*, *Gaussian processes (GPs) for SSL*, and *neural processes (NPs)*.

SSL for image classification. Most methods for semi-supervised image classification in the past few years are based on pseudo-labeling and consistency regularization. Pseudo-labeling approaches rely on the high confidence of pseudo-labels, which can be added to the training data set as labeled data, and those approaches can be classified into two classes, namely, disagreement-based models and self-training models. The former models aim to train multiple learners and exploit the disagreement during the learning process (Qiao et al., 2018; Chen et al., 2018), while the latter models aim at training the model on a small amount of labeled data, and then using its predictions on the unlabeled data as pseudo-labels (Lee, 2013; Zhai et al., 2019; Wang et al., 2020; Pham et al., 2021). Consistency-regularization-based approaches work by performing different transformations on an input image and adding a regularization term to make their predictions consistent (Bachman et al., 2014; Sajjadi et al., 2016; Laine & Aila, 2017; Berthelot et al., 2019; Xie et al., 2020). Based on these two approaches, FixMatch

(Sohn et al., 2020) is proposed, which achieves new state-of-the-art (SOTA) results on the most commonly-studied SSL benchmarks. FixMatch (Sohn et al., 2020) combines the merits of these two approaches: given an unlabeled image, weak data augmentation and strong data augmentation are performed on the image, leading to two versions of the image, and then FixMatch produces a pseudo-label based on its weakly-augmented version and a preset confidence threshold, which is used as the true label for its strongly augmented version to train the whole framework. The success of FixMatch inspired several subsequent methods (Li et al., 2021; Rizve et al., 2021; Zhang et al., 2021; Nassar et al., 2021; Pham et al., 2021; Hu et al., 2021). For instance, Li et al. (2021) additionally design the classification head and the projection head for generating a class probability and a low-dimensional embedding, respectively. The projection head and the classification head are jointly optimized during training. Specifically, the former is learnt with contrastive learning on pseudo-label graphs to encourage the embeddings of samples with similar pseudo-labels to be close, and the latter is trained with pseudo-labels that are smoothed by aggregating information from nearby samples in the embedding space. Zhang et al. (2021) propose to use dynamic confidence thresholds that are automatically adjusted according to the model’s learning status of each class. Rizve et al. (2021) propose an uncertainty-aware pseudo-label selection (UPS) framework for semi-supervised image classification. The UPS framework introduces MC dropout to obtain uncertainty estimates, which are then leveraged as a tool for selecting pseudo-labels. This is the first work using MC dropout for semi-supervised image classification.

GPs for SSL. Since NPs are also closely related to GPs, we review the application of GPs to different SSL tasks in this part. GPs, which are non-parametric models, have been preliminarily investigated in different semi-supervised learning tasks. For example, Sindhvani et al. (2007) introduce a semi-supervised GP classifier, which incorporates the information of relationships among labeled and unlabeled data into the kernel. Their approach, however, has high computational costs and is thus only evaluated for a simple binary classification task on small datasets. Deep kernel learning (Wilson et al., 2016) also lies on the spectrum between NNs and GPs, and has been integrated into a new framework for the semi-supervised regression task, named semi-supervised deep kernel learning (Jean et al., 2018), which aims to minimize the predictive variance for unlabeled data, encouraging unlabeled embeddings to be near labeled embeddings. Semi-supervised deep kernel learning, however, has not been applied to SSL image classification, and (similarly to semi-supervised GPs) also comes with a high (cubic) computational complexity. Recently, Yasarla et al. (2020) proposed to combine GPs with UNet (Ronneberger et al., 2015) for SSL image deraining. Here, GPs are used to get

pseudo-labels for unlabeled samples based on the feature representations of labeled and unlabeled images. GPs have also been combined with graph convolutional networks for semi-supervised learning on graphs (Ng et al., 2018; Walker & Glocker, 2019; Liu et al., 2020). Although many previous works explore GPs in different semi-supervised learning tasks, none of them investigates the application of GPs to semi-supervised large-scale image classification.

NPs. The origin of NPs can be traced back to conditional NPs (Garnelo et al., 2018a), which define conditional distributions over functions given a set of observations. Conditional NPs, however, do not introduce global latent variables for observations, which led to the birth of NPs (Garnelo et al., 2018b). In NPs, the mean-aggregator is used to summarize the encoded inputs of a task into a global latent variable, which is then used to make predictions on targets in the task. In recent years, several NP variants have been proposed to better approximate stochastic processes. For example, Kim et al. (2019) consider that the mean-aggregator may cause difficulties for the decoder to pick relevant information with regard to making predictions, and they introduce a differentiable attention mechanism to solve this issue, resulting in new attentive NPs. Gordon et al. (2020) consider that the translation equivariance is important for prediction problems, which should be considered. Therefore, they incorporate translation equivariance into NPs and design a new model, called convolutional conditional NPs. Besides, Louizos et al. (2019) consider that using global latent variables is not flexible for encoding inductive biases. Thus, they propose to use local latent variables along with a dependency structure among them, resulting in new functional NPs. Lee et al. (2020) point out the limitation of using a single Gaussian latent variable to model functional uncertainty. To solve the limitation, they propose to use bootstrapping for inducing functional uncertainty, leading to a new NP variant, called Bootstrapping Neural Processes (BNP). Bruinsma et al. (2021) introduce a novel member of the NP family that incorporates translation equivariance and models the predictive distributions directly with Gaussian processes, called Gaussian NP (GNP). GNPs do not allow for correlations in the predictive distribution, but also provide universal approximation guarantees. Currently, NPs and their variants have been widely used in many different settings, including meta-learning (Singh et al., 2019; Yoon et al., 2020; Requeima et al., 2019) and sequential data modelling (Qin et al., 2019), but they have not been studied in SSL. Our work is the first to leverage NPs for semi-supervised large-scale image recognition. We choose the most basic model from (Garnelo et al., 2018b) (i.e., the original NPs), and we expect that future works can further study the application of other variants to this task.

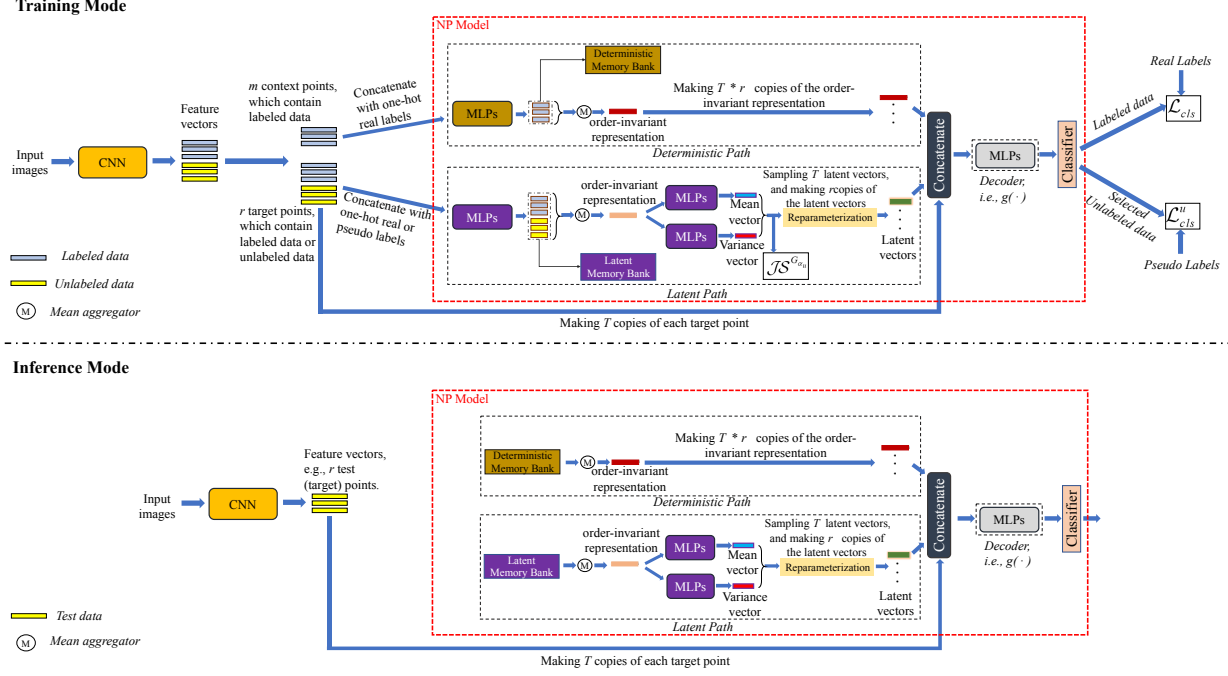


Figure 1. Overview of NP-Match: it contains a convolutional neural network (CNN) and an NP model that is shown in the red dotted box. The feature vectors come from the global average pooling layer in the CNN.

3. Methodology

In this section, we provide a brief introduction to neural processes (NPs) and a detailed description of our NP-Match.

3.1. NPs

NPs approximate stochastic processes via finite-dimensional marginal distributions (Garnelo et al., 2018b). Formally, given a probability space (Ω, Σ, Π) and an index set \mathcal{X} , a stochastic process can be written as $\{F(x, \omega) : x \in \mathcal{X}\}$, where $F(\cdot, \omega)$ is a sample function mapping \mathcal{X} to another space \mathcal{Y} for any point $\omega \in \Omega$. For each finite sequence $x_{1:n}$, a marginal joint distribution function can be defined on the function values $F(x_1, \omega), F(x_2, \omega), \dots, F(x_n, \omega)$, which satisfies two conditions given by the Kolmogorov Extension Theorem (Øksendal, 2003), namely, exchangeability and consistency. Assuming that a density π , where $d\Pi = \pi d\mu$, and the likelihood density $p(y_{1:n}|F(\cdot, \omega), x_{1:n})$ exist, the marginal joint distribution function can be written as:

$$p(y_{1:n}|x_{1:n}) = \int \pi(\omega) p(y_{1:n}|F(\cdot, \omega), x_{1:n}) d\mu(\omega). \quad (1)$$

The exchangeability condition requires the joint distributions to be invariant to permutations of the elements, i.e., $p(y_{1:n}|x_{1:n}) = p(\varphi(y_{1:n})|\varphi(x_{1:n}))$, where φ is a permutation of $\{1, \dots, n\}$. The consistency condition expresses that if a part of the sequence is marginalised out, then the resulting marginal distribution is consistent with that defined on

the original sequence.

Letting (Ω, Σ) be $(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$, where $\mathcal{B}(\mathbb{R}^d)$ denotes the Borel σ -algebra of \mathbb{R}^d , NPs parameterize the function $F(\cdot, \omega)$ with a high-dimensional random vector z sampled from a multivariate Gaussian distribution. Then, $F(x_i, \omega)$ can be replaced by $g(x_i, z)$, where $g(\cdot)$ denotes a neural network, and Eq. (1) becomes:

$$p(y_{1:n}|x_{1:n}) = \int \pi(z) p(y_{1:n}|g(x_{1:n}, z), x_{1:n}) d\mu(z). \quad (2)$$

The training objective of NPs is to maximize $p(y_{1:n}|x_{1:n})$, and the learning procedure reflects the NPs' property that they have the capability to make predictions for target points conditioned on context points (Garnelo et al., 2018b).

3.2. NP-Match

As Figure 1 shows, NP-Match is mainly composed of two parts: a deep neural network and an NP model. The deep neural network is leveraged for obtaining feature representations of input images, while the NP model is built upon the network to receive the representations for classification.

3.2.1. NP MODEL FOR SEMI-SUPERVISED IMAGE CLASSIFICATION

Since we extend the original NPs (Garnelo et al., 2018b) to the classification task, $p(y_{1:n}|g(x_{1:n}, z), x_{1:n})$ in Eq. (2) should define a categorical distribution rather than a Gaus-

sian distribution. Therefore, we parameterize the categorical distribution by probability vectors from a classifier that contains a weight matrix (\mathcal{W}) and a softmax function (Φ):

$$p(y_{1:n}|g(x_{1:n}, z), x_{1:n}) = \text{Categorical}(\Phi(\mathcal{W}g(x_{1:n}, z))). \quad (3)$$

Note that $g(\cdot)$ can be learned via amortised variational inference, and to use this method, two steps need to be done: (1) parameterize a variational distribution over z , and (2) find the evidence lower bound (ELBO) as the learning objective. For the first step, we let $q(z|x_{1:n}, y_{1:n})$ be a variational distribution defined on the same measure space, which can be parameterized by a neural network. For the second step, given a finite sequence with length n , we assume that there are m context points ($x_{1:m}$) and r target points ($x_{m+1:m+r}$) in it, i.e., $m+r=n$. Then, the ELBO is given by (with proof in the appendix):

$$\begin{aligned} \log p(y_{1:n}|x_{1:n}) \geq \\ \mathbb{E}_{q(z|x_{m+1:m+r}, y_{m+1:m+r})} \left[\sum_{i=m+1}^{m+r} \log p(y_i|z, x_i) - \right. \\ \left. \log \frac{q(z|x_{m+1:m+r}, y_{m+1:m+r})}{q(z|x_{1:m}, y_{1:m})} \right] + \text{const}. \end{aligned} \quad (4)$$

To learn the NP model, one can maximize this ELBO. Under the setting of SSL, we consider that only labeled data can be treated as context points, and either labeled or unlabeled data can be treated as target points, since the target points are what the NP model makes predictions for.

3.2.2. NP-MATCH PIPELINE

We now introduce the NP-Match pipeline. We first focus on the configuration of the NP model, which is shown in the red dotted box in Figure 1. The NP model is mainly constructed by MLPs, memory banks, and a classifier. Specifically, the classifier is composed of the weight matrix (\mathcal{W}) and the softmax function (Φ). Similarly to the original implementation of NPs, we build two paths with the memory banks and MLPs, namely, the latent path and the deterministic path. The decoder $g(\cdot)$ is also implemented with MLPs. The workflow of NP-Match at the training stage and the inference stage are different, which are shown in Figure 1, and they are introduced separately as follows.

Training mode. Given a batch of B labeled images $\mathcal{L} = \{(x_i, y_i) : i \in \{1, \dots, B\}\}$ and a batch of unlabeled images $\mathcal{U} = \{x_i^u : i \in \{1, \dots, \mu B\}\}$ at each iteration, where μ determines the relative size of \mathcal{U} to \mathcal{L} , we apply weak augmentation (i.e., crop-and-flip) on the labeled and unlabeled samples, and strong augmentation (i.e., RandAugment (Cubuk et al., 2020)) on only the unlabeled samples. After the augmentation is applied, the images are passed through the deep neural network, and the features are input to the NP model, which finally outputs the predictions and associated uncertainties. The detailed process can be summarized as

follows. At the start of each iteration, NP-Match is switched to inference mode, and it makes predictions for the weakly-augmented unlabeled data. Then, inference mode is turned off, and those predictions are treated as pseudo-labels for unlabeled data. After receiving the features, real labels, and pseudo-labels, the NP model first duplicates the labeled samples and treats them as context points, and all the labeled and unlabeled samples in the original batches are then treated as target points, since the NP model needs to make a prediction for them. Thereafter, the target points and context points are separately fed to the latent path and the deterministic path. As for the latent path, target points are concatenated with their corresponding real labels or pseudo labels, and processed by MLPs to get new representations. Then, the representations are averaged by a mean aggregator along the batch dimension, leading to an order-invariant representation, which implements the exchangeability and the consistency condition, and they are simultaneously stored in the latent memory bank, which is updated with a first-in-first-out strategy. After the mean aggregator, the order-invariant representation is further processed by other two MLPs in order to get the mean vector and the variance vector, which are used for sampling latent vectors via the reparameterization trick, and the number of latent vectors sampled at each feed-forward pass is denoted T . As for the deterministic path, context points are input to this path and are processed in the same way as the target points, until an order-invariant representation is procured from the mean aggregator. We also introduce a memory bank to the deterministic path for storing representations. Subsequently, each target point is concatenated with the T latent vectors and the order-invariant representations from the deterministic path (note that, practically, the target point and the order-invariant representations from the deterministic path must be copied T times). After the concatenation operation, the $T * r$ feature representations are fed into the decoder $g(\cdot)$ and then the classifier, which outputs T probability distributions over classes for each target point. The final prediction for each target point can be obtained by averaging the T predictions, and the uncertainty is computed as the entropy of the average prediction (Kendall & Gal, 2017). The ELBO (Eq. (4)) shows the learning objective. Specifically, the first term can be achieved by using the cross-entropy loss on the labeled and unlabeled data with their corresponding real labels and pseudo-labels, while the second term is the KL divergence between $q(z|x_{m+1:m+r}, y_{m+1:m+r})$ and $q(z|x_{1:m}, y_{1:m})$.

Inference mode. Concerning a set of test images, they are also passed through the deep neural network at first to obtain their feature representations. Then, they are treated as target points and are fed to the NP model. Since the labels of test data are not available, it is impossible to obtain the order-invariant representation from test data. In this case, the stored features in the two memory banks can be directly

used. As the bottom diagram of Figure 1 shows, after the order-invariant representations are obtained from the memory banks, the target points are leveraged in the same way as in the training mode to generate concatenated feature representations for the decoder $g(\cdot)$ and then the classifier.

3.2.3. UNCERTAINTY-GUIDED SKEW-GEOMETRIC JS DIVERGENCE

NP-Match, like many SSL approaches, relies on the use of pseudo-labels for the unlabeled samples. Pseudo-labels, however, are sometimes inaccurate and can lead to the neural network learning poor feature representations. In our pipeline, this can go on to impact the representation procured from the mean-aggregator and hence the model’s estimated mean vector, variance vector, and global latent vectors (see “Latent Path” in Figure 1). To remedy this, similarly to how the KL divergence term in the ELBO (Eq. (4)) is used to learn global latent variables (Garnelo et al., 2018b), we propose a new distribution divergence, called the uncertainty-guided skew-geometric JS divergence ($JS^{G_{\alpha_u}}$). We first formalize the definition of $JS^{G_{\alpha_u}}$:

Definition 1. Let (Ω, Σ) be a measurable space, where Ω denotes the sample space, and Σ denotes the σ -algebra of measurable events. P and Q are two probability measures defined on the measurable space. Concerning a positive measure¹, which is denoted as μ , the uncertainty-guided skew-geometric JS divergence ($JS^{G_{\alpha_u}}$) can be defined as:

$$JS^{G_{\alpha_u}}(p, q) = (1 - \alpha_u) \int p \log \frac{p}{G(p, q)_{\alpha_u}} d\mu + \alpha_u \int q \log \frac{q}{G(p, q)_{\alpha_u}} d\mu, \quad (5)$$

where p and q are the Radon-Nikodym derivatives of P and Q with respect to μ , the scalar $\alpha_u \in [0, 1]$ is calculated based on the uncertainty, and $G(p, q)_{\alpha_u} = p^{1-\alpha_u} q^{\alpha_u} / (\int_{\Omega} p^{1-\alpha_u} q^{\alpha_u} d\mu)$. The dual form of $JS^{G_{\alpha_u}}$ is given by:

$$JS_*^{G_{\alpha_u}}(p, q) = (1 - \alpha_u) \int G(p, q)_{\alpha_u} \log \frac{G(p, q)_{\alpha_u}}{p} d\mu + \alpha_u \int G(p, q)_{\alpha_u} \log \frac{G(p, q)_{\alpha_u}}{q} d\mu. \quad (6)$$

The proposed $JS^{G_{\alpha_u}}$ is an extension of the skew-geometric JS divergence first proposed by Nielsen (2020). Specifically, Nielsen (2020) generalizes the JS divergence with abstract means (quasi-arithmetic means (Niculescu & Persson, 2006)), in which a scalar α is defined to control the degree of divergence skew.² By selecting the weighted geometric mean $p^{1-\alpha} q^{\alpha}$, such generalized JS divergence becomes the

¹Specifically, the positive measure is usually the Lebesgue measure with the Borel σ -algebra $\mathcal{B}(\mathbb{R}^d)$ or the counting measure with the power set σ -algebra 2^{Ω} .

²The divergence skew means how closely related the intermediate distribution (the abstract mean of p and q) is to p or q .

skew-geometric JS divergence, which can be easily applied to the Gaussian distribution because of its property that the weighted product of exponential family distributions stays in the exponential family (Nielsen & Garcia, 2009). Our $JS^{G_{\alpha_u}}$ extends such divergence by incorporating the uncertainty into the scalar α to dynamically adjust the divergence skew. We assume the real variational distribution of the global latent variable under the supervised learning to be q^* . If the framework is trained with real labels, the condition $q(z|x_{m+1:m+r}, y_{m+1:m+r}) = q(z|x_{1:m}, y_{1:m}) = q^*$ will hold after training, since they are all the marginal distributions of the same stochastic process. However, as for SSL, $q(z|x_{m+1:m+r}, y_{m+1:m+r})$ and $q(z|x_{1:m}, y_{1:m})$ are no longer equal to q^* , as some low-quality representations are involved during training, which affect the estimation of $q(z|x_{m+1:m+r}, y_{m+1:m+r})$ and $q(z|x_{1:m}, y_{1:m})$. Our proposed $JS^{G_{\alpha_u}}$ solves this issue by introducing an intermediate distribution that is calculated via $G(q(z|x_{1:m}, y_{1:m}), q(z|x_{m+1:m+r}, y_{m+1:m+r}))_{\alpha_u}$, where $\alpha_u = u_{c_{avg}} / (u_{c_{avg}} + u_{t_{avg}})$. Here, $u_{c_{avg}}$ denotes the average value over the uncertainties of the predictions of context points, and $u_{t_{avg}}$ represents the average value over that of target points. With this setting, the intermediate distribution is usually close to q^* . For example, when $u_{c_{avg}}$ is large, and $u_{t_{avg}}$ is small, which means that there are many low-quality feature presentations involved for calculating $q(z|x_{1:m}, y_{1:m})$, and $q(z|x_{m+1:m+r}, y_{m+1:m+r})$ is closer to q^* , then $G(q(z|x_{1:m}, y_{1:m}), q(z|x_{m+1:m+r}, y_{m+1:m+r}))_{\alpha_u}$ will be close to $q(z|x_{m+1:m+r}, y_{m+1:m+r})$, and as a result, the network is optimized to learn the distribution of the global latent variable in the direction to q^* , which mitigates the issue to some extent.³ Concerning the variational distribution being supposed to be a Gaussian distribution, we introduce the following theorem (with proof in the appendix) for calculating $JS^{G_{\alpha_u}}$ on Gaussian distributions:

Theorem 1. Given two multivariate Gaussians $\mathcal{N}_1(\mu_1, \Sigma_1)$ and $\mathcal{N}_2(\mu_2, \Sigma_2)$, the following holds:

$$\begin{aligned} JS^{G_{\alpha_u}}(\mathcal{N}_1, \mathcal{N}_2) &= \frac{1}{2} (\text{tr}(\Sigma_{\alpha_u}^{-1}((1 - \alpha_u)\Sigma_1 + \alpha_u\Sigma_2)) + \\ &(1 - \alpha_u)(\mu_{\alpha_u} - \mu_1)^T \Sigma_{\alpha_u}^{-1}(\mu_{\alpha_u} - \mu_1) + \\ &\alpha_u(\mu_{\alpha_u} - \mu_2)^T \Sigma_{\alpha_u}^{-1}(\mu_{\alpha_u} - \mu_2) + \\ &\log[\frac{\det[\Sigma_{\alpha_u}]}{\det[\Sigma_1]^{1-\alpha_u} \det[\Sigma_2]^{\alpha_u}}] - D) \quad (7) \\ JS_*^{G_{\alpha_u}}(\mathcal{N}_1, \mathcal{N}_2) &= \\ &\frac{1}{2} (\log[\frac{\det[\Sigma_1]^{1-\alpha_u} \det[\Sigma_2]^{\alpha_u}}{\det[\Sigma_{\alpha_u}]}] + \alpha_u \mu_2^T \Sigma_2^{-1} \mu_2 \\ &- \mu_{\alpha_u}^T \Sigma_{\alpha_u}^{-1} \mu_{\alpha_u} + (1 - \alpha_u) \mu_1^T \Sigma_1^{-1} \mu_1), \end{aligned}$$

³As long as one of $q(z|x_{1:m}, y_{1:m})$ and $q(z|x_{m+1:m+r}, y_{m+1:m+r})$ is close to q^* , the proposed $JS^{G_{\alpha_u}}$ mitigates the issue, but $JS^{G_{\alpha_u}}$ still has difficulties to solve the problem when both of their calculations involve many low-quality representations.

where $\Sigma_{\alpha_u} = ((1 - \alpha_u)\Sigma_1^{-1} + \alpha_u\Sigma_2^{-1})^{-1}$ and $\mu_{\alpha_u} = \Sigma_{\alpha_u}((1 - \alpha_u)\Sigma_1^{-1}\mu_1 + \alpha_u\Sigma_2^{-1}\mu_2)$, D denotes the number of dimension, and $\det[\cdot]$ represents the determinant.

With Theorem 1, one can calculate $JS^{G\alpha_u}$ or its dual form $JS_*^{G\alpha_u}$ based on the mean vector and the variance vector, and use $JS^{G\alpha_u}$ or $JS_*^{G\alpha_u}$ to replace the original KL divergence term in the ELBO (Eq. (4)) for training the whole framework. When the two distributions are diagonal Gaussians, Σ_1 and Σ_2 can be implemented by diagonal matrices with the variance vectors for calculating $JS^{G\alpha_u}$ or $JS_*^{G\alpha_u}$.

3.2.4. LOSS FUNCTIONS

To calculate loss functions, reliable pseudo-labels are required for unlabeled data. In practice, to select reliable unlabeled samples from \mathcal{U} and their corresponding pseudo-labels, we preset a confidence threshold (τ_c) and an uncertainty threshold (τ_u). In particular, as for unlabeled data x_i^u , NP-Match gives its prediction $p(y|Aug_w(x_i^u))$ and associated uncertainty estimate under the inference mode, where $Aug_w(\cdot)$ denotes the weak augmentation. When the highest prediction score $\max(p(y|Aug_w(x_i^u)))$ is higher than τ_c , and the uncertainty is smaller than τ_u , the sample will be chosen, and we denote the selected sample as x_i^{uc} , since the model is certain about his prediction, and the pseudo-label of x_i^{uc} is $\hat{y}_i = \arg \max(p(y|Aug_w(x_i^{uc})))$. Concerning μB unlabeled samples in \mathcal{U} , we assume B_c unlabeled samples are selected from them in each feedforward pass. According to the ELBO (Eq. (4)), three loss terms are used for training, namely, L_{cls} , L_{cls}^u , and $JS^{G\alpha_u}$. For each input (labeled or unlabeled), the NP model can give T predictions, and hence L_{cls} and L_{cls}^u are defined as:

$$L_{cls} = \frac{1}{B \times T} \sum_{i=1}^B \sum_{j=1}^T H(y_i^*, p_j(y|Aug_w(x_i))),$$

$$L_{cls}^u = \frac{1}{B_c \times T} \sum_{i=1}^{B_c} \sum_{j=1}^T H(\hat{y}_i, p_j(y|Aug_s(x_i^{uc}))),$$
(8)

where $Aug_s(\cdot)$ denotes the strong augmentation, y_i^* represents the real label for the labeled sample x_i , and $H(\cdot, \cdot)$ denotes the cross-entropy between two distributions. Thus, the total loss function is given by:

$$L_{total} = L_{cls} + \lambda_u L_{cls}^u + \beta JS^{G\alpha_u},$$
(9)

where λ_u and β are coefficients. During training, we followed previous work (Sohn et al., 2020; Zhang et al., 2021; Rizve et al., 2021; Li et al., 2021) to utilize the exponential moving average (EMA) technique. It is worth noting that, in the real implementation, NP-Match only preserves the averaged representation over all representations in each memory bank after training, which just takes up negligible storage space.

4. Experiments

We now report on our experiments of NP-Match on four public image classification benchmarks. To save space, implementation details are given in the appendix.

4.1. Datasets

We conducted our experiments on four widely used public SSL benchmarks, including CIFAR-10 (Krizhevsky et al., 2009), CIFAR-100 (Krizhevsky et al., 2009), STL-10 (Coates et al., 2011), and ImageNet (Deng et al., 2009). CIFAR-10 and CIFAR-100 contain 50,000 images of size 32×32 from 10 and 100 classes, respectively. We evaluated NP-match on these two datasets following the evaluation settings used in previous works (Sohn et al., 2020; Zhang et al., 2021; Li et al., 2021). The STL-10 dataset has 5000 labeled samples with size 96×96 from 10 classes and 100,000 unlabeled samples, and it is more difficult than CIFAR, since STL-10 has a number of out-of-distribution images in the unlabeled set. We follow the experimental settings for STL-10 as detailed in (Zhang et al., 2021). Finally, ImageNet contains around 1.2 million images from 1000 classes. Following the experimental settings in (Zhang et al., 2021), we used 100K labeled data, namely, 100 labels per class.

4.2. Main Results

In the following, we report the main experimental results on the accuracy, the average uncertainty, the expected uncertainty calibration error, and the running time of NP-Match compared with SOTA approaches.

First, in Table 1, we compare NP-Match with SOTA SSL image classification methods on CIFAR-10, CIFAR-100, and STL-10. We see that NP-Match outperforms SOTA results or achieves competitive results under different SSL settings. We highlight two key observations. First, NP-Match outperforms all other methods by a wide margin on all three benchmarks under the most challenging settings, where the number of labeled samples is smallest. Second, NP-Match is compared to UPS⁴, since the UPS framework is the MC-dropout-based probabilistic model for semi-supervised image classification, and NP-Match completely outperforms them on all three benchmarks. This suggests that NPs can be a good alternative to MC dropout in probabilistic approaches to semi-supervised learning tasks.

Second, we analyse the relationship between the average class-wise uncertainty and accuracy at test phase on CIFAR-10 and STL10. From Figure 2, we empirically observe that: (1) when more labeled data are used for training, the average uncertainty of samples' predictions for each class

⁴Note that UPS (Rizve et al., 2021) does not use strong augmentations, thus we re-implemented it with RandAugment (Cubuk et al., 2020) for fair comparisons.

NP-Match: When Neural Processes meet Semi-Supervised Learning

Dataset	CIFAR-10			CIFAR-100			STL-10			
	Label Amount	40	250	4000	400	2500	10000	40	250	1000
MixMatch (Berthelot et al., 2019)		36.19 (± 6.48)	13.63 (± 0.59)	6.66 (± 0.26)	67.59 (± 0.66)	39.76 (± 0.48)	27.78 (± 0.29)	54.93 (± 0.96)	34.52 (± 0.32)	21.70 (± 0.68)
ReMixMatch (Berthelot et al., 2020)		9.88 (± 1.03)	6.30 (± 0.05)	4.84 (± 0.01)	42.75 (± 1.05)	26.03 (± 0.35)	20.02 (± 0.27)	32.12 (± 6.24)	12.49 (± 1.28)	6.74 (± 0.14)
UDA (Xie et al., 2020)		10.62 (± 3.75)	5.16 (± 0.06)	4.29 (± 0.07)	46.39 (± 1.59)	27.73 (± 0.21)	22.49 (± 0.23)	37.42 (± 8.44)	9.72 (± 1.15)	6.64 (± 0.17)
CoMatch (Li et al., 2021)		6.88 (± 0.92)	4.90 (± 0.35)	4.06 (± 0.03)	40.02 (± 1.11)	27.01 (± 0.21)	21.83 (± 0.23)	31.77 (± 2.56)	11.56 (± 1.27)	8.66 (± 0.41)
SemCo (Nassar et al., 2021)		7.87 (± 0.22)	5.12 (± 0.27)	3.80 (± 0.08)	44.11 (± 1.18)	31.93 (± 0.33)	24.45 (± 0.12)	34.17 (± 2.78)	12.23 (± 1.40)	7.49 (± 0.29)
Meta Pseudo Labels (Pham et al., 2021)		6.93 (± 0.17)	4.94 (± 0.04)	3.89 (± 0.07)	44.23 (± 0.99)	27.68 (± 0.22)	22.48 (± 0.18)	34.29 (± 3.29)	9.90 (± 0.96)	6.45 (± 0.26)
FlexMatch (Zhang et al., 2021)		4.96 (± 0.06)	4.98 (± 0.09)	4.19 (± 0.01)	39.94 (± 1.62)	26.49 (± 0.20)	21.90 (± 0.15)	29.15 (± 4.16)	8.23 (± 0.39)	5.77 (± 0.18)
UPS (Rizve et al., 2021)		5.26 (± 0.29)	5.11 (± 0.08)	4.25 (± 0.05)	41.07 (± 1.66)	27.14 (± 0.24)	21.97 (± 0.23)	30.82 (± 2.16)	9.77 (± 0.44)	6.02 (± 0.28)
FixMatch (Sohn et al., 2020)		7.47 (± 0.28)	4.86 (± 0.05)	4.21 (± 0.08)	46.42 (± 0.82)	28.03 (± 0.16)	22.20 (± 0.12)	35.96 (± 4.14)	9.81 (± 1.04)	6.25 (± 0.33)
NP-Match (ours)		4.91 (± 0.04)	4.96 (± 0.06)	4.11 (± 0.02)	38.91 (± 0.99)	26.03 (± 0.26)	21.22 (± 0.13)	14.20 (± 0.67)	9.51 (± 0.37)	5.59 (± 0.24)

Table 1. Comparison with SOTA results on CIFAR-10, CIFAR-100, and STL-10. The error rates are reported with standard deviation.




Figure 2. Analysis of average class-wise uncertainty and accuracy.

Dataset	CIFAR-10			STL-10			
	Label Amount	40	250	4000	40	250	1000
UPS (MC Dropout)		7.96	7.02	5.82	17.23	9.65	5.69
NP-Match		7.23	6.85	5.89	12.45	8.72	5.28

Table 2. Expected UCEs (%) of the MC-dropout-based model (i.e., UPS (Rizve et al., 2021)) and of NP-Match on the test sets of CIFAR-10 and STL-10.

	Method	Top-1	Top-5
		Deterministic Methods	FixMatch (Sohn et al., 2020)
	FlexMatch (Zhang et al., 2021)	41.85	19.48
	CoMatch (Li et al., 2021)	42.17	19.64
Probabilistic Methods	UPS (Rizve et al., 2021)	42.69	20.23
	NP-Match	41.78	19.33

Table 3. Error rates of SOTA methods on ImageNet.




Figure 3. Time consumption of estimating uncertainty for the MC-dropout-based model (i.e., UPS (Rizve et al., 2021)) and NP-Match. The horizontal axis refers to the number of predictions used for the uncertainty quantification, and the vertical axis indicates the time consumption (sec).

decreases. This is consistent with the property of NPs and GPs where the model is less uncertain with regard to its prediction when more real and correct labels are leveraged; (2) the classes with higher average uncertainties have lower accuracy, meaning that the uncertainty is a good standard for choosing unlabeled samples.

Third, the expected uncertainty calibration error (UCE) of our method is also calculated to evaluate the uncertainty

estimation. The expected UCE is used to measure the miscalibration of uncertainty (Laves et al., 2020), which is an analogue to the expected calibration error (ECE) (Guo et al., 2017; Naeini et al., 2015). The low expected UCE indicates that the model is certain when making accurate predictions and that the model is uncertain when making inaccurate predictions. More details about the expected UCE can be found in previous works (Laves et al., 2020; Krishnan & Tickoo, 2020). The results of NP-Match and the MC-dropout-based model (i.e., UPS (Rizve et al., 2021)) are shown in Table 2; their comparison shows that NP-Match can output more reliable and well-calibrated uncertainty estimates.

Furthermore, we compare the running time of NP-Match and the MC dropout-based model (i.e., UPS (Rizve et al., 2021)). We use a batch of 16 samples and two network architectures that are widely used in previous works (Zagoruyko & Komodakis, 2016; Zhang et al., 2021; Sohn et al., 2020; Li et al., 2021), namely, WRN-28-2 on CIFAR-10 (Figure 3 (a)) and WRN-28-8 on CIFAR-100 (Figure 3 (b)). In (a), we observe that when the number of predictions (T) increases, the time cost of the UPS framework rises quickly, but the time cost of NP-Match grows slowly. In (b), we observe that the time cost gap between these two methods is even

Dataset	CIFAR-10			CIFAR-100			STL-10		
	Label Amount	40	250	4000	400	2500	10000	40	250
NP-Match with KL	5.32 (± 0.06)	5.20 (± 0.02)	4.36 (± 0.03)	39.15 (± 1.53)	26.48 (± 0.23)	21.51 (± 0.17)	14.67 (± 0.38)	9.92 (± 0.24)	6.21 (± 0.23)
NP-Match with $JS_*^{G_{\alpha u}}$	4.93 (± 0.02)	4.87 (± 0.03)	4.19 (± 0.04)	38.67 (± 1.29)	26.24 (± 0.17)	21.33 (± 0.10)	14.45 (± 0.55)	9.48 (± 0.28)	5.47 (± 0.19)
NP-Match with $JS^{G_{\alpha u}}$	4.91 (± 0.04)	4.96 (± 0.06)	4.11 (± 0.02)	38.91 (± 0.99)	26.03 (± 0.26)	21.22 (± 0.13)	14.20 (± 0.67)	9.51 (± 0.37)	5.59 (± 0.24)

Table 4. Ablation studies of the proposed uncertainty-guided skew-geometric JS divergence and its dual form.

larger when a larger model is tested on a larger dataset. This demonstrates that NP-Match is significantly more computationally efficient than MC dropout-based methods.

Finally, Table 3 shows the experiments conducted on ImageNet. Here, NP-Match achieves a SOTA performance, suggesting that it is effective at handling challenging large-scale datasets. Note that previous works usually evaluate their frameworks under distinct SSL settings, and thus it is hard to compare different methods directly. Therefore, we re-evaluate another two methods proposed recently under the same SSL setting with the same training details, namely, UPS and CoMatch.

4.3. Ablation Studies

We now report our ablation studies on CIFAR-10, CIFAR-100, and STL-10. We also present further experiments, including a hyperparameter exploration, in the appendix.

We evaluate our uncertainty-guided skew-geometric JS divergence ($JS^{G_{\alpha u}}$) as well as its dual form ($JS_*^{G_{\alpha u}}$), and compare them to the original KL divergence in NPs. In Table 4, we see that NP-Match with KL divergence consistently underperforms relative to our proposed $JS^{G_{\alpha u}}$ and $JS_*^{G_{\alpha u}}$. This suggests that our uncertainty-guided skew-geometric JS divergence can mitigate the problem caused by low-quality feature representations. Between the two, $JS^{G_{\alpha u}}$ and $JS_*^{G_{\alpha u}}$ achieve a comparable performance across the three benchmarks, and thus we select $JS^{G_{\alpha u}}$ to replace the original KL divergence in the ELBO (Eq. (4)) for the comparisons to previous SOTA methods in Section 4.2.

5. Summary and Outlook

In this work, we proposed the application of neural processes (NPs) to semi-supervised learning (SSL), designing a new framework called NP-Match, and explored its use in semi-supervised large-scale image classification. To our knowledge, this is the first such work. To better adapt NP-Match to the SSL task, we proposed a new divergence term, which we call uncertainty-guided skew-geometric JS divergence, to replace the original KL divergence in NPs. We demonstrated the effectiveness of NP-Match and the proposed divergence term for SSL in extensive experiments, and also showed that NP-Match could be a good alternative to MC dropout in SSL.

Future works will explore the following two directions. First, due to the successful application of NPs to semi-supervised image classification, it is valuable to explore NPs in other SSL tasks, such as object detection and segmentation. Second, many successful NPs variants have been proposed since the original NPs (Garnelo et al., 2018b) (see Section 2). We will also explore these in SSL for image classification.

6. Acknowledgements

This work was partially supported by the Alan Turing Institute under the EPSRC grant EP/N510129/1, by the AXA Research Fund, and by the EPSRC grant EP/R013667/1. We also acknowledge the use of the EPSRC-funded Tier 2 facility JADE (EP/P020275/1) and GPU computing support by Scan Computers International Ltd.

References

Bachman, P., Alsharif, O., and Precup, D. Learning with pseudo-ensembles. *Advances in Neural Information Processing Systems*, 27:3365–3373, 2014.

Berthelot, D., Carlini, N., Goodfellow, I., Papernot, N., Oliver, A., and Raffel, C. MixMatch: A holistic approach to semi-supervised learning. *Advances in Neural Information Processing Systems*, 2019.

Berthelot, D., Carlini, N., Cubuk, E. D., Kurakin, A., Sohn, K., Zhang, H., and Raffel, C. ReMixMatch: Semi-supervised learning with distribution alignment and augmentation anchoring. *International Conference on Learning Representations*, 2020.

Bruinsma, W. P., Requeima, J., Foong, A. Y., Gordon, J., and Turner, R. E. The Gaussian neural process. *Advances in Approximate Bayesian Inference*, 2021.

Chen, D.-D., Wang, W., Gao, W., and Zhou, Z.-H. Tri-net for semi-supervised deep learning. In *Proceedings of 27th International Joint Conference on Artificial Intelligence*, pp. 2014–2020, 2018.

Coates, A., Ng, A., and Lee, H. An analysis of single-layer networks in unsupervised feature learning. In *International Conference on Artificial Intelligence and Statistics*, pp. 215–223. JMLR Workshop and Conference Proceedings, 2011.

- Cubuk, E. D., Zoph, B., Shlens, J., and Le, Q. V. RandAugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pp. 702–703, 2020.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. ImageNet: A large-scale hierarchical image database. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 248–255, 2009.
- Gal, Y. and Ghahramani, Z. Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In *International Conference on Machine Learning*, pp. 1050–1059. PMLR, 2016.
- Garnelo, M., Rosenbaum, D., Maddison, C., Ramalho, T., Saxton, D., Shanahan, M., Teh, Y. W., Rezende, D., and Eslami, S. A. Conditional neural processes. In *International Conference on Machine Learning*, pp. 1704–1713. PMLR, 2018a.
- Garnelo, M., Schwarz, J., Rosenbaum, D., Viola, F., Rezende, D. J., Eslami, S., and Teh, Y. W. Neural processes. *arXiv:1807.01622*, 2018b.
- Gordon, J., Bruinsma, W. P., Foong, A. Y., Requeima, J., Dubois, Y., and Turner, R. E. Convolutional conditional neural processes. *International Conference on Learning Representations*, 2020.
- Guo, C., Pleiss, G., Sun, Y., and Weinberger, K. Q. On calibration of modern neural networks. In *International Conference on Machine Learning*, pp. 1321–1330. PMLR, 2017.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 770–778, 2016.
- Hu, Z., Yang, Z., Hu, X., and Nevatia, R. SimPLE: Similar pseudo label exploitation for semi-supervised classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 15099–15108, 2021.
- Jean, N., Xie, S. M., and Ermon, S. Semi-supervised deep kernel learning: Regression with unlabeled data by minimizing predictive variance. *Advances in Neural Information Processing Systems*, 2018.
- Kendall, A. and Gal, Y. What uncertainties do we need in bayesian deep learning for computer vision? *Advances in Neural Information Processing Systems*, 2017.
- Kim, H., Mnih, A., Schwarz, J., Garnelo, M., Eslami, A., Rosenbaum, D., Vinyals, O., and Teh, Y. W. Attentive neural processes. *International Conference on Learning Representations*, 2019.
- Krishnan, R. and Tickoo, O. Improving model calibration with accuracy versus uncertainty optimization. *Advances in Neural Information Processing Systems*, 2020.
- Krizhevsky, A., Hinton, G., et al. Learning multiple layers of features from tiny images. 2009.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. ImageNet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems*, 25: 1097–1105, 2012.
- Laine, S. and Aila, T. Temporal ensembling for semi-supervised learning. *International Conference on Learning Representations*, 2017.
- Laves, M.-H., Ihler, S., Kortmann, K.-P., and Ortmaier, T. Calibration of model uncertainty for dropout variational inference. *arXiv:2006.11584*, 2020.
- Lee, D.-H. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on Challenges in Representation Learning, International Conference on Machine Learning*, volume 3, pp. 896, 2013.
- Lee, J., Lee, Y., Kim, J., Yang, E., Hwang, S. J., and Teh, Y. W. Bootstrapping neural processes. *Advances in Neural Information Processing Systems*, pp. 6606–6615, 2020.
- Li, J., Xiong, C., and Hoi, S. C. CoMatch: Semi-supervised learning with contrastive graph regularization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 9475–9484, 2021.
- Li, Y.-F. and Zhou, Z.-H. Towards making unlabeled data never hurt. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(1):175–188, 2014.
- Liu, W., He, J., and Chang, S.-F. Large graph construction for scalable semi-supervised learning. In *International Conference on Machine Learning*, 2010.
- Liu, Z.-Y., Li, S.-Y., Chen, S., Hu, Y., and Huang, S.-J. Uncertainty aware graph Gaussian process for semi-supervised learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 4957–4964, 2020.
- Loshchilov, I. and Hutter, F. SGDR: Stochastic gradient descent with warm restarts. *arXiv:1608.03983*, 2016.
- Louizos, C., Shi, X., Schutte, K., and Welling, M. The functional neural process. *Advances in Neural Information Processing Systems*, 2019.

- Naeini, M. P., Cooper, G., and Hauskrecht, M. Obtaining well calibrated probabilities using bayesian binning. In *29th AAAI Conference on Artificial Intelligence*, 2015.
- Nassar, I., Herath, S., Abbasnejad, E., Buntine, W., and Haffari, G. All labels are not created equal: Enhancing semi-supervision via label grouping and co-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7241–7250, 2021.
- Ng, Y. C., Colombo, N., and Silva, R. Bayesian semi-supervised learning with graph Gaussian processes. *Advances in Neural Information Processing Systems*, 2018.
- Niculescu, C. and Persson, L.-E. *Convex functions and their applications*. Springer, 2006.
- Nielsen, F. On a generalization of the Jensen-Shannon divergence and the Jensen-Shannon centroid. *Entropy*, 22(2):221, 2020.
- Nielsen, F. and Garcia, V. Statistical exponential families: A digest with flash cards. *arXiv:0911.4863*, 2009.
- Øksendal, B. Stochastic differential equations. In *Stochastic Differential Equations*, pp. 65–84. Springer, 2003.
- Pham, H., Dai, Z., Xie, Q., and Le, Q. V. Meta pseudo labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11557–11568, 2021.
- Qiao, S., Shen, W., Zhang, Z., Wang, B., and Yuille, A. Deep co-training for semi-supervised image recognition. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 135–152, 2018.
- Qin, S., Zhu, J., Qin, J., Wang, W., and Zhao, D. Recurrent attentive neural process for sequential data. *LIRE Workshop NeurIPS*, 2019.
- Requeima, J., Gordon, J., Bronskill, J., Nowozin, S., and Turner, R. E. Fast and flexible multi-task classification using conditional neural adaptive processes. *Advances in Neural Information Processing Systems*, 32:7959–7970, 2019.
- Rizve, M. N., Duarte, K., Rawat, Y. S., and Shah, M. In defense of pseudo-labeling: An uncertainty-aware pseudo-label selection framework for semi-supervised learning. *International Conference on Learning Representations*, 2021.
- Ronneberger, O., Fischer, P., and Brox, T. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 234–241. Springer, 2015.
- Sajjadi, M., Javanmardi, M., and Tasdizen, T. Regularization with stochastic transformations and perturbations for deep semi-supervised learning. *Advances in Neural Information Processing Systems*, 29:1163–1171, 2016.
- Sedai, S., Antony, B., Rai, R., Jones, K., Ishikawa, H., Schuman, J., Gadi, W., and Garnavi, R. Uncertainty guided semi-supervised segmentation of retinal layers in OCT images. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 282–290. Springer, 2019.
- Shi, Y., Zhang, J., Ling, T., Lu, J., Zheng, Y., Yu, Q., Qi, L., and Gao, Y. Inconsistency-aware uncertainty estimation for semi-supervised medical image segmentation. *IEEE Transactions on Medical Imaging*, 2021.
- Simonyan, K. and Zisserman, A. Very deep convolutional networks for large-scale image recognition. *International Conference on Learning Representations*, 2014.
- Sindhwani, V., Chu, W., and Keerthi, S. S. Semi-supervised Gaussian process classifiers. In *International Joint Conference on Artificial Intelligence*, pp. 1059–1064, 2007.
- Singh, G., Yoon, J., Son, Y., and Ahn, S. Sequential neural processes. *Advances in Neural Information Processing Systems*, 2019.
- Sohn, K., Berthelot, D., Li, C.-L., Zhang, Z., Carlini, N., Cubuk, E. D., Kurakin, A., Zhang, H., and Raffel, C. FixMatch: Simplifying semi-supervised learning with consistency and confidence. *Advances in Neural Information Processing Systems*, 2020.
- Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., and Rabinovich, A. Going deeper with convolutions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1–9, 2015.
- Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., and Wojna, Z. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2818–2826, 2016.
- Walker, I. and Glocker, B. Graph convolutional Gaussian processes. In *International Conference on Machine Learning*, pp. 6495–6504. PMLR, 2019.
- Wang, K., Zhan, B., Zu, C., Wu, X., Zhou, J., Zhou, L., and Wang, Y. Triple uncertainty guided mean teacher model for semi-supervised medical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 450–460. Springer, 2021.

- Wang, X., Kihara, D., Luo, J., and Qi, G.-J. Enaet: A self-trained framework for semi-supervised and supervised learning with ensemble transformations. *IEEE Transactions on Image Processing*, 30:1639–1647, 2020.
- Wilson, A. G., Hu, Z., Salakhutdinov, R., and Xing, E. P. Deep kernel learning. In *Artificial Intelligence and Statistics*, pp. 370–378. PMLR, 2016.
- Xie, Q., Dai, Z., Hovy, E., Luong, M.-T., and Le, Q. V. Unsupervised data augmentation for consistency training. *Advances in Neural Information Processing Systems*, 2020.
- Yasarla, R., Sindagi, V. A., and Patel, V. M. Syn2Real transfer learning for image deraining using Gaussian processes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2726–2736, 2020.
- Yoon, J., Singh, G., and Ahn, S. Robustifying sequential neural processes. In *International Conference on Machine Learning*, pp. 10861–10870. PMLR, 2020.
- Yu, L., Wang, S., Li, X., Fu, C.-W., and Heng, P.-A. Uncertainty-aware self-ensembling model for semi-supervised 3D left atrium segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 605–613. Springer, 2019.
- Zagoruyko, S. and Komodakis, N. Wide residual networks. *British Machine Vision Conference*, 2016.
- Zhai, X., Oliver, A., Kolesnikov, A., and Beyer, L. S4l: Self-supervised semi-supervised learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 1476–1485, 2019.
- Zhang, B., Wang, Y., Hou, W., Wu, H., Wang, J., Okumura, M., and Shinozaki, T. FlexMatch: Boosting semi-supervised learning with curriculum pseudo labeling. *Advances in Neural Information Processing Systems*, 34, 2021.
- Zhu, H., Li, Y., Bai, F., Chen, W., Li, X., Ma, J., Teo, C. S., Tao, P. Y., and Lin, W. Grasping detection network with uncertainty estimation for confidence-driven semi-supervised domain adaptation. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 9608–9613. IEEE, 2020.

Appendix

A. Derivation of ELBO (Eq. (4))

Proof. As for the marginal joint distribution $p(y_{1:n}|x_{1:n})$ over n data points in which there are m context points and r target points (i.e., $m + r = n$), we assume a variation distribution q , and then:

$$\begin{aligned}
\log p(y_{1:n}|x_{1:n}) &= \log \int_z p(z, y_{1:n}|x_{1:n}) \\
&= \log \int_z \frac{p(z, y_{1:n}|x_{1:n})}{q(z|x_{m+1:m+r}, y_{m+1:m+r})} q(z|x_{m+1:m+r}, y_{m+1:m+r}) \\
&\geq \mathbb{E}_{q(z|x_{m+1:m+r}, y_{m+1:m+r})} \left[\log \frac{p(z, y_{1:n}|x_{1:n})}{q(z|x_{m+1:m+r}, y_{m+1:m+r})} \right] \\
&= \mathbb{E}_{q(z|x_{m+1:m+r}, y_{m+1:m+r})} \left[\log \frac{p(y_{1:m})p(z|x_{1:m}, y_{1:m}) \prod_{i=m+1}^{m+r} p(y_i|z, x_i)}{q(z|x_{m+1:m+r}, y_{m+1:m+r})} \right] \\
&= \mathbb{E}_{q(z|x_{m+1:m+r}, y_{m+1:m+r})} \left[\sum_{i=m+1}^{m+r} \log p(y_i|z, x_i) + \log \frac{p(z|x_{1:m}, y_{1:m})}{q(z|x_{m+1:m+r}, y_{m+1:m+r})} + \log p(y_{1:m}) \right] \\
&= \mathbb{E}_{q(z|x_{m+1:m+r}, y_{m+1:m+r})} \left[\sum_{i=m+1}^{m+r} \log p(y_i|z, x_i) - \log \frac{q(z|x_{m+1:m+r}, y_{m+1:m+r})}{p(z|x_{1:m}, y_{1:m})} \right] + \text{const},
\end{aligned} \tag{10}$$

where “const” refers to $\mathbb{E}_{q(z|x_{m+1:m+r}, y_{m+1:m+r})} [\log p(y_{1:m})]$, which is a constant term. Concerning that $p(z|x_{1:m}, y_{1:m})$ is unknown, we replace it with $q(z|x_{1:m}, y_{1:m})$, and then we get:

$$\log p(y_{1:n}|x_{1:n}) \geq \mathbb{E}_{q(z|x_{m+1:m+r}, y_{m+1:m+r})} \left[\sum_{i=m+1}^{m+r} \log p(y_i|z, x_i) - \log \frac{q(z|x_{m+1:m+r}, y_{m+1:m+r})}{q(z|x_{1:m}, y_{1:m})} \right] + \text{const}. \tag{11}$$

□

B. Proof of Theorem 1

Proof. Let us first show that $\Sigma_{\alpha_u} = ((1 - \alpha_u)\Sigma_1^{-1} + \alpha_u\Sigma_2^{-1})^{-1}$ and $\mu_{\alpha_u} = \Sigma_{\alpha_u}((1 - \alpha_u)\Sigma_1^{-1}\mu_1 + \alpha_u\Sigma_2^{-1}\mu_2)$. Concerning two Gaussian distributions $\mathcal{N}_1(\mu_1, \Sigma_1)$ and $\mathcal{N}_2(\mu_2, \Sigma_2)$, the weighted geometric mean of them ($\mathcal{N}_1^{1-\alpha_u}\mathcal{N}_2^{\alpha_u}$) is given by:

$$\begin{aligned}
&(2\pi)^{-\frac{D}{2}} \det[\Sigma_1]^{-\frac{1-\alpha_u}{2}} \det[\Sigma_2]^{-\frac{\alpha_u}{2}} e^{-\frac{1-\alpha_u}{2}(x-\mu_1)^T \Sigma_1^{-1}(x-\mu_1) - \frac{\alpha_u}{2}(x-\mu_2)^T \Sigma_2^{-1}(x-\mu_2)} \\
&= (2\pi)^{-\frac{D}{2}} \det[\Sigma_1]^{-\frac{1-\alpha_u}{2}} \det[\Sigma_2]^{-\frac{\alpha_u}{2}} e^{-\frac{1}{2}((x-\mu_1)^T((1-\alpha_u)\Sigma_1^{-1})(x-\mu_1) + (x-\mu_2)^T(\alpha_u\Sigma_2^{-1})(x-\mu_2))}.
\end{aligned} \tag{12}$$

Now, we let $\Sigma_{1_u}^{-1} = (1 - \alpha_u)\Sigma_1^{-1}$ and $\Sigma_{2_u}^{-1} = \alpha_u\Sigma_2^{-1}$, then:

$$\begin{aligned}
&(2\pi)^{-\frac{D}{2}} \det[\Sigma_1]^{-\frac{1-\alpha_u}{2}} \det[\Sigma_2]^{-\frac{\alpha_u}{2}} e^{-\frac{1}{2}((x-\mu_1)^T \Sigma_{1_u}^{-1}(x-\mu_1) + (x-\mu_2)^T \Sigma_{2_u}^{-1}(x-\mu_2))} \\
&= C_1 e^{-\frac{1}{2}(x^T(\Sigma_{1_u}^{-1} + \Sigma_{2_u}^{-1})x - x^T(\Sigma_{1_u}^{-1}\mu_1 + \Sigma_{2_u}^{-1}\mu_2) - (\mu_1^T \Sigma_{1_u}^{-1} + \mu_2^T \Sigma_{2_u}^{-1})x + (\mu_1^T \Sigma_{1_u}^{-1}\mu_1 + \mu_2^T \Sigma_{2_u}^{-1}\mu_2))} \\
&= C_1 e^{-\frac{1}{2}(x^T(\Sigma_{1_u}^{-1} + \Sigma_{2_u}^{-1})x - x^T(\Sigma_{1_u}^{-1} + \Sigma_{2_u}^{-1})(\Sigma_{1_u}^{-1} + \Sigma_{2_u}^{-1})^{-1}(\Sigma_{1_u}^{-1}\mu_1 + \Sigma_{2_u}^{-1}\mu_2) - (\mu_1^T \Sigma_{1_u}^{-1} + \mu_2^T \Sigma_{2_u}^{-1})x + (\mu_1^T \Sigma_{1_u}^{-1}\mu_1 + \mu_2^T \Sigma_{2_u}^{-1}\mu_2))} \\
&= C_1 e^{-\frac{1}{2}(x^T(\Sigma_{1_u}^{-1} + \Sigma_{2_u}^{-1})(x - (\Sigma_{1_u}^{-1} + \Sigma_{2_u}^{-1})^{-1}(\Sigma_{1_u}^{-1}\mu_1 + \Sigma_{2_u}^{-1}\mu_2)) - (\mu_1^T \Sigma_{1_u}^{-1} + \mu_2^T \Sigma_{2_u}^{-1})x + (\mu_1^T \Sigma_{1_u}^{-1}\mu_1 + \mu_2^T \Sigma_{2_u}^{-1}\mu_2))}
\end{aligned} \tag{13}$$

$$\begin{aligned}
 &= C_1 e^{-\frac{1}{2}(x^T(\Sigma_{1_u}^{-1} + \Sigma_{2_u}^{-1})(x - (\Sigma_{1_u}^{-1} + \Sigma_{2_u}^{-1})^{-1}(\Sigma_{1_u}^{-1}\mu_1 + \Sigma_{2_u}^{-1}\mu_2)) - (\mu_1^T \Sigma_{1_u}^{-1} + \mu_2^T \Sigma_{2_u}^{-1})x + (\mu_1^T \Sigma_{1_u}^{-1} + \mu_2^T \Sigma_{2_u}^{-1})(\Sigma_{1_u}^{-1} + \Sigma_{2_u}^{-1})^{-1}(\Sigma_{1_u}^{-1}\mu_1 + \Sigma_{2_u}^{-1}\mu_2) + C_2)} \\
 &= C_1 e^{-\frac{1}{2}(x^T(\Sigma_{1_u}^{-1} + \Sigma_{2_u}^{-1})(x - (\Sigma_{1_u}^{-1} + \Sigma_{2_u}^{-1})^{-1}(\Sigma_{1_u}^{-1}\mu_1 + \Sigma_{2_u}^{-1}\mu_2)) - (\mu_1^T \Sigma_{1_u}^{-1} + \mu_2^T \Sigma_{2_u}^{-1})(x - (\Sigma_{1_u}^{-1} + \Sigma_{2_u}^{-1})^{-1}(\Sigma_{1_u}^{-1}\mu_1 + \Sigma_{2_u}^{-1}\mu_2)) + C_2)} \\
 &= C_1 e^{-\frac{1}{2}(x^T(\Sigma_{1_u}^{-1} + \Sigma_{2_u}^{-1})(x - (\Sigma_{1_u}^{-1} + \Sigma_{2_u}^{-1})^{-1}(\Sigma_{1_u}^{-1}\mu_1 + \Sigma_{2_u}^{-1}\mu_2)) - (\mu_1^T \Sigma_{1_u}^{-1} + \mu_2^T \Sigma_{2_u}^{-1})(\Sigma_{1_u}^{-1} + \Sigma_{2_u}^{-1})^{-1}(\Sigma_{1_u}^{-1}\mu_1 + \Sigma_{2_u}^{-1}\mu_2)) + C_2)} \\
 &= C_1 e^{-\frac{1}{2}((x^T - (\mu_1^T \Sigma_{1_u}^{-1} + \mu_2^T \Sigma_{2_u}^{-1})(\Sigma_{1_u}^{-1} + \Sigma_{2_u}^{-1})^{-1})(\Sigma_{1_u}^{-1} + \Sigma_{2_u}^{-1})(x - (\Sigma_{1_u}^{-1} + \Sigma_{2_u}^{-1})^{-1}(\Sigma_{1_u}^{-1}\mu_1 + \Sigma_{2_u}^{-1}\mu_2)) + C_2)} \\
 &= C_1 e^{-\frac{1}{2}((x - (\Sigma_{1_u}^{-1} + \Sigma_{2_u}^{-1})^{-1}(\Sigma_{1_u}^{-1}\mu_1 + \Sigma_{2_u}^{-1}\mu_2))^T (\Sigma_{1_u}^{-1} + \Sigma_{2_u}^{-1})(x - (\Sigma_{1_u}^{-1} + \Sigma_{2_u}^{-1})^{-1}(\Sigma_{1_u}^{-1}\mu_1 + \Sigma_{2_u}^{-1}\mu_2)) + C_2)} \\
 &= C_3 e^{-\frac{1}{2}(x - (\Sigma_{1_u}^{-1} + \Sigma_{2_u}^{-1})^{-1}(\Sigma_{1_u}^{-1}\mu_1 + \Sigma_{2_u}^{-1}\mu_2))^T (\Sigma_{1_u}^{-1} + \Sigma_{2_u}^{-1})(x - (\Sigma_{1_u}^{-1} + \Sigma_{2_u}^{-1})^{-1}(\Sigma_{1_u}^{-1}\mu_1 + \Sigma_{2_u}^{-1}\mu_2))},
 \end{aligned} \tag{14}$$

where $C_1 = (2\pi)^{-\frac{D}{2}} \det[\Sigma_1]^{-\frac{1-\alpha_u}{2}} \det[\Sigma_2]^{-\frac{\alpha_u}{2}}$, C_2 is a constant for aborting the terms used for completing the square relative to x , and $C_3 = C_1 e^{-\frac{1}{2}C_2}$. The last formula of Eq. (13) is an unnormalized Gaussian curve with covariance $(\Sigma_{1_u}^{-1} + \Sigma_{2_u}^{-1})^{-1}$ and mean $(\Sigma_{1_u}^{-1} + \Sigma_{2_u}^{-1})^{-1}(\Sigma_{1_u}^{-1}\mu_1 + \Sigma_{2_u}^{-1}\mu_2)$. Therefore, we can get $\Sigma_{\alpha_u} = ((1-\alpha_u)\Sigma_1^{-1} + \alpha_u\Sigma_2^{-1})^{-1}$ and $\mu_{\alpha_u} = \Sigma_{\alpha_u}((1-\alpha_u)\Sigma_1^{-1}\mu_1 + \alpha_u\Sigma_2^{-1}\mu_2)$. After the normalization step, we can get a Gaussian distribution $\mathcal{N}_{\alpha_u}(\mu_{\alpha_u}, \Sigma_{\alpha_u})$.

As for $JS^{G_{\alpha_u}}$, we first calculate $\mathbb{E}_{\mathcal{N}_1}[\log \mathcal{N}_1 - \log \mathcal{N}_{\alpha_u}]$ as follows:

$$\begin{aligned}
 &\mathbb{E}_{\mathcal{N}_1}[\log \mathcal{N}_1 - \log \mathcal{N}_{\alpha_u}] \\
 &= \frac{1}{2} \mathbb{E}_{\mathcal{N}_1}[-\log \det[\Sigma_1] - (x - \mu_1)^T \Sigma_1^{-1}(x - \mu_1) + \log \det[\Sigma_{\alpha_u}] + (x - \mu_{\alpha_u})^T \Sigma_{\alpha_u}^{-1}(x - \mu_{\alpha_u})] \\
 &= \frac{1}{2} \left(\log \frac{\det[\Sigma_{\alpha_u}]}{\det[\Sigma_1]} + \mathbb{E}_{\mathcal{N}_1}[-(x - \mu_1)^T \Sigma_1^{-1}(x - \mu_1) + (x - \mu_{\alpha_u})^T \Sigma_{\alpha_u}^{-1}(x - \mu_{\alpha_u})] \right) \\
 &= \frac{1}{2} \left(\log \frac{\det[\Sigma_{\alpha_u}]}{\det[\Sigma_1]} + \mathbb{E}_{\mathcal{N}_1}[-\text{tr}[\Sigma_1^{-1}\Sigma_1] + \text{tr}[\Sigma_{\alpha_u}^{-1}(xx^T - 2x\mu_{\alpha_u}^T + \mu_{\alpha_u}\mu_{\alpha_u}^T)] \right) \\
 &= \frac{1}{2} \log \frac{\det[\Sigma_{\alpha_u}]}{\det[\Sigma_1]} - \frac{D}{2} + \frac{1}{2} \mathbb{E}_{\mathcal{N}_1}[\text{tr}[\Sigma_{\alpha_u}^{-1}(xx^T - 2x\mu_{\alpha_u}^T + \mu_{\alpha_u}\mu_{\alpha_u}^T)]] \\
 &= \frac{1}{2} \log \frac{\det[\Sigma_{\alpha_u}]}{\det[\Sigma_1]} - \frac{D}{2} + \frac{1}{2} \mathbb{E}_{\mathcal{N}_1}[\text{tr}[\Sigma_{\alpha_u}^{-1}((x - \mu_1)(x - \mu_1)^T + 2\mu_1 x^T - \mu_1 \mu_1^T - 2x\mu_{\alpha_u}^T + \mu_{\alpha_u}\mu_{\alpha_u}^T)]] \\
 &= \frac{1}{2} \log \frac{\det[\Sigma_{\alpha_u}]}{\det[\Sigma_1]} - \frac{D}{2} + \frac{1}{2} \text{tr}[\Sigma_{\alpha_u}^{-1}(\Sigma_1 + \mu_1 \mu_1^T - 2\mu_{\alpha_u} \mu_1^T + \mu_{\alpha_u} \mu_{\alpha_u}^T)] \\
 &= \frac{1}{2} \log \frac{\det[\Sigma_{\alpha_u}]}{\det[\Sigma_1]} - \frac{D}{2} + \frac{1}{2} \text{tr}[\Sigma_{\alpha_u}^{-1}\Sigma_1] + \frac{1}{2} \text{tr}[\mu_1^T \Sigma_{\alpha_u}^{-1} \mu_1 - 2\mu_1^T \Sigma_{\alpha_u}^{-1} \mu_{\alpha_u} + \mu_{\alpha_u}^T \Sigma_{\alpha_u}^{-1} \mu_{\alpha_u}] \\
 &= \frac{1}{2} \log \frac{\det[\Sigma_{\alpha_u}]}{\det[\Sigma_1]} - \frac{D}{2} + \frac{1}{2} \text{tr}[\Sigma_{\alpha_u}^{-1}\Sigma_1] + \frac{1}{2} (\mu_{\alpha_u} - \mu_1)^T \Sigma_{\alpha_u}^{-1} (\mu_{\alpha_u} - \mu_1).
 \end{aligned} \tag{15}$$

The calculation of $\mathbb{E}_{\mathcal{N}_2}[\log \mathcal{N}_2 - \log \mathcal{N}_{\alpha_u}]$ is the same, and then, $JS^{G_{\alpha_u}}$ is given by:

$$\begin{aligned}
 JS^{G_{\alpha_u}} &= \frac{1 - \alpha_u}{2} \log \frac{\det[\Sigma_{\alpha_u}]}{\det[\Sigma_1]} - \frac{D(1 - \alpha_u)}{2} + \frac{1 - \alpha_u}{2} \text{tr}[\Sigma_{\alpha_u}^{-1}\Sigma_1] + \frac{1 - \alpha_u}{2} (\mu_{\alpha_u} - \mu_1)^T \Sigma_{\alpha_u}^{-1} (\mu_{\alpha_u} - \mu_1) + \\
 &\frac{\alpha_u}{2} \log \frac{\det[\Sigma_{\alpha_u}]}{\det[\Sigma_2]} - \frac{D\alpha_u}{2} + \frac{\alpha_u}{2} \text{tr}[\Sigma_{\alpha_u}^{-1}\Sigma_2] + \frac{\alpha_u}{2} (\mu_{\alpha_u} - \mu_2)^T \Sigma_{\alpha_u}^{-1} (\mu_{\alpha_u} - \mu_2) \\
 &= \frac{1}{2} \left(\log \frac{\det[\Sigma_{\alpha_u}]^{1-\alpha_u}}{\det[\Sigma_1]^{1-\alpha_u}} + \log \frac{\det[\Sigma_{\alpha_u}]^{\alpha_u}}{\det[\Sigma_2]^{\alpha_u}} \right) - \frac{D}{2} + \frac{1}{2} \text{tr}(\Sigma_{\alpha_u}^{-1}((1-\alpha_u)\Sigma_1 + \alpha_u\Sigma_2)) + \\
 &\frac{1 - \alpha_u}{2} (\mu_{\alpha_u} - \mu_1)^T \Sigma_{\alpha_u}^{-1} (\mu_{\alpha_u} - \mu_1) + \frac{\alpha_u}{2} (\mu_{\alpha_u} - \mu_2)^T \Sigma_{\alpha_u}^{-1} (\mu_{\alpha_u} - \mu_2) \\
 &= \frac{1}{2} \left(\log \left[\frac{\det[\Sigma_{\alpha_u}]}{\det[\Sigma_1]^{1-\alpha_u} \det[\Sigma_2]^{\alpha_u}} \right] - D + \text{tr}(\Sigma_{\alpha_u}^{-1}((1-\alpha_u)\Sigma_1 + \alpha_u\Sigma_2)) + (1 - \alpha_u)(\mu_{\alpha_u} - \mu_1)^T \Sigma_{\alpha_u}^{-1} (\mu_{\alpha_u} - \mu_1) + \right. \\
 &\left. \alpha_u (\mu_{\alpha_u} - \mu_2)^T \Sigma_{\alpha_u}^{-1} (\mu_{\alpha_u} - \mu_2) \right).
 \end{aligned} \tag{16}$$

As to the dual form $JS_*^{G_{\alpha_u}}$, we calculate $\mathbb{E}_{\mathcal{N}_{\alpha_u}}[\log \mathcal{N}_{\alpha_u} - \log \mathcal{N}_1]$, which is given by:

$$\frac{1}{2} \log \frac{\det[\Sigma_1]}{\det[\Sigma_{\alpha_u}]} - \frac{D}{2} + \frac{1}{2} \text{tr}[\Sigma_1^{-1}\Sigma_{\alpha_u}] + \frac{1}{2} (\mu_1 - \mu_{\alpha_u})^T \Sigma_1^{-1} (\mu_1 - \mu_{\alpha_u}). \tag{17}$$

Then, the calculation of $\mathbb{E}_{\mathcal{N}_{\alpha_u}}[\log \mathcal{N}_{\alpha_u} - \log \mathcal{N}_2]$ is the same, and $JS_*^{G_{\alpha_u}}$ is given by:

$$\begin{aligned}
 JS_*^{G_{\alpha_u}} &= \frac{1 - \alpha_u}{2} \log \frac{\det[\Sigma_1]}{\det[\Sigma_{\alpha_u}]} - \frac{D(1 - \alpha_u)}{2} + \frac{1 - \alpha_u}{2} \text{tr}[\Sigma_1^{-1} \Sigma_{\alpha_u}] + \frac{1 - \alpha_u}{2} (\mu_1 - \mu_{\alpha_u})^T \Sigma_1^{-1} (\mu_1 - \mu_{\alpha_u}) + \\
 &\quad \frac{\alpha_u}{2} \log \frac{\det[\Sigma_2]}{\det[\Sigma_{\alpha_u}]} - \frac{D\alpha_u}{2} + \frac{\alpha_u}{2} \text{tr}[\Sigma_2^{-1} \Sigma_{\alpha_u}] + \frac{\alpha_u}{2} (\mu_2 - \mu_{\alpha_u})^T \Sigma_2^{-1} (\mu_2 - \mu_{\alpha_u}) \\
 &= \frac{1}{2} \left(\log \frac{\det[\Sigma_1]^{1 - \alpha_u}}{\det[\Sigma_{\alpha_u}]^{1 - \alpha_u}} + \log \frac{\det[\Sigma_2]^{\alpha_u}}{\det[\Sigma_{\alpha_u}]^{\alpha_u}} \right) - \frac{D}{2} + \frac{1}{2} \text{tr} \left((1 - \alpha_u) \Sigma_1^{-1} \Sigma_{\alpha_u} + \alpha_u \Sigma_2^{-1} \Sigma_{\alpha_u} \right) + \frac{1 - \alpha_u}{2} \mu_1^T \Sigma_1^{-1} \mu_1 - \\
 &\quad \frac{1 - \alpha_u}{2} \mu_1^T \Sigma_1^{-1} \mu_{\alpha_u} - \frac{1 - \alpha_u}{2} \mu_{\alpha_u}^T \Sigma_1^{-1} \mu_1 + \frac{1 - \alpha_u}{2} \mu_{\alpha_u}^T \Sigma_1^{-1} \mu_{\alpha_u} + \frac{\alpha_u}{2} \mu_2^T \Sigma_2^{-1} \mu_2 - \frac{\alpha_u}{2} \mu_2^T \Sigma_2^{-1} \mu_{\alpha_u} \\
 &\quad - \frac{\alpha_u}{2} \mu_{\alpha_u}^T \Sigma_2^{-1} \mu_2 + \frac{\alpha_u}{2} \mu_{\alpha_u}^T \Sigma_2^{-1} \mu_{\alpha_u} \\
 &= \frac{1}{2} \log \frac{\det[\Sigma_1]^{1 - \alpha_u} \det[\Sigma_2]^{\alpha_u}}{\det[\Sigma_{\alpha_u}]} - \frac{D}{2} + \frac{1}{2} \text{tr} \left(\underbrace{((1 - \alpha_u) \Sigma_1^{-1} + \alpha_u \Sigma_2^{-1})}_{\Sigma_{\alpha_u}^{-1}} \Sigma_{\alpha_u} \right) + \frac{1 - \alpha_u}{2} \mu_1^T \Sigma_1^{-1} \mu_1 - \\
 &\quad (1 - \alpha_u) \mu_1^T \Sigma_1^{-1} \mu_{\alpha_u} + \frac{1 - \alpha_u}{2} \mu_{\alpha_u}^T \Sigma_1^{-1} \mu_{\alpha_u} + \frac{\alpha_u}{2} \mu_2^T \Sigma_2^{-1} \mu_2 - \alpha_u \mu_2^T \Sigma_2^{-1} \mu_{\alpha_u} + \frac{\alpha_u}{2} \mu_{\alpha_u}^T \Sigma_2^{-1} \mu_{\alpha_u} \\
 &= \frac{1}{2} \log \frac{\det[\Sigma_1]^{1 - \alpha_u} \det[\Sigma_2]^{\alpha_u}}{\det[\Sigma_{\alpha_u}]} + \frac{1 - \alpha_u}{2} \mu_1^T \Sigma_1^{-1} \mu_1 + \frac{\alpha_u}{2} \mu_2^T \Sigma_2^{-1} \mu_2 - \underbrace{((1 - \alpha_u) \mu_1^T \Sigma_1^{-1} + \alpha_u \mu_2^T \Sigma_2^{-1})}_{\mu_{\alpha_u}^T \Sigma_{\alpha_u}^{-1}} \mu_{\alpha_u} + \\
 &\quad \frac{1}{2} \mu_{\alpha_u}^T \underbrace{((1 - \alpha_u) \Sigma_1^{-1} + \alpha_u \Sigma_2^{-1})}_{\Sigma_{\alpha_u}^{-1}} \mu_{\alpha_u} \\
 &= \frac{1}{2} \left(\log \frac{\det[\Sigma_1]^{1 - \alpha_u} \det[\Sigma_2]^{\alpha_u}}{\det[\Sigma_{\alpha_u}]} + (1 - \alpha_u) \mu_1^T \Sigma_1^{-1} \mu_1 + \alpha_u \mu_2^T \Sigma_2^{-1} \mu_2 - \mu_{\alpha_u}^T \Sigma_{\alpha_u}^{-1} \mu_{\alpha_u} \right).
 \end{aligned} \tag{18}$$

□

C. Implementation Details

The deep neural network configuration and training details are summarized in Table 5. As for the NP-Match related hyperparameters, we set the lengths of both memory banks (\mathcal{Q}) to 2560. The coefficient (β) is set to 0.01, and we sample $T = 10$ latent vectors for each target point. The uncertainty threshold (τ_u) is set to 0.4 for CIFAR-10, CIFAR-100, and STL-10, and it is set to 1.2 for ImageNet. NP-Match is trained by using stochastic gradient descent (SGD) with a momentum of 0.9. The initial learning rate is set to 0.03 for CIFAR-10, CIFAR-100, and STL-10, and it is set to 0.05 for ImageNet. The learning rate is decayed with a cosine decay schedule (Loshchilov & Hutter, 2016), and NP-Match is trained for 2^{20} iterations. The MLPs used in the NP model all have two layers with \mathcal{M} hidden units for each layer. For WRN, \mathcal{M} is a quarter of the channel dimension of the last convolutional layer, and as for ResNet-50, \mathcal{M} is equal to 256. To compete with the most recent SOTA method (Zhang et al., 2021), we followed this work to use the Curriculum Pseudo Labeling (CPL) strategy in our method and UPS (Rizve et al., 2021). We initialize each memory bank with a random vector.

Dataset	CIFAR-10	CIFAR-100	STL-10	ImageNet
Model	WRN-28-2	WRN-28-8	WRN-37-2	ResNet-50
Weight Decay	5e-4	1e-3	5e-4	1e-4
Batch Size (B)	64			256
μ	7			1
Confidence Threshold (τ_c)	0.95			0.7
EMA Momentum	0.999			
λ_u	1.0			

Table 5. Details of the training setting.

We ran each label amount setting for three times using different random seeds to obtain the error bars on CIFAR-10, CIFAR-100, and STL-10, but on ImageNet, we only ran for once. GeForce GTX 1080 Ti GPUs were used for the experiments on CIFAR-10, CIFAR-100, and STL-10, while Tesla V100 SXM2 GPUs were used for the experiments on ImageNet.

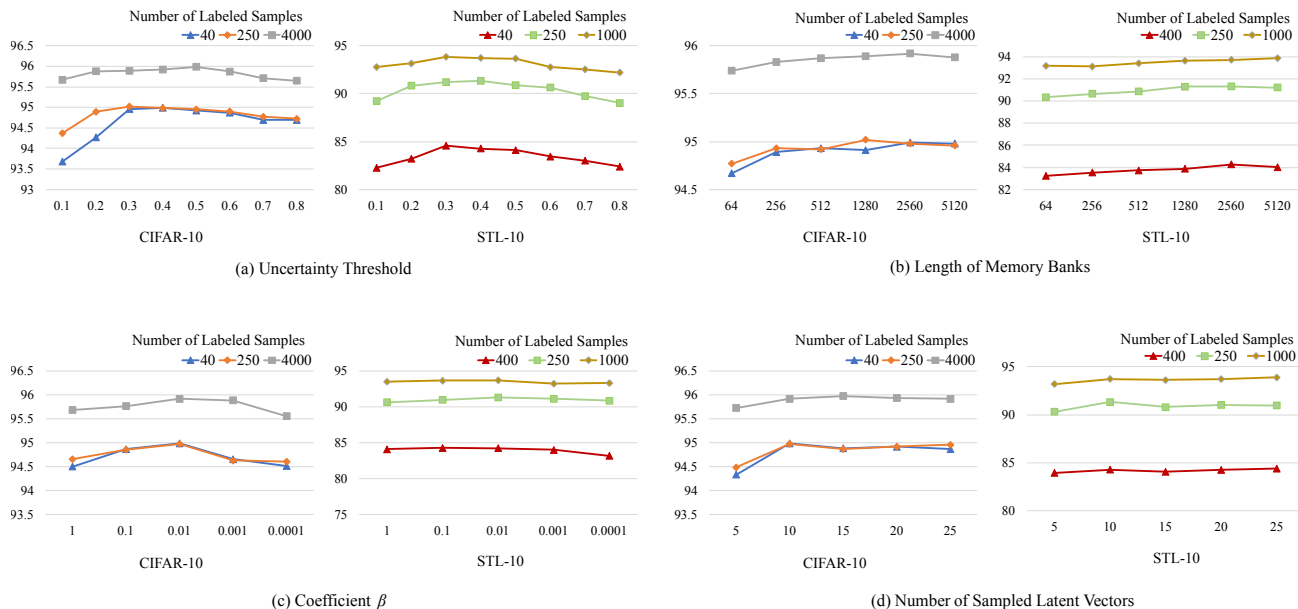


Figure 4. Performance for different hyperparameters.

D. Hyperparameter Exploration

We did additional experiments on CIFAR-10 and STL-10 in terms of the hyperparameters related to the NP model in NP-Match, in order to explore how performance is affected by the changes of hyperparameters, which may provide some hints to readers for applying NP-Match to other datasets. We consider four hyperparameters in total, including the uncertainty threshold (τ_u), the length of memory banks (\mathcal{Q}), the coefficient of $JS^{G_{\alpha_u}}$ (β), and the number of sampled latent vectors (T). By Figure 4(a), a reasonable τ_u is important. Specifically, lower τ_u usually leads to worse performance, because lower τ_u enforces NP-Match to select a limited number of unlabeled data during training, which is equivalent to training the whole framework with a small dataset. Conversely, when τ_u is too large, more uncertain unlabeled samples are chosen, whose pseudo-labels might be incorrect, and using these uncertain samples to train the framework can also lead to a poor performance. Furthermore, the difficulty of a training set also affects the setting of τ_u , as a more difficult dataset usually has more classes and hard samples (e.g., ImageNet), which makes the uncertainties of predictions large, so that τ_u should be adjusted accordingly. From Figure 4(b), the performance becomes better with the increase of \mathcal{Q} . When more context points are used, the more information is involved for inference, and then the NP model can better estimate the global latent variable and make predictions. This observation is consistent with the experimental results where the original NPs are used for image completion (Garnelo et al., 2018b). Figure 4(c) shows the ablation study of β that controls the contribution of $JS^{G_{\alpha_u}}$ to the total loss function, and when $\beta = 0.01$, we can obtain the best accuracy on both datasets. By Figure 4(d), if T is smaller than 5, then the performance will go down, but when T is further increased, then the performance of NP-Match is not influenced greatly.