

Effect of Random Learning Rate: Theoretical Analysis of SGD Dynamics in Non-Convex Optimization via Stationary Distribution

author names withheld

Under Review for the Workshop on High-dimensional Learning Dynamics, 2024

Abstract

We consider a variant of the stochastic gradient descent (SGD) with a random learning rate and reveal its convergence properties. SGD is a widely used stochastic optimization algorithm in machine learning, especially deep learning. Numerous studies reveal the convergence properties of SGD and its simplified variants. Among these, the analysis of convergence using a stationary distribution of updated parameters provides generalizable results. However, to obtain a stationary distribution, the update direction of the parameters must not degenerate, which limits the applicable variants of SGD. In this study, we consider a novel SGD variant, Poisson SGD, which has degenerated parameter update directions and instead utilizes a random learning rate. Consequently, we demonstrate that a distribution of a parameter updated by Poisson SGD converges to a stationary distribution under weak assumptions on a loss function. Based on this, we further show that Poisson SGD finds global minima in non-convex optimization problems and also evaluate the generalization error using this method. As a proof technique, we approximate the distribution by Poisson SGD with that of the bouncy particle sampler (BPS) and derive its stationary distribution, using the theoretical advance of the piece-wise deterministic Markov process (PDMP).

1. Introduction

Stochastic gradient descent (SGD) stands out as a widely employed optimization algorithm in machine learning. It falls under the category of stochastic optimization, where parameters are updated with randomness from the mini-batch sampling. SGD is valued for two main reasons in optimization: (i) it is memory-efficient and requires only low computational resources by updating parameters from a fraction of the training data at each iteration [5], and (ii) models optimized with SGD have less generalization error than those optimized by other algorithms such as gradient descent (GD) for neural networks. Owing to these advantages, SGD has been one of the standard methods for training deep learning models [20, 22, 33].

Among the methods analyzing the properties of SGD, one of the most general approaches is to study a *stationary distribution* of parameters updated by SGD and its variants. The stationary distribution is a distribution that remains unchanged when the parameter is updated by one step. It is useful in theoretical analysis, because (i) it can analyze the global dynamics of the optimization algorithm, and (ii) it can be applied to a wide range of loss functions regardless of its shape. For these reasons, we can use it to investigate the optimization of complex loss functions such as deep neural networks. For example, [14] studied the stationary distribution of the parameter optimized by SGD when the loss function is strongly convex, and [29] studied the stationary distribution of SGLD when the loss function is non-convex.

Despite the above advantages, there are not many SGD variants to which stationary distribution analysis can be applied. This is because, to use the analysis by a stationary distribution, it is required that the direction of parameter updates by an algorithm does not degenerate; in other words, there must be no directions that are not being explored. Examples of such variants are (i) SGLD [10, 15, 32], which adds a Gaussian noise to the parameter update of SGD and (ii) Gaussian SGD [19, 21, 23], which assumes that the noise of SGD on the gradient induced by the mini-batch sampling is non-degenerate Gaussian. In contrast, the parameter update of SGD degenerates in many practical cases, such as deep learning [25, 30, 34]. Hence, there is a gap between the variants of SGD considered in the theoretical analysis and the empirical facts about SGD. This gap fosters the following question:

Do parameters optimized by a variant of SGD have a stationary distribution even if the update direction degenerates - and if so, what is the form of it?

1.1. Our Contribution

We theoretically prove that a variant of SGD has a stationary distribution even if the update direction degenerates. Specifically, we develop a novel SGD variant with a *random learning rate*, which follows the Poisson process depending on a mini-batch gradient. We call the variant *Poisson SGD*, and prove that the distribution of a parameter updated by Poisson SGD converges to a stationary distribution. As a result, we provide a positive answer to the question posed above: even with a degenerated parameter update, it is possible to construct a variant of SGD that reaches a stationary distribution by using a random learning rate.

Our specific contributions are overviewed in Appendix A.1.

1.2. Notation

See Appendix A.2.

2. Preliminary

2.1. Problem Setup: Empirical Risk Minimization

We consider the following stochastic optimization problem. Let \mathcal{Z} be a compact sample space, and consider a probability measure P_* on \mathcal{Z} . Suppose that we observe n samples $\mathbf{z} = \{z_1, \dots, z_n\} \subset \mathcal{Z}$, that are independently and identically generated from the measure P_* . Using the samples, we consider an empirical risk with a loss function. Let $\Theta \subset \mathbb{R}^d$ be a compact parameter space such as torus or sphere, and define $W := \text{diam}(\Theta)$. With a (potentially non-convex) loss function $\ell : \mathcal{Z} \times \Theta \rightarrow \mathbb{R}$, we consider the following empirical risk with the samples as $L_{\mathbf{z}}(\theta) = \frac{1}{n} \sum_{i=1}^n \ell(z_i; \theta)$, $\theta \in \Theta$.

Our goal is to find a global minimum of the empirical risk $L_{\mathbf{z}}(\cdot)$, which is defined as a parameter $\theta^* \in \Theta$ which satisfies the Empirical Risk Minimization (ERM) problem: $L_{\mathbf{z}}(\theta^*) = \min_{\theta' \in \Theta} L_{\mathbf{z}}(\theta')$.

2.2. Gradient Descent Algorithm and Variants

To find the global minimum θ^* in ERM problem, we often use the optimization algorithm called stochastic gradient descent (SGD) with momentum.

2.2.1. GENERAL FORM OF STOCHASTIC GRADIENT DESCENT

We give a formal definition of SGD with a momentum term associated with empirical risk $L_{\mathbf{z}}(\theta)$. Let $K \in \mathbb{N}$ be the number of iterations. The SGD with momentum generates a sequence of Θ -valued random parameters $\theta_1, \dots, \theta_K$ and \mathbb{R}^d -valued random vectors v_1, \dots, v_K , by the following procedure.

Let $\theta_0 \in \Theta$ be an arbitrary parameter for the initialization, $v_0 \in \mathbb{R}^d$ as an initial velocity vector, and $m \in [n]$ be a number of sub-samples, i.e., the batch size. Suppose that we observe the n samples $z := \{z_1, \dots, z_n\}$, i.e., the full-batch. For $k = 1, \dots, K$, we uniformly sample m integers $I^{(k)} = \{i_1, i_2, \dots, i_m\}$ from $[n]$, which is called mini-batch sampling with the batch-size m . We define an associated mini-batch risk as

$$\widehat{L}_{\mathbf{z}}^{(k)}(\theta) := \frac{1}{m} \sum_{i \in I^{(k)}} \ell(z_i; \theta). \quad (1)$$

Then, with initial values $\theta_0 \in \Theta$ and $v_0 \in \mathbb{R}^d$, the SGD with momentum generates the parameter and the velocity vector by the following recursive formula for $k = 1, \dots, K$:

$$\theta_k = \theta_{k-1} + \eta_k v_{k-1}, \text{ and } v_k = v_{k-1} - \alpha_k \nabla \widehat{L}_{\mathbf{z}}^{(k)}(\theta_k), \quad (2)$$

where $\eta_k > 0$ is a learning rate and $\alpha_k \in \mathbb{R}$ is a momentum coefficient. This form is generic and can be identical to other forms of SGD with momentum [28, 31] by adjusting the parameters η and α .

3. Our SGD Variant: Poisson SGD

In this section, we introduce our algorithm, *Poisson SGD*, which is a variant of SGD with a random learning rate η and momentum coefficient α . We design our method so that the parameter can search the whole parameter space owing to the design.

We describe the random learning rate. In preparation, we define the following exponential distribution function with a function $f : \Theta \rightarrow \mathbb{R}^d$ and parameters $\theta \in \Theta, v \in \mathbb{S}^{d-1}$:

$$E(f(\cdot), \theta, v) := \exp \left(- \int_0^t \{ \max \{ \langle f(\theta + rv), v \rangle, 0 \} + C_P \} dr \right),$$

where $C_P > 0$ is some constant. Then, for each update $k = 1, \dots, K$, we design the random learning rate η_k following the exponential distribution:

$$P(\eta_k \geq t) = E(\beta \nabla \widehat{L}_{\mathbf{z}}^{(k)}(\cdot), \theta_{k-1}, v_{k-1}), \quad (3)$$

where $\beta > 0$ is the hyper-parameter of Poisson SGD, called an inverse temperature parameter.

Second, we select the momentum coefficient α_k for each $k = 1, \dots, K$ as

$$\alpha_k = 2 \frac{\langle \nabla \widehat{L}_{\mathbf{z}}^{(k)}(\theta_k), v_{k-1} \rangle}{\|\nabla \widehat{L}_{\mathbf{z}}^{(k)}(\theta_k)\|^2}. \quad (4)$$

This setup keeps the length of the velocity vector constant as $\|v_k\| = 1$ for every k (See Proposition 7 in Appendix), and only uses its angle to update the parameters. We update the parameter by changing η_k and α_k in every iteration. The pseudo-code of Poisson SGD is shown in Algorithm 1 in Appendix A.5.

4. Convergence Theory for Poisson SGD

We provide theoretical results on the convergence of Poisson SGD (Algorithm 1). Our main interest is a distribution of the generated parameter θ_K by Poisson SGD associated with the empirical risk minimization problem.

4.1. Stationary Distribution of Poisson SGD

In this section, we show that the parameter θ_K by the Poisson SGD follows a stationary distribution. Formally, we define the stationary distribution of the Markov process. In preparation, we utilize the notion of *transition probability* $Q(\theta, dw)$ from a distribution $p_0(\theta)$ to another $p_1(\theta)$ on Θ , that is, $p_1(w) = \int_{\Theta} Q(\theta, dw)p_0(d\theta)$ holds.

Definition 1 *Let $Q(\theta, dw)$ be the transition probability of a Markov process in Θ . If the following equation holds, we call the probability distribution $\pi(\theta)$ stationary distribution of the Markov process:*

$$\pi(dw) = \int_{\Theta} Q(\theta, dw)\pi(d\theta).$$

A stationary distribution is an useful notion to represent a limit of the parameter distribution, and it enables us to analyze where the parameter converges by algorithms. For example, see the theoretical framework to analyze stochastic optimization algorithms by [29].

We provide several principal assumptions. First, we consider the basis assumptions on the loss function $\ell(\cdot; \cdot)$. The following conditions are fairly general for the analysis of stochastic optimization algorithms, e.g. [2].

Assumption 1 (Loss function) *The loss function $\ell : \mathcal{Z} \times \Theta \rightarrow \mathbb{R}_{\geq 0}$ satisfies the following conditions:*

- $\ell(z; \theta)$ is absolutely continuous and differentiable with respect to $\theta \in \Theta$ for every $z \in \mathcal{Z}$.
- $\nabla_{\theta}\ell(z; \theta)$ is continuous in θ and z for all $\theta \in \Theta$ and $z \in \mathcal{Z}$.

The first condition is satisfied by a large class of models, such as linear regression model, or deep neural networks whose activation function is ReLU or sigmoid function. From the second condition, we define an upper bound $M_{\ell} := \max_{\theta \in \Theta, z \in \mathcal{Z}} \|\nabla_{\theta}\ell(z; \theta)\|$ since Θ and \mathcal{Z} are compact.

In preparation, we define a probability measure on Θ for arbitrary $\beta, \varepsilon > 0$, whose density is written as follows:

$$\mu_{\mathbf{z}}^{(\beta, \varepsilon)}(d\theta) \propto \left(\beta M_{\ell} + \frac{1}{\varepsilon} + a_d \beta \|\nabla L_{\mathbf{z}}(\theta)\| \right) \exp(-\beta L_{\mathbf{z}}(\theta)) d\theta, \quad (5)$$

where $a_d := \Gamma(d/2)/(\sqrt{\pi}\Gamma(d/2 + 1/2))$. The probability measure (5) is concentrated around the global minima of $L_{\mathbf{z}}(\theta)$, since the dominant exponential term $\exp(-\beta L_{\mathbf{z}}(\theta))$ in (5) increases in $L_{\mathbf{z}}(\theta)$. In addition, as the inverse temperature parameter β increases, the measure $\mu_{\mathbf{z}}^{(\beta, \varepsilon)}$ concentrates more around the global minimum.

We show our results on the convergence of the stationary distribution. The discrepancy is measured by the Wasserstein distance $\mathcal{W}_1(\cdot, \cdot)$. We remark that this theorem is the integration of Theorem 8 and Theorem 10 appearing later in Section D. Recall that we defined $W := \text{diam}(\Theta)$. Let $\mu_{\mathbf{z}, K}$ be a distribution of the output θ_K from the Poisson SGD in Algorithm 1 with the given dataset \mathbf{z} . We discuss the convergence of $\mu_{\mathbf{z}, K}$ as K increases.

Theorem 2 (Stationary distribution of Poisson SGD) Fix arbitrary $\beta, \varepsilon > 0$. Suppose Assumption 1 holds. We set the hyper-parameter of Poisson SGD as $C_P = 1/\varepsilon$. Then, for any $K \in \mathbb{N}$, the distribution $\mu_{z,K}$ satisfies

$$\mathcal{W}_1(\mu_{z,K}, \mu_z^{(\beta, \varepsilon)}) \leq 4\sqrt{d}K\varepsilon + W \cdot \kappa(\beta, \varepsilon, d)^K, \quad (6)$$

where $0 < \kappa(\beta, \varepsilon, d) < 1$ is a constant.

Moreover, if $\kappa(\beta, \varepsilon, d)$ satisfies $\lim_{K \rightarrow \infty} \kappa(\beta, \delta/K, d)^K = 0$ with some $\delta > 0$, there exists a sequence $\varepsilon = \varepsilon_K \searrow 0$ as $K \rightarrow \infty$ such that $\mathcal{W}_1(\mu_{z,K}, \mu_z^{(\beta, \varepsilon)}) = o(1)$ as $K \rightarrow \infty$ holds.

This theorem shows that the parameter distribution $\mu_{z,K}$ by Poisson SGD converges to the stationary distribution $\mu_z^{(\beta, \varepsilon)}$ owing to the random learning rate (3). This is contrast to ordinary SGD, which is not shown to converge to a stationary distribution. Further, Poisson SGD does not make any assumptions on the gradient noise $\xi_k^{(n,m)}$ in Appendix A.4, unlike SGLD, which converges to a stationary distribution by introducing Gaussianity in the gradient noise. Detail is in Appendix A.7.

4.2. Global Convergence

We discuss the global convergence statement, that is, the empirical risk $L_z(\theta_K)$ with Poisson SGD is minimized with high probability. We consider the additional assumption for the loss function ℓ :

Assumption 2 With some $c_1 > 0$, $\sup_{z \in \mathcal{Z}} \|\nabla \ell(z; \theta_1) - \nabla \ell(z; \theta_2)\| \leq c_1 \|\theta_1 - \theta_2\|$ holds for every $\theta_1 \neq \theta_2 \in \Theta$.

Then, we obtain the following global convergence theorem. We define $B := \sup_{z \in \mathcal{Z}} \|\nabla \ell(z; 0)\|$ by following Assumption 1.

Theorem 3 (Global convergence of Poisson SGD on empirical risk) Fix arbitrary $\beta, \varepsilon > 0$. Consider Poisson SGD in which $C_P = 1/\varepsilon$. Let the upper bound of $\mathcal{W}_1(\mu_{z,K}, \mu_z^{(\beta, \varepsilon)})$ obtained in Theorem 2 be $d_K(\beta, \varepsilon, d)$. Also, suppose that Assumption 1 and 2 hold. Then, it holds that

$$\begin{aligned} & \mathbb{E}_{\theta_K \sim \mu_{z,K}} [L_z(\theta_K)] - \min_{\theta \in \Theta} L_z(\theta) \\ & \leq (c_1 W + B) \sqrt{W d_K(\beta, \varepsilon, d)} + \frac{1}{\beta} \left(\frac{d}{2} \log \frac{e W^2 c_1 \beta}{d} + \log \left(1 + \frac{a_d (c_1 W + B)}{M_\ell} \right) \right). \end{aligned}$$

Theorem 3 states that we can make $\mathbb{E}[L_z(\theta_K)]$ be arbitrarily close to $\min_{\theta \in \Theta} L_z(\theta)$ by selecting large β , provided that we can make $d_K(\beta, \varepsilon, d)$ arbitrarily small by the choice of ε and K in spite of β . Intuitively, Poisson SGD achieves global convergence by appropriately adjusting the learning rate and momentum coefficient based on the shape of the loss function at the current location. Poisson SGD achieves the global convergence by the similar approach of global convergence of SGLD by [29].

5. Conclusion

We developed a new variant of SGD, Poisson SGD, whose search direction degenerates and derived its stationary distribution by incorporating a modification on the learning rate. The parameters trained by Poisson SGD are close enough to the global minima to take advantage of convergence to the stationary distribution. The generalization error is also evaluated. We believe that our work leads to the analysis of the actual SGD dynamics, not variants of it in the future.

References

- [1] Dominique Bakry, Ivan Gentil, and Michel Ledoux. *Analysis and Geometry of Markov Diffusion Operators*. Springer, 2014.
- [2] Andrea Bertazzi, Joris Bierkens, and Paul Dobson. Approximations of piecewise deterministic markov processes and their convergence properties. *Stochastic Processes and their Applications*, 154:91–153, 2022.
- [3] Léonard Blier, Pierre Wolinski, and Yann Ollivier. Learning with random learning rates. In *ECML PKDD*, 2019.
- [4] François Bolley and Cédric Villani. Weighted csiszár-kullback-pinsker inequalities and applications to transportation inequalities. *Annales de la Faculté des sciences de Toulouse: Mathématiques*, 14(3):331–352, 2005.
- [5] Léon Bottou. Stochastic gradient learning in neural networks. *Proceedings of Neuro-Nimes*, 91(8):12, 1991.
- [6] Alexandre Bouchard-Côté, Sebastian J Vollmer, and Arnaud Doucet. The bouncy particle sampler: A nonreversible rejection-free markov chain monte carlo method. *Journal of the American Statistical Association*, 113(522):855–867, 2018.
- [7] Zaiwei Chen, Shancong Mou, and Siva Theja Maguluri. Stationary behavior of constant stepsize sgd type algorithms: An asymptotic characterization. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 6(1)(19):1–24, 2022.
- [8] Xiang Cheng, Dong Yin, Peter Bartlett, and Michael Jordan. Stochastic gradient and langevin processes. In *International Conference on Machine Learning*, 2020.
- [9] Alexandre H. Thiery Chris Sherlock. A discrete bouncy particle sampler. *Biometrika*, 109: 335–349, 2022.
- [10] Arnak Dalalyan. Further and stronger analogy between sampling and optimization: Langevin monte carlo and gradient descent. In *Conference on Learning Theory*, 2017.
- [11] M.H.A Davis. Piecewise-deterministic markov processes: A general class of non-diffusion stochastic models. *Journal of the Royal Statistical Society, Series B (Methodological)*(46): 353–388, 1984.
- [12] M.H.A Davis. *Markov Models and Optimization*. Chapman & Hall/CRC Monographs on Statistics & Applied Probability. Taylor & Francis, 1993.
- [13] George Deligiannidis, Alexandre Bouchard-Côté, and Arnaud Doucet. Exponential ergodicity of the bouncy particle sampler. *The Annals of Statistics*, 47:1268–1287, 2019.
- [14] Aymeric Dieuleveut, Alain Durmus, and Francis Bach. Bridging the gap between constant step size stochastic gradient descent and markov chains. *The Annals of Statistics*, 48(3):1348–1382, 2020.

- [15] Alain Durmus and Éric Moulines. Sampling from a strongly log-concave distribution with the unadjusted langevin algorithm. *HAL preprint hal-01304430v1*, 2016.
- [16] Alain Durmus, Arnaud Guillin, and Pierre Monmarché. Geometric ergodicity of the bouncy particle sampler. *Annals of Applied Probability*, 30:2069–2098, 2020.
- [17] Alison L. Gibbs and Francis Edward Su. On choosing and bounding probability metrics. *International Statistical Review / Revue Internationale de Statistique*, 70(3):419–435, 2002.
- [18] Jeff Z. HaoChen, Colin Wei, Jason D. Lee, and Tengyu Ma. Shape matters: Understanding the implicit bias of the noise covariance. In *Conference On Learning Theory*, 2021.
- [19] Fengxiang He, Tongliang Liu, and Dacheng Tao. Control batch size and learning rate to generalize well: Theoretical and empirical evidence. In *Advances in Neural Information Processing Systems*, 2019.
- [20] Elad Hoffer, Itay Hubara, and Daniel Soudry. Train longer, generalize better: closing the generalization gap in large batch training of neural networks. *Advances in neural information processing systems*, 30, 2017.
- [21] Stanislaw Jastrzebski, Zachary Kenton, Devansh Arpit, Nicolas Ballas, Asja Fischer, Yoshua Bengio, and Amos Storkey. Three factors influencing minima in sgd. *arXiv preprint arXiv:1711.04623*, 2017.
- [22] Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. On large-batch training for deep learning: Generalization gap and sharp minima. In *International Conference on Learning Representations*, 2016.
- [23] Stephan Mandt, Matthew D. Hoffman, and David M. Blei. Stochastic gradient descent as approximate bayesian inference. *Journal of Machine Learning Research*, 18:1–35, 2017.
- [24] Daniele Musso. Stochastic gradient descent with random learning rate. *arXiv preprint arXiv:2003.06926*, 2020.
- [25] Thanh Huy Nguyen, Umut Simsekli, Mert Gürbüzbalaban, and Gaël Richard. First exit time analysis of stochastic gradient descent under heavy-tailed gradient noise. *Advances in neural information processing systems*, 32, 2019.
- [26] E. A. J. F. Peters and G. de With. Rejection-free monte carlo sampling for general potentials. *Physical Review E*, 2012.
- [27] Feng Qi and Qiu-Ming Luo. Bounds for the ratio of two gamma functions: from Wendel’s asymptotic relation to Elezović-Giordano-Pečarić’s theorem. *Journal of Inequalities and Applications*, 2013(1):542, November 2013.
- [28] Ning Qian. On the momentum term in gradient descent learning algorithms. *Neural Networks*, pages 145–151, 1999.
- [29] Maxim Raginsky, Alexander Rakhlin, and Matus Telgarsky. Non-convex learning via stochastic gradient langevin dynamics: a nonasymptotic analysis. In *Conference On Learning Theory*, 2017.

- [30] Umut Simsekli, Levent Sagun, and Mert Gurbuzbalaban. A tail-index analysis of stochastic gradient noise in deep neural networks. In *International Conference on Machine Learning*, pages 5827–5837. PMLR, 2019.
- [31] Ilya Sutskever, James Martens, George Dahl, and Geoffrey Hinton. On the importance of initialization and momentum in deep learning. In *International Conference on Machine Learning*, 2013.
- [32] Max Welling and Yee Whye Teh. Bayesian learning via stochastic gradient langevin dynamics. In *International Conference on Machine Learning*, 2011.
- [33] Zhanxing Zhu, Jingfeng Wu, Bing Yu, Lei Wu, and Jinwen Ma. The anisotropic noise in stochastic gradient descent: Its behavior of escaping from sharp minima and regularization effects. In *International Conference on Machine Learning*, 2019.
- [34] Zhanxing Zhu, Jingfeng Wu, Bing Yu, Lei Wu, and Jinwen Ma. The anisotropic noise in stochastic gradient descent: Its behavior of escaping from sharp minima and regularization effects. In *International Conference on Machine Learning*, pages 7654–7663. PMLR, 2019.

Appendix A. Supportive Information

A.1. Our Specific Contribution

We consider the empirical risk minimization problem and prove the following results under weak assumptions on the loss function such as absolute continuity: (i) the distribution of the parameters updated by Poisson SGD converges to a stationary distribution, and (ii) an output of Poisson SGD converges to the global minima of the empirical risk, applying the stationary distribution while controlling the inverse-temperature parameter. Furthermore, we evaluate the generalization error of the updated parameter for prediction with unseen data by studying an expectation of the risk function in terms of the obtained stationary distribution.

On the technical side, we utilize an algorithm called the Bouncy Particle Sampler (BPS) to demonstrate the convergence to the stationary distribution by Poisson SGD. BPS is a piecewise deterministic Markov process (PDMP) that achieves ergodicity using stochastically occurring jumps [11, 12]. In our proof, we show that the distribution of parameters updated by Poisson SGD can be well approximated by that of BPS, and we concretely construct the stationary distribution using the theory of BPS.

A.2. Notation

For a natural number $a \in \mathbb{N}$, we define $[a] := \{1, 2, \dots, a\}$. For a real $z \in \mathbb{R}$, $\lfloor z \rfloor$ denotes the largest integer which is no more than z . I_d is a d -dimensional identity matrix. $\langle a, b \rangle$ means the inner product in Euclid space, *i.e.*, sum of the product of each component. $\|\cdot\|$ and $\|\cdot\|_1$ mean the vector norms which represent 1-norm and 2-norm respectively. \mathbb{S}^{d-1} is a unit sphere in \mathbb{R}^d . For probability measures P, P' on \mathbb{R}^d and $p \in [1, \infty]$, the p -Wasserstein distance is defined as $\mathcal{W}_p(P, P') := \inf_{\pi \in \Pi(P, P')} (\int_{\mathbb{R}^d} \|z - z'\|_p^p d\pi(z, z'))^{1/p}$, where $\Pi(P, P')$ is a set of coupling measure between P and P' . $\|P - P'\|_{\text{TV}}$ denotes the total variation of $P - P'$. $\Gamma : \mathbb{R} \rightarrow \mathbb{R}$ denotes the gamma function, *i.e.*, $\Gamma(z) = \int_0^\infty t^{z-1} e^{-t} dt$. $\text{B} : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ denotes the beta function, *i.e.*, $\text{B}(x, y) = \int_0^1 t^{x-1} (1-t)^{y-1} dt$. For a compact set Θ , we denote $\text{diam}(\Theta) = \sup_{\theta_1, \theta_2 \in \Theta} \|\theta_1 - \theta_2\|$. For a random variable $X \in \mathcal{X}$, $\mathbb{E}_X[X]$ denotes the expected value with regard to X , *i.e.*, $\int_{\mathcal{X}} x d\mu_X(dx)$, where μ_X is the probability measure of X . $\mathbb{1}[\cdot]$ denotes an indicator function, which takes 1 if the condition in the bracket is satisfied, and 0 otherwise. We denote $a_+ = \max\{0, a\}$.

A.3. Related Work

There are many works which investigate the stationary distribution of SGD or its variants. [7, 14] derived the stationary distribution of the parameters obtained by SGD when the loss function is strongly-convex, through the theories about Markov processes. The parameters obtained through the SGLD algorithm are theoretically proven to converge to the Gibbs distribution and generalize well [29]. [19] and [23] assumed the noise of SGD is Gaussian whose covariance matrix is constant and approximate the process of optimization through SGD by Ornstein-Uhlenbeck process and derive its stationary distribution. Gradient Langevin dynamics (GLD), which is a full-batch version of SGLD, can also be seen as a variant of SGD which assumes that the noise of SGD is Gaussian with a covariance matrix of constant multiples of the identity matrix. Like SGLD, it converges to a stationary distribution even in non-convex scenarios [10, 15].

In terms of the random learning rate, there are several empirical studies. [24] investigated the dynamics of SGD with a random learning rate through analyzing the stochastic differential equation

Algorithm 1 Poisson SGD

-
- 1: Initialize (θ_0, v_0) as $\|v_0\| = 1$.
 - 2: **for** $k = 1, 2, \dots, K$ **do**
 - 3: Sample $I^{(k)} \subset [n]$ and obtain $\nabla \widehat{L}_{\mathbf{z}}^{(k)}(\theta_k)$ as (1).
 - 4: Sample η_k as (3).
 - 5: Obtain θ_k as $\theta_k = \theta_{k-1} + \eta_k v_{k-1}$.
 - 6: Obtain v_k as $v_k = v_{k-1} - \alpha_k \nabla \widehat{L}_{\mathbf{z}}^{(k)}(\theta_k)$ with α_k as (4).
 - 7: **end for**
 - 8: **Return** (θ_K, v_K) .
-

and its Fokker-Planck equation. [3] showed experimentally that SGD with random learning rates performs well in the optimization of deep neural networks.

As for BPS, [13, 16] proved that the parameters updated by continuous-time BPS converge to a stationary distribution and derived the concrete form of the stationary distribution and its convergence rate. [9] clarified the relation between discrete-time BPS and continuous-time BPS.

A.4. Gradient Noise

For the sake of technical discussions below, we define a notion of *gradient noise* $\xi_k^{(m,n)}(\theta) := \nabla \widehat{L}_{\mathbf{z}}^{(k)}(\theta) - \nabla L_{\mathbf{z}}(\theta)$ for $k = 1, \dots, N$ and $\theta \in \Theta$, which is caused by the sub-sampling of the SGD. If one assumes that $\xi_k^{(m,n)}(\theta)$ follows a centered Gaussian distribution with an identity covariance, the SGD corresponds to the gradient Langevin dynamics (GLD). However, since it is empirically observed that it is not realistic [8, 18, 33], we do not assume that the gradient noise follows a particular distribution.

A.5. Pseudo Code of Poisson SGD

Pseudo code of Poisson SGD is shown in Algorithm 1.

A.6. Intuition of Poisson SGD

The algorithm is designed to effectively explore large regions of the parameter space Θ . Specifically, the update direction is determined by the velocity vector v_k normalised by α_k as (4), and the size of the update is randomly set by the random learning rate η_k as (3). When the gradient $\nabla \widehat{L}_{\mathbf{z}}^{(k)}(\cdot)$ is small, the learning rate η_k is chosen to be large, thus the updated parameter tends to escape local minima or saddle points. Figure 1 illustrates that Poisson SGD explores a wider parameter space and discovers the global minimum owing to the random learning rate, while the parameters updated by SGD converge to the local minimum.

A.7. Discussion in Theorem 2

The right-hand side in (6) shows an approximation-complexity trade-off of Poisson SGD described as follows. In preparation, we will introduce a certain stochastic process to achieve the stationary distribution $\mu_{\mathbf{z}}^{(\beta, \varepsilon)}$ (detail is in Section D). The first term of (6) describes an approximation error




Figure 1: The comparison of the trajectories of SGD (with a fixed learning rate) and Poisson SGD (with the random learning rate) in optimizing the function $z = x^4 - 4x^3 - 36x^2 + y^2$. The point $(-3, 0)$ represents a local minimum and the point $(6, 0)$ is identified as the global minimum. A green point indicates the initial position, a black line represents the trajectory of SGD, and a red line represents the trajectory of Poisson SGD.

of Poisson SGD to the stochastic process. The second term of (6) denotes a convergence error of the stochastic process to the stationary distribution $\mu_{\mathbf{z}}^{(\beta, \varepsilon)}$, which reflects the complexity of the stochastic process. ε is a parameter for the stochastic process and controls the balance between the approximation error and the complexity error.

We further discuss the additional assumption $\lim_{K \rightarrow \infty} \kappa(\beta, \delta/K, d)^K = 0$. This condition is related to the convergence rate of the approximated stochastic process of Poisson SGD. Although the explicit form of $\kappa(\beta, \delta/K, d)$ is not clarified in our case, there is a common example having its explicit form. One example is SGLD: [29] shows that a form of $\kappa(\beta, \delta/K, d)$ can be calculated, because SGLD is reduced to the Langevin process.

Remark 4 (Comparison with SGLD) *We discuss the difference between Poisson SGD and SGLD, which is another method achieving a stationary distribution. First, while SGLD adds a Gaussian noise to the update formula of SGD, Poisson SGD does not have an additive noise. The second difference is the form of the stationary distribution. A stationary distribution of SGLD is the Gibbs distribution, and that of Poisson SGD has the different form (5). This difference is derived from the random learning rate of Poisson SGD.*

Appendix B. Generalization Error Analysis

We define an expected risk of $\theta \in \Theta$, also known as the generalization error

$$L(\theta) := \mathbb{E}_{z \sim P_*} [\ell(z; \theta)],$$

which measures a prediction performance with unseen data. We calculate the generalization error of the parameter obtained by the Poisson SGD, using the discussion in [29].

Now, we give our results. We define $A := \sup_{z \in \mathcal{Z}} |\ell(z; 0)|$ by following Assumption 1.

Theorem 5 (Generalization Error of Poisson SGD) *Suppose that Assumption 1 and 2 hold. Let θ_K be the parameter obtained by Poisson SGD with $C_P = 1/\varepsilon$. Then, we obtain the following bound:*

$$\begin{aligned} & \mathbb{E}_{\mathbf{z} \sim P_*^n} [\mathbb{E}_{\theta_K \sim \mu_{\mathbf{z}, K}} [L(\theta_K)]] - \min_{\theta \in \Theta} L(\theta) \\ & \leq (c_1 W + B) \left(\sqrt{W d_K(\beta, \varepsilon, d)} + 2W \left(\left(\frac{C_d + \beta C}{n} \right)^{\frac{1}{2}} + \left(\frac{C_d + \beta C}{n} \right)^{\frac{1}{4}} \right) \right) \\ & \quad + \frac{1}{\beta} \left(\frac{d}{2} \log \frac{e W^2 c_1 \beta}{d} + \log \left(1 + \frac{a_d (c_1 W + B)}{M_\ell} \right) \right), \end{aligned}$$

where $d_K(\beta, \varepsilon, d)$ is the upper bound of the Wasserstein distance in Theorem 2, $C_d = 4a_d(c_1 W + B)/M_\ell$, and $C = c_1 W^2 + 2BW + 2A$.

Theorem 5 states that the expected value of the generalization error of Poisson SGD can be arbitrarily close to its global optima in $\theta \in \Theta$, by selecting small ε , large K , large β , and large n , provided that $d_K(\beta, \varepsilon, d)$ can be arbitrarily small only by the choice of ε and K .

We further discuss a way of improve an order of the generalization bound in Theorem 5. While our bound has the order $O((1/n)^{1/4})$, we can obtain an order $O(1/n)$ by using the *dissipativity* condition of the loss function, which is used in [29] for SGLD. The dissipativity condition allows us to derive log-Sobolev inequality for $L_{\mathbf{z}}(\theta)$, which leads the improved sample complexity. We state this fact in the following proposition.

Proposition 6 *Suppose that the same condition and setting as Theorem 5 hold. In addition, we assume that the Gibbs distribution $v_{\mathbf{z}}^{(\beta)} \propto \exp(-\beta L_{\mathbf{z}}(\theta))$ satisfies the log-Sobolev inequality for any dataset $\mathbf{z} = \{z_1, \dots, z_n\}$, that is,*

$$\mathbb{E}[f(\theta)^2 \log f(\theta)^2] - \mathbb{E}[f(\theta)^2] \log \mathbb{E}[f(\theta)^2] \leq c_{\text{LS}}^{(\beta)} \mathbb{E}[\|\nabla f(\theta)\|^2]$$

holds for all smooth functions f and any data $\mathbf{z} = \{z_1, \dots, z_n\}$, where $\theta \sim v_{\mathbf{z}}^{(\beta)}$ and $c_{\text{LS}}^{(\beta)} < \infty$ is a constant. Then, the following holds:

$$\begin{aligned} & \mathbb{E}_{\mathbf{z} \sim P_*^n} [\mathbb{E}_{\theta_K \sim \mu_{\mathbf{z}, K}} [L(\theta_K)]] - \min_{\theta \in \Theta} L(\theta) \\ & \leq (c_1 W + B) \left(\sqrt{W d_K(\beta, \varepsilon, d)} + \frac{2c_{\text{LS}}^{(\beta)} \beta M_\ell}{n} \right) + W \sqrt{2c_{\text{LS}}^{(\beta)} \log(1 + a_d \beta \varepsilon M_\ell)} + \frac{d}{2\beta} \log \left(\frac{e W^2 c_1 \beta}{d} \right). \end{aligned}$$

Appendix C. Auxiliary Propositions

We verify that the velocity vector is normalized by the choice of the momentum coefficient for Poisson SGD and BPS.

Proposition 7 *Consider the update (2) for v_k with its momentum coefficient (4). Then, for $\forall k \in \{1, 2, \dots, K\}$, we have $\|v_k\| = 1$. Further, for \widehat{v}_k defined in Algorithm 2, we obtain $\|\widehat{v}_k\| = 1$ for every $k = 1, \dots, K$.*

Proof We first consider v_k with the Poisson SGD case. Simply, we have

$$\begin{aligned}
\|v_k\| &= \left\| v_{k-1} - 2 \frac{\langle \nabla \widehat{L}_{\mathbf{z}}^{(k)}(\theta_k), v_{k-1} \rangle}{\|\nabla \widehat{L}_{\mathbf{z}}^{(k)}(\theta_k)\|^2} \nabla \widehat{L}_{\mathbf{z}}^{(k)}(\theta_k) \right\| \\
&= \left\| \left(I_d - 2 \frac{\nabla \widehat{L}_{\mathbf{z}}^{(k)}(\theta_k) \nabla \widehat{L}_{\mathbf{z}}^{(k)}(\theta_k)^\top}{\|\nabla \widehat{L}_{\mathbf{z}}^{(k)}(\theta_k)\|^2} \right) v_{k-1} \right\| \\
&= \sqrt{v_{k-1}^\top \left(I_d - 2 \frac{\nabla \widehat{L}_{\mathbf{z}}^{(k)}(\theta_k) \nabla \widehat{L}_{\mathbf{z}}^{(k)}(\theta_k)^\top}{\|\nabla \widehat{L}_{\mathbf{z}}^{(k)}(\theta_k)\|^2} \right)^2 v_{k-1}} \\
&= \|v_{k-1}\|.
\end{aligned}$$

Since we set $\|v_0\| = 1$ for initialization, the statement holds.

For \widehat{v}_k with the BPS case, the reflection does not change the norm of \widehat{v}_k in the same way, and the refreshment also keeps $\|\widehat{v}_k\| = 1$, which completes the proof. \blacksquare

Appendix D. Proof Outline

D.1. Overview

We give an overview of a proof of Theorem 2. In preparation, we present several key concepts: (i) the property of the *piece-wise deterministic Markov process* (PDMP) [11, 12], and (ii) the ergodicity of *bouncy particle sampler* (BPS) [26]. The PDMP is a class of Markov processes that behave deterministically for some period and jumps randomly, which easily converges to a stationary distribution. BPS is a stochastic algorithm in the class of the PDMP.

We show the statement by the following steps:

- (I) We show that the distribution of the parameter by Poisson SGD is sufficiently close to that of a parameter by BPS. We show this claim by using the approximation theory on PDMP (Theorem 8).
- (II) We derive a stationary distribution and the ergodicity of BPS, following previous researches (Theorem 10).

D.2. Design of BPS

We introduce BPS, which is one of the most popular algorithms in PDMPs, and actively studied in terms of MCMC algorithm [6, 13]. BPS generates a sequence of parameters $\{\widehat{\theta}_k\}_{k=1}^K \subset \Theta$ and velocity vectors $\{\widehat{v}_k\}_{k=1}^K \subset \mathbb{R}^d$ in its recursive manner, as shown in Algorithm 2. Let $(\widehat{\theta}_0, \widehat{v}_0)$ be the initialization. For the k -th iteration, BPS generates a learning rate $\widehat{\eta}_k$ from an exponential distribution whose intensity depends on the previous pair $(\widehat{\theta}_{k-1}, \widehat{v}_{k-1})$ and the positive constants Λ_{ref} and C_B . After obtaining the parameter $\widehat{\theta}_k$, we consider the stochastic update of the velocity vector. That is, with the probability

$$\widehat{p}_k := \frac{\beta \langle \nabla L_{\mathbf{z}}(\widehat{\theta}_k), \widehat{v}_{k-1} \rangle_+ + C_B}{\beta \langle \nabla L_{\mathbf{z}}(\widehat{\theta}_k), \widehat{v}_{k-1} \rangle_+ + \Lambda_{\text{ref}} + C_B}, \quad (7)$$

Algorithm 2 Bouncy Particle Sampler

-
- 1: Initialize $(\widehat{\theta}_0, \widehat{v}_0)$ as $\|\widehat{v}_0\| = 1$.
 - 2: **for** $k = 1, 2, \dots, K$ **do**
 - 3: Sample $\widehat{\eta}_k$ as $\widehat{\eta}_k \sim P(\widehat{\eta}_k \geq t) = \exp\left(-\int_0^t \{\beta \langle \nabla L_{\mathbf{z}}(\widehat{\theta}_{k-1} + r\widehat{v}_{k-1}), \widehat{v}_{k-1} \rangle_+ + \Lambda_{\text{ref}} + C_B\} dr\right)$
 - 4: Update $\widehat{\theta}_k$ as $\widehat{\theta}_k = \widehat{\theta}_{k-1} + \widehat{\eta}_k \widehat{v}_{k-1}$
 - 5: With probability \widehat{p}_k as (7), update \widehat{v}_k as

$$\widehat{v}_k = \widehat{v}_{k-1} - 2 \frac{\langle \nabla L_{\mathbf{z}}(\widehat{\theta}_k), \widehat{v}_{k-1} \rangle}{\|\nabla L_{\mathbf{z}}(\widehat{\theta}_k)\|^2} \nabla L_{\mathbf{z}}(\widehat{\theta}_k)$$

Otherwise, update \widehat{v}_k as

$$\widehat{v}_k \sim \text{Unif}(\mathbb{S}^{d-1})$$

- 6: **end for**
 - 7: Return $(\widehat{\theta}_K, \widehat{v}_K)$
-

we update the velocity vector with the gradient of the full-batch loss $\nabla L_{\mathbf{z}}$, otherwise with the sample from the uniform distribution on \mathbb{S}^{d-1} . The former update is called *reflection*, and the latter is *refreshment*. We remark that $\|\widehat{v}_k\|$ is constant for $k = 1, 2, \dots, K$ in the same way as Poisson SGD (See Proposition 7 in Appendix).

D.3. Connect Poisson SGD and BPS

We show that the output distribution of Poisson SGD and that of BPS are sufficiently close as in the following statement:

Theorem 8 (Distance between Poisson SGD and BPS) *Fix arbitrary $\beta, \varepsilon > 0$. As for Poisson SGD, we set $C_P = 1/\varepsilon$. As for BPS, we set Λ_{ref} and C_B as $\Lambda_{\text{ref}} + C_B = \beta M_\ell + 1/\varepsilon$. Let the distribution of the obtained parameter by Poisson SGD and BPS be $\mu_{z,K}$ and $\widehat{\mu}_{z,K}$ respectively. We set the same initial value between Poisson SGD and BPS. Then, the following holds:*

$$\mathcal{W}_1(\mu_{z,K}, \widehat{\mu}_{z,K}) \leq 4\sqrt{d}K\varepsilon.$$

For proving this theorem, we calculate the distance between Poisson SGD and BPS by a one-step update. Then, we simply accumulate this error for K times. In this discussion, we mainly use the property that if learning rate η_k and $\widehat{\eta}_k$ are small, the difference of v_k and \widehat{v}_k is also made to be small. This type of discussion is also used in [29].

D.4. The Stationary Distribution and Ergodicity of BPS

In this section, we investigate the stationary distribution and ergodicity of BPS. First, we define the term *ergodicity*.

Definition 9 (Ergodicity) *We consider the discrete-time Markov process. If the process converges to a unique stationary distribution, we call the process has the ergodicity. Especially, if the ergodic process converges to its stationary distribution by the exponential rate about the number of iteration k , the process is called exponentially ergodic.*

Without ergodicity, the stochastic process may converge to more than one stationary distribution, or not converge to any stationary distribution due to stacking to a saddle point in the parameter space. So we have to prove this property when we try to analyze the stationary distribution of a stochastic process.

Now, we show our result about BPS.

Theorem 10 (Stationary Distribution of BPS) *Suppose that Assumption 1 holds. Set the parameter of BPS, Λ_{ref} and C_B as in Theorem 8. Then, the distribution $\widehat{\mu}_{z,K}$ of the obtained parameters $\widehat{\theta}_K$ by BPS satisfies the following:*

$$\|\widehat{\mu}_{z,K} - \mu_z^{(\beta,\varepsilon)}\|_{\text{TV}} \leq \kappa(\beta, \varepsilon, d)^K,$$

where $\kappa(\beta, \varepsilon, d)$ is a positive constant less than 1.

In the proof of this theorem, we use the discussion in [13] which showed that continuous-time BPS converges to the unique stationary distribution $\pi(\theta) \propto \exp(-U(\theta))$ by the exponential rate in TV distance.

Appendix E. Proof of Theorem 2

Proof By Theorem 8 and 10, we can bound the approximation error

$$\mathcal{W}_1(\mu_{z,K}, \widehat{\mu}_{z,K}) \leq 4\sqrt{d}K\varepsilon,$$

and the convergence error of BPS as

$$\|\widehat{\mu}_{z,K} - \mu_z^{(\beta,\varepsilon)}\|_{\text{TV}} \leq \kappa(\beta, \varepsilon, d)^K.$$

From Theorem 4 in [17], we can bound the Wasserstein distance by the total variation, then obtain

$$\mathcal{W}_1(\widehat{\mu}_{z,K}, \mu_z^{(\beta,\varepsilon)}) \leq W\|\widehat{\mu}_{z,K} - \mu_z^{(\beta,\varepsilon)}\|_{\text{TV}} \leq W\kappa(\beta, \varepsilon, d)^K.$$

The triangle inequality completes the proof. ■

Appendix F. Proof of Theorem 8

Proof From the definition of Wasserstein distance,

$$\mathcal{W}_1(\mu_{z,k}, \widehat{\mu}_{z,k}) = \inf_{\pi \in \Pi(\mu_{z,k}, \widehat{\mu}_{z,k})} \mathbb{E}_\pi[\|\theta_k - \widehat{\theta}_k\|_1]$$

holds, so we study the distance between θ_k and $\widehat{\theta}_k$ in terms of the norm $\|\cdot\|_1$. Since $\|v_k\| = \|v_{k-1}\| = \|\widehat{v}_k\| = \|\widehat{v}_{k-1}\| = 1$ holds by Proposition 7, we have

$$\begin{aligned} \mathbb{E}_\pi[\|\theta_k - \widehat{\theta}_k\|_1] &= \mathbb{E}_\pi[\|\theta_{k-1} + \eta_k v_{k-1} - (\widehat{\theta}_{k-1} + \widehat{\eta}_k \widehat{v}_{k-1})\|_1] \\ &\leq \mathbb{E}_\pi[\|\theta_{k-1} - \widehat{\theta}_{k-1}\|_1] + \mathbb{E}_\pi[\|(\widehat{\eta}_k - \eta_k)\widehat{v}_{k-1} + \eta_k(\widehat{v}_{k-1} - v_{k-1})\|_1] \\ &\leq \mathbb{E}_\pi[\|\theta_{k-1} - \widehat{\theta}_{k-1}\|_1] + \mathbb{E}_\pi[\|(\widehat{\eta}_k - \eta_k)\widehat{v}_{k-1}\|_1] + \mathbb{E}_\pi[\|\eta_k(\widehat{v}_{k-1} - v_{k-1})\|_1] \end{aligned}$$

$$\leq \mathbb{E}_\pi[\|\theta_{k-1} - \widehat{\theta}_{k-1}\|_1] + \sqrt{d}\mathbb{E}_\pi[|\eta_k - \widehat{\eta}_k|] + 2\sqrt{d}\mathbb{E}_\pi[\eta_k], \quad (8)$$

where we use $\|\cdot\|_1 \leq \sqrt{d}\|\cdot\|$ in the last inequality.

We first evaluate the second term of (8). There exists a coupling π such that

$$\mathbb{E}_\pi[|\eta_k - \widehat{\eta}_k|] = \mathcal{W}_1(P_{\eta_k}, P_{\widehat{\eta}_k})$$

holds, where P_{η_k} and $P_{\widehat{\eta}_k}$ denote the distribution of η_k and $\widehat{\eta}_k$ respectively. We use such a coupling as π . In evaluating $\mathcal{W}_1(P_{\eta_k}, P_{\widehat{\eta}_k})$, we consider the following analysis. η_k and $\widehat{\eta}_k$ are 1-dimensional and their cumulative distribution function is written as

$$F_1(t) = 1 - \exp\left(-\int_0^t (\beta\langle \nabla \widehat{L}_{\mathbf{z}}^{(k)}(\theta + rv), v \rangle_+ + C_P) dr\right),$$

$$F_2(t) = 1 - \exp\left(-\int_0^t (\beta\langle \nabla L_{\mathbf{z}}(\theta + rv), v \rangle_+ + C_B + \Lambda_{\text{ref}}) dr\right),$$

respectively, and we also have

$$\begin{aligned} \beta\langle \nabla \widehat{L}_{\mathbf{z}}^{(k)}(\theta + rv), v \rangle_+ + C_P &\geq C_P, \\ \beta\langle \nabla L_{\mathbf{z}}(\theta + rv), v \rangle_+ + C_B + \Lambda_{\text{ref}} &\geq C_B + \Lambda_{\text{ref}}, \text{ and} \\ |(\beta\langle \nabla \widehat{L}_{\mathbf{z}}^{(k)}(\theta + rv), v \rangle_+ + C_P) - (\beta\langle \nabla L_{\mathbf{z}}(\theta + rv), v \rangle_+ + C_B + \Lambda_{\text{ref}})| \\ &\leq \max\{|-\beta M_\ell + C_P - (C_B + \Lambda_{\text{ref}})|, |\beta M_\ell + C_P - (C_B + \Lambda_{\text{ref}})|\}. \end{aligned}$$

Hence, we can use Lemma 11 and obtain

$$\mathcal{W}_1(P_{\eta_k}, P_{\widehat{\eta}_k}) \leq \frac{\max\{|-\beta M_\ell + C_P - (C_B + \Lambda_{\text{ref}})|, |\beta M_\ell + C_P - (C_B + \Lambda_{\text{ref}})|\}}{C_P(C_B + \Lambda_{\text{ref}})}. \quad (9)$$

Next, we evaluate the third term of (8). We have

$$\begin{aligned} \mathbb{E}[\eta_k] &= \int_0^\infty P(\eta_k \geq t) dt \\ &= \int_0^\infty \exp\left(-\int_0^t \{\beta\langle \nabla \widehat{L}_{\mathbf{z}}^{(k)}(\theta_{k-1} + rv_{k-1}), v_{k-1} \rangle_+ + C_P\} dr\right) dt \\ &\leq \int_0^\infty \exp(-C_P t) dt \\ &= \frac{1}{C_P}. \end{aligned} \quad (10)$$

Substituting (9) and (10) into (8), we have

$$\begin{aligned} \mathbb{E}_\pi[\|\theta_k - \widehat{\theta}_k\|_1] &\leq \mathbb{E}_\pi[\|\theta_{k-1} - \widehat{\theta}_{k-1}\|_1] \\ &\quad + \frac{\sqrt{d} \max\{|-\beta M_\ell + C_P - (C_B + \Lambda_{\text{ref}})|, |\beta M_\ell + C_P - (C_B + \Lambda_{\text{ref}})|\}}{C_P(C_B + \Lambda_{\text{ref}})} + \frac{2\sqrt{d}}{C_P}. \end{aligned} \quad (11)$$

Since we take C_P in Poisson SGD as $C_P = 1/\varepsilon$ and C_B and Λ_{ref} in BPS as $C_B + \Lambda_{\text{ref}} = \beta M_\ell + 1/\varepsilon$, (11) can be written as

$$\mathbb{E}_\pi[\|\theta_k - \widehat{\theta}_k\|_1] \leq \mathbb{E}_\pi[\|\theta_{k-1} - \widehat{\theta}_{k-1}\|_1] + 4\sqrt{d}\varepsilon.$$

Hence, solving this recursive inequality with $\theta_0 = \widehat{\theta}_0$, we have

$$\mathbb{E}_\pi[\|\theta_K - \widehat{\theta}_K\|_1] \leq 4\sqrt{d}K\varepsilon,$$

which is the desired conclusion. ■

Lemma 11 *Let a_1 and a_2 be \mathbb{R} -valued random variables whose cumulative distribution functions are*

$$F_1(t) = 1 - \exp\left(-\int_0^t f_1(r)dr\right), \text{ and } F_2(t) = 1 - \exp\left(-\int_0^t f_2(r)dr\right),$$

respectively, where $f_1, f_2 : \mathbb{R} \rightarrow \mathbb{R}$ are continuous functions. Let the distributions of a_1 and a_2 be P_1 and P_2 respectively. Suppose that there exists $M, m_1, m_2 > 0$ such that $|f_2(t) - f_1(t)| \leq M$, $m_1 \leq f_1(t)$, and $m_2 \leq f_2(t)$ hold for $\forall t \in \mathbb{R}$. Then, the Wasserstein distance between P_1 and P_2 satisfies

$$\mathcal{W}_1(P_1, P_2) \leq \frac{M}{m_1 m_2}.$$

Proof Since a_1 and a_2 are 1-dimensional, we have

$$\mathcal{W}_1(P_1, P_2) = \int_0^1 |F_1^{-1}(q) - F_2^{-1}(q)| dq.$$

We introduce several notation $\delta(r) = f_2(r) - f_1(r)$, $t = F_1^{-1}(q)$, and $t' = F_2^{-1}(q)$, then

$$\begin{aligned} \int_0^t f_1(r)dr &= \log \frac{1}{1-q} \\ \int_0^{t'} (f_1(r) + \delta(r))dr &= \log \frac{1}{1-q} \end{aligned}$$

holds. So, we obtain

$$\int_{t'}^t f_1(r)dr = \int_0^{t'} \delta(r)dr.$$

Hence, we have

$$\left| \int_{t'}^t f_1(r)dr \right| = \int_{\min\{t, t'\}}^{\max\{t, t'\}} f_1(r)dr \leq M t'.$$

In addition, $\int_{\min\{t, t'\}}^{\max\{t, t'\}} f_1(r)dr \geq m_1 |t - t'|$ holds, so we have

$$|t - t'| \leq \frac{M t'}{m_1}.$$

We have the upper bound of t' as

$$\log \frac{1}{1-q} = \int_0^{t'} f_2(r)dr \geq m_2 t',$$

so we have

$$|t - t'| \leq \frac{M}{m_1 m_2} \log \frac{1}{1 - q}.$$

Since $\int_0^1 |\log(1 - q)| dq = 1$ holds, we obtain

$$\int_0^1 |F_1^{-1}(q) - F_2^{-1}(q)| dq \leq \frac{M}{m_1 m_2}.$$

■

Appendix G. Proof of Theorem 10

We prove this theorem by two steps. First, we prove that BPS has $\mu_{\mathbf{z}}^{(\beta, \varepsilon)}$ as one of its stationary distributions in section G.1. At this stage, BPS may have other forms of stationary distribution or may not converge to its stationary distribution. Second, we prove that BPS has a unique stationary distribution and converges to its stationary distribution at exponential rate, in other words, it has the exponential ergodicity, in section G.2.

G.1. The form of the stationary distribution

In this section, we check that BPS has $\mu_{\mathbf{z}}^{(\beta, \varepsilon)}$ as a stationary distribution. In the proof, we define $\lambda(\theta, v) := \beta \langle \nabla L_{\mathbf{z}}(\theta), v \rangle_+$, $\bar{\lambda}(\theta, v) := \lambda(\theta, v) + \Lambda_{\text{ref}}$, and $R_{\mathbf{z}}(\theta) := I_d - 2 \frac{\nabla L_{\mathbf{z}}(\theta) \nabla L_{\mathbf{z}}(\theta)^\top}{\|\nabla L_{\mathbf{z}}(\theta)\|^2}$. We remark that $R_{\mathbf{z}}$ is a symmetric matrix and satisfies $R_{\mathbf{z}}(\theta)^2 = I_d$, so it is also an orthogonal matrix.

From the proof of Lemma 1 in the supplementary material of [13], we can write the transition probability \widehat{Q} of BPS as following for arbitrary measurable sets $A \subset \Theta$ and $B \subset \mathbb{S}^{d-1}$:

$$\begin{aligned} \widehat{Q}((\theta, v), A \times B) &= \int_0^\infty \exp \left\{ - \int_0^s (\bar{\lambda}(\theta + uv, v) + C_B) du \right\} \\ &\quad \times (\bar{\lambda}(\theta + uv, v) + C_B) K((\theta + sv, v), A \times B) ds, \end{aligned} \quad (12)$$

where a transition kernel K is expressed as

$$\begin{aligned} K((\theta, v), A \times B) &= \frac{\lambda(\theta, v) + C_B}{\bar{\lambda}(\theta, v) + C_B} \mathbb{1}[\theta \in A] \mathbb{1}[R_{\mathbf{z}}(\theta)v \in B] \\ &\quad + \frac{\Lambda_{\text{ref}}}{\bar{\lambda}(\theta, v) + C_B} \mathbb{1}[\theta \in A] \mu_{\text{unif}}(B), \end{aligned} \quad (13)$$

where μ_{unif} is the uniform probability measure on \mathbb{S}^{d-1} .

Lemma 12 *Under Assumption 1, a probability measure on $\Theta \times \mathbb{S}^{d-1}$*

$$\widehat{\mu}_{\mathbf{z}}(A \times B) \propto \int_{A \times B} (\bar{\lambda}(\theta, -v) + C_B) \exp(-\beta L_{\mathbf{z}}(\theta)) d\theta \mu_{\text{unif}}(dv)$$

is the stationary distribution induced from the transition probability \widehat{Q} as (12).

Proof Our proof is almost the same as the proof of Lemma 1 in [13]. Let $\pi_{\mathbf{z}}(d\theta, dv) = \exp(-\beta L_{\mathbf{z}}(\theta))d\theta\mu_{\text{unif}}(dv)$.

First, we prove

$$\int (\bar{\lambda}(\theta, v) + C_B)\pi_{\mathbf{z}}(d\theta, dv)K((\theta, v), A \times B) \propto \widehat{\mu}_{\mathbf{z}}(A \times B). \quad (14)$$

Substituting (13), the left side of (14) is rewritten as

$$\int \pi_{\mathbf{z}}(d\theta, dv)(\lambda(\theta, v) + C_B)\mathbb{1}[\theta \in A]\mathbb{1}[R_{\mathbf{z}}(\theta)v \in B] + \int \pi_{\mathbf{z}}(d\theta, dv)\Lambda_{\text{ref}}\mathbb{1}[\theta \in A]\mu_{\text{unif}}(B).$$

We consider changing the variable as $v' = R_{\mathbf{z}}(\theta)v$. Since $R_{\mathbf{z}}(\theta)^{-1} = R_{\mathbf{z}}(\theta)$ holds, we get $\lambda(\theta, R_{\mathbf{z}}(\theta)^{-1}v') = \lambda(\theta, -v')$. In addition, since $|\det(R_{\mathbf{z}}(\theta))| = 1$, and $\mu_{\text{unif}}(R_{\mathbf{z}}(\theta)^{-1}dv') = \mu_{\text{unif}}(dv')$ hold due to the rotational invariance of μ_{unif} , we obtain

$$\int_{A \times B} \pi_{\mathbf{z}}(d\theta, dv')(\lambda(\theta, -v') + C_B) + \int_{A \times B} \pi_{\mathbf{z}}(d\theta, dv')\Lambda_{\text{ref}},$$

which is proportional to the right side of (14) from the definition of $\widehat{\mu}_{\mathbf{z}}$.

Second, we prove $\int \widehat{Q}((\theta, v), (dy, dw))\widehat{\mu}_{\mathbf{z}}(d\theta, dv) = \widehat{\mu}_{\mathbf{z}}(dy, dw)$. We have

$$\begin{aligned} & \int \widehat{Q}((\theta, v), (dy, dw))\widehat{\mu}_{\mathbf{z}}(d\theta, dv) \\ & \propto \int_0^\infty \exp\left(-\int_0^s \{\bar{\lambda}(\theta + uv, v) + C_B\}du\right) \{\bar{\lambda}(\theta + sv, v) + C_B\} \\ & \quad \times K((\theta + sv, v), (dy, dw))\{\bar{\lambda}(\theta, -v) + C_B\}\pi_{\mathbf{z}}(d\theta, dv)ds. \end{aligned}$$

If we change θ as $t = \theta + sv$, then this integral becomes

$$\begin{aligned} & \int_0^\infty \exp\left(-\int_0^s \{\bar{\lambda}(t + (u-s)v, v) + C_B\}du\right) \{\bar{\lambda}(t, v) + C_B\} \\ & \quad \times K((t, v), (dy, dw))\{\bar{\lambda}(t - sv, -v) + C_B\}\pi_{\mathbf{z}}(d\theta, dv)ds. \end{aligned}$$

Since $L_{\mathbf{z}}(\theta)$ is absolutely continuous,

$$\exp(-\beta L_{\mathbf{z}}(t - sv)) = \exp\left(-\beta L_{\mathbf{z}}(t) - \int_0^s \lambda(t - wv, -v)dw + \int_0^s \lambda(t - wv, v)dw\right)$$

holds in the same way as [13]. Substituting it into $\pi_{\mathbf{z}}(dx, dv)$ and changing u as $u - s = -w$,

$$\begin{aligned} & \int_0^\infty \exp\left(-\int_0^s \{\bar{\lambda}(t - wv, -v) + C_B\}dw\right) \{\bar{\lambda}(t - sv, -v) + C_B\}ds \\ & \quad \times \{\bar{\lambda}(t, v) + C_B\}K((t, v), (dy, dw))\pi_{\mathbf{z}}(dt, dv) \end{aligned}$$

holds. The first line can be calculated as $\left[-\exp\left(-\int_0^s \{\bar{\lambda}(t-wv, -v) + C_B\}dw\right)\right]_0^\infty = 1$, so it is equal to

$$\int \{\bar{\lambda}(t, v) + C_B\}K((t, v), (dy, dw))\pi_{\mathbf{z}}(dt, dv).$$

Using (14), it is proportional to $\widehat{\mu}_{\mathbf{z}}(dy, dw)$, which completes the proof. \blacksquare

By the following proposition, we prove that $\mu_{\mathbf{z}}^{(\beta, \varepsilon)}$ is one of the stationary distributions of BPS. Recall that we defined $a_d := \Gamma(d/2)/(\sqrt{\pi}\Gamma(d/2 + 1/2))$.

Proposition 13 *The marginal distribution of the stationary distribution expressed in Lemma 12 is written as*

$$\widehat{\mu}_{\mathbf{z}}(d\theta) \propto (\Lambda_{\text{ref}} + C_B + a_d\beta\|\nabla L_{\mathbf{z}}(\theta)\|) \exp(-\beta L_{\mathbf{z}}(\theta))d\theta.$$

Hence, if we put Λ_{ref} and C_B as $\Lambda_{\text{ref}} + C_B = \beta M_\ell + 1/\varepsilon$, it corresponds to $\mu_{\mathbf{z}}^{(\beta, \varepsilon)}$.

Proof We only need to integrate with v the distribution $\widehat{\mu}_{\mathbf{z}}$ expressed in Lemma 12. We have

$$\begin{aligned} \widehat{\mu}_{\mathbf{z}}(d\theta) &\propto \int_{v \in \mathbb{S}^{d-1}} (\Lambda_{\text{ref}} + C_B + \beta\langle \nabla L_{\mathbf{z}}(\theta), -v \rangle_+) \exp(-\beta L_{\mathbf{z}}(\theta))d\theta \mu_{\text{unif}}(dv) \\ &= (\Lambda_{\text{ref}} + C_B) \exp(-\beta L_{\mathbf{z}}(\theta))d\theta + \exp(-\beta L_{\mathbf{z}}(\theta))d\theta \beta \mathbb{E}_{v \sim \mu_{\text{unif}}}[\langle \nabla L_{\mathbf{z}}(\theta), v \rangle_+]. \end{aligned}$$

We can calculate the expected value in the last term as

$$\mathbb{E}_{v \sim \mu_{\text{unif}}}[\langle \nabla L_{\mathbf{z}}(\theta), v \rangle_+] = \mathbb{E}_{v \sim \mu_{\text{unif}}}[\|\nabla L_{\mathbf{z}}(\theta)\|(\cos \phi)_+] = \|\nabla L_{\mathbf{z}}(\theta)\| \mathbb{E}_{v \sim \mu_{\text{unif}}}[(\cos \phi)_+],$$

where $\phi \in \mathbb{R}$ is a random variable dependent on v which satisfies

$$\cos \phi = \left\langle \frac{\nabla L_{\mathbf{z}}(\theta)}{\|\nabla L_{\mathbf{z}}(\theta)\|}, v \right\rangle. \quad (15)$$

From the symmetry of the uniform distribution, we can calculate $\mathbb{E}_{v \sim \mu_{\text{unif}}}[(\cos \phi)_+]$ by replacing $\frac{\nabla L_{\mathbf{z}}(\theta)}{\|\nabla L_{\mathbf{z}}(\theta)\|}$ in (15) by $(1, 0, \dots, 0)$. Hence,

$$\mathbb{E}_{v \sim \mu_{\text{unif}}}[(\cos \phi)_+] = \mathbb{E}_{v \sim \mu_{\text{unif}}}[(v_1)_+] = \mathbb{E}\left[\left(\frac{x_1}{\sqrt{x_1^2 + \dots + x_d^2}}\right)_+\right]$$

holds, where v_1 is the first component of v and $x_i (i = 1, \dots, d)$ is *i.i.d.* standard Gaussian variables.

For $(x_1, \dots, x_d) \sim \mathcal{N}(\mathbf{0}, I_d)$, we have

$$\begin{aligned} \mathbb{E}\left[\sqrt{\frac{x_1^2}{x_1^2 + \dots + x_d^2}}\right] &= \int_{\mathbb{R}^d} \sqrt{\frac{z_1^2}{z_1^2 + \dots + z_d^2}} \frac{1}{(2\pi)^{d/2}} \exp\left(-\frac{z_1^2 + \dots + z_d^2}{2}\right) dz_1 \dots dz_d \\ &= \int_{[0, \infty)^2} \sqrt{\frac{r}{r+s}} \frac{r^{-1/2} \exp(-r/2)}{\sqrt{2\pi}} \frac{s^{(d-1)/2-1} \exp(-s/2)}{\Gamma((d-1)/2)2^{(d-1)/2}} dr ds \\ &= \int_{[0, 1]} \frac{t^{1/2} t^{1/2-1} (1-t)^{(d-1)/2-1}}{\mathbf{B}(1/2, (d-1)/2)} dt \end{aligned}$$

$$\begin{aligned}
 &= \frac{\mathbf{B}(1, (d-1)/2)}{\mathbf{B}(1/2, (d-1)/2)} \\
 &= \frac{\Gamma(1)\Gamma((d-1)/2)\Gamma(d/2)}{\Gamma(1/2)\Gamma((d-1)/2)\Gamma(d/2+1/2)} \\
 &= \frac{\Gamma(d/2)}{\sqrt{\pi}\Gamma(d/2+1/2)}.
 \end{aligned}$$

Note that for all $d \geq 2$,

$$\frac{1}{\sqrt{d/2}} \leq \frac{\Gamma(d/2)}{\Gamma(d/2+1/2)} \leq \frac{1}{\sqrt{d/2-1/2}}$$

holds (e.g., see [27]). Therefore, for all $d \geq 2$, we have

$$\mathbb{E} \left[\left(\frac{x_1}{\sqrt{x_1^2 + \dots + x_d^2}} \right)_+ \right] = \frac{\Gamma(d/2)}{2\sqrt{\pi}\Gamma(d/2+1/2)} \in \left[\frac{1}{\sqrt{2\pi d}}, \frac{1}{\sqrt{2\pi(d-1)}} \right].$$

■

G.2. The exponential ergodicity of BPS

The next proposition is on the minorization condition of the 2-skeletons of BPS on the restricted domains. In short, minorization means that the stochastic process can go from any measurable set to any measurable set in the parameter space, which is a sufficient condition for the exponential ergodicity in the compact parameter space. 2-Skeleton means 2 step of the stochastic process. This proposition completes the proof of Theorem 10.

Proposition 14 *Under Assumption 1, the 2-skeletons of BPS satisfies the minorization condition; that is, for some $c > 0$, for all $(\theta, v) \in \Theta \times \mathbb{S}^{d-1}$ and all measurable $E \subset \Theta \times \mathbb{S}^{d-1}$, we have*

$$\widehat{\mathcal{Q}}^2((\theta, v), E) \geq c \int_{\Theta} \int_{\mathbb{S}^{d-1}} \mathbb{1}[(\theta, v) \in E] d\theta \mu_{\text{unif}}(dv).$$

Moreover, BPS is exponentially ergodic in total variation distance.

Proof We partially follow the proof of Lemma 4 of [13].

Let $f : \Theta \times \mathbb{S}^{d-1} \rightarrow [0, \infty)$ be a non-negative and bounded function. We also use the notation $M' = \sup_{(\theta, v) \in \Theta \times \mathbb{S}^{d-1}} (\bar{\lambda}(\theta, v) + C_B) < \infty$. By considering the event where the first update of v is *refreshment* from $\text{Unif}(\mathbb{S}^{d-1})$, we see that for any $(\theta_0, v_0) \in \Theta \times \mathbb{S}^{d-1}$,

$$\begin{aligned}
 &\int_{\Theta \times \mathbb{S}^{d-1}} f(\theta, v) \widehat{\mathcal{Q}}^2((\theta_0, v_0), (d\theta, dv)) \\
 &= \int_{\Theta \times \mathbb{S}^{d-1}} \int_{\Theta \times \mathbb{S}^{d-1}} f(\theta, v) \widehat{\mathcal{Q}}((\theta_1, v_1), (d\theta, dv)) \widehat{\mathcal{Q}}((\theta_0, v_0), (d\theta_1, dv_1)) \\
 &\geq \frac{\Lambda_{\text{ref}}}{M'} \inf_{\theta_1 \in \Theta} \int_{\Theta \times \mathbb{S}^{d-1}} f(\theta, v) \widehat{\mathcal{Q}}((\theta_1, v_1), (d\theta, dv)) \mu_{\text{unif}}(dv_1)
 \end{aligned}$$

holds. We also obtain that for $T \sim \text{Exp}(M')$, $V_1, V_2 \sim \text{i.i.d. Unif}(\mathbb{S}^{d-1})$, we have

$$\begin{aligned}
 & \inf_{\theta_1 \in \Theta} \int_{\Theta \times \mathbb{S}^{d-1}} f(\theta, v) \widehat{Q}((\theta_1, v_1), d\theta dv) \mu_{\text{unif}}(dv_1) \\
 & \geq \inf_{\theta_1 \in \Theta} \frac{\Lambda_{\text{ref}}}{M'} \mathbb{E} [\mathbb{1}[\theta_1 + TV_1 \in \Theta] f(\theta_1 + TV_1, V_2)] \\
 & \geq \inf_{\theta_1 \in \Theta} \frac{\Lambda_{\text{ref}}^2}{M'} \int_{[0, \infty) \times \mathbb{S}^{d-1}} \mathbb{1}[\theta_1 + tv_1 \in \Theta] e^{-M't} f(\theta_1 + tv_1, v) dt \mu_{\text{unif}}(dv_1) \mu_{\text{unif}}(dv) \\
 & \geq \inf_{\theta_1 \in \Theta} \frac{\Lambda_{\text{ref}}^2 e^{-M' \text{diam}(\Theta)}}{M'} \int_{[0, \infty) \times \mathbb{S}^{d-1}} \mathbb{1}[\theta_1 + tv_1 \in \Theta] f(\theta_1 + tv_1, v) dt \mu_{\text{unif}}(dv_1) \mu_{\text{unif}}(dv) \\
 & = \inf_{\theta_1 \in \Theta} \frac{\Lambda_{\text{ref}}^2 e^{-M' \text{diam}(\Theta)}}{M'} \int_{\Theta \times \mathbb{S}^{d-1}} \mathbb{1}[\theta \in \Theta] f(\theta, v) \|\theta - \theta_1\|^{1-d} d\theta \mu_{\text{unif}}(dv) \\
 & \geq \frac{\Lambda_{\text{ref}}^2 e^{-M' \text{diam}(\Theta)}}{M' \text{diam}(\Theta)^{d-1}} \int_{\Theta \times \mathbb{S}^{d-1}} f(\theta, v) d\theta \mu_{\text{unif}}(dv),
 \end{aligned}$$

where the second last equality uses a change of coordinates. Since f is generic, the minorization condition holds. Harris's theorem thus gives the exponential ergodicity of BPS. \blacksquare

Appendix H. Proof of Theorem 5

Proof We prove in the same way as the proof of Theorem 2.1 in [29]. Let θ_μ be a random variable satisfying $\theta_\mu \sim \mu_{\mathbf{z}}^{(\beta, \varepsilon)}$, where $\mu_{\mathbf{z}}^{(\beta, \varepsilon)}$ is defined in (5). We denote $\theta_K \sim \mu_{\mathbf{z}, K}$ as the output of Poisson SGD (Algorithm 1). We have

$$\begin{aligned}
 & \mathbb{E}_{\mathbf{z}}[\mathbb{E}_{\theta_K}[L(\theta_K)]] - \inf_{\theta \in \Theta} L(\theta) \\
 & = \mathbb{E}_{\mathbf{z}}[\mathbb{E}_{\theta_K}[L(\theta_K)] - \mathbb{E}_{\theta_\mu}[L(\theta_\mu)]] + \{\mathbb{E}_{\mathbf{z}}[\mathbb{E}_{\theta_\mu}[L(\theta_\mu)]] - \inf_{\theta \in \Theta} L(\theta)\},
 \end{aligned}$$

and the second term of right-hand side is written as

$$\begin{aligned}
 & \mathbb{E}_{\mathbf{z}}[\mathbb{E}_{\theta_\mu}[L(\theta_\mu)]] - \inf_{\theta \in \Theta} L(\theta) \\
 & = \mathbb{E}_{\mathbf{z}}[\mathbb{E}_{\theta_\mu}[L(\theta_\mu)] - \mathbb{E}_{\theta_\mu}[L_{\mathbf{z}}(\theta_\mu)]] + \left(\mathbb{E}_{\mathbf{z}}[\mathbb{E}_{\theta_\mu}[L_{\mathbf{z}}(\theta_\mu)]] - \inf_{\theta \in \Theta} L(\theta) \right).
 \end{aligned}$$

Letting $\theta^\circ = \text{argmin}_{\theta \in \Theta} L(\theta)$, the second part of the right-hand side in the equation above is

$$\begin{aligned}
 & \mathbb{E}_{\mathbf{z}}[\mathbb{E}_{\theta_\mu}[L_{\mathbf{z}}(\theta_\mu)]] - \inf_{\theta \in \Theta} L(\theta) = \mathbb{E}_{\mathbf{z}}[\mathbb{E}_{\theta_\mu}[L_{\mathbf{z}}(\theta_\mu)] - \inf_{\theta \in \Theta} L_{\mathbf{z}}(\theta)] + \left(\mathbb{E}_{\mathbf{z}} \left[\inf_{\theta \in \Theta} L_{\mathbf{z}}(\theta) - L_{\mathbf{z}}(\theta^\circ) \right] \right) \\
 & \leq \mathbb{E}_{\mathbf{z}}[\mathbb{E}_{\theta_\mu}[L_{\mathbf{z}}(\theta_\mu)] - \inf_{\theta \in \Theta} L_{\mathbf{z}}(\theta)].
 \end{aligned}$$

As a result, we have

$$\mathbb{E}_{\mathbf{z}}[\mathbb{E}_{\theta_K}[L(\theta_K)]] - \inf_{\theta \in \Theta} L(\theta) \leq \mathbb{E}_{\mathbf{z}}[\mathbb{E}_{\theta_K}[L(\theta_K)] - \mathbb{E}_{\theta_\mu}[L(\theta_\mu)]] \quad (16)$$

$$+\mathbb{E}_{\mathbf{z}}[\mathbb{E}_{\theta_{\mu}}[L(\theta_{\mu})] - \mathbb{E}_{\theta_{\mu}}[L_{\mathbf{z}}(\theta_{\mu})]] \quad (17)$$

$$+\mathbb{E}_{\mathbf{z}}[\mathbb{E}_{\theta_{\mu}}[L_{\mathbf{z}}(\theta_{\mu})] - \inf_{\theta \in \Theta} L_{\mathbf{z}}(\theta)]. \quad (18)$$

To evaluate the terms (16), (17), and (18), we prepare the following lemma to calculate the upper bound of the difference between two expected value by the Wasserstein distance.

Lemma 15 *Consider probability measures μ and ν on Θ . Suppose that $\sup_{z \in \mathcal{Z}} |\ell(z; 0)| \leq A$ and $\sup_{z \in \mathcal{Z}} \|\nabla \ell(z; 0)\| \leq B$ hold. Then, we obtain*

$$|\mathbb{E}_{\theta_1 \sim \mu}[\ell(z; \theta_1)] - \mathbb{E}_{\theta_2 \sim \nu}[\ell(z; \theta_2)]| \leq (c_1 W + B) \sqrt{W \mathcal{W}_1(\mu, \nu)}, \text{ and} \quad (19)$$

$$|\mathbb{E}_{\theta_1 \sim \mu}[L(\theta_1)] - \mathbb{E}_{\theta_2 \sim \nu}[L(\theta_2)]| \leq (c_1 W + B) \sqrt{W \mathcal{W}_1(\mu, \nu)}. \quad (20)$$

Proof

Under the assumption, Lemma 3.1 in [29] holds. Hence, we have

$$\|\nabla \ell(z; \theta)\| \leq c_1 \|\theta\| + B, \forall \theta \in \Theta, \forall z \in \mathcal{Z} \quad (21)$$

$$\ell(z; \theta) \leq \frac{c_1}{2} \|\theta\|^2 + B \|\theta\| + A, \forall \theta \in \Theta, \forall z \in \mathcal{Z}. \quad (22)$$

Moreover, from Lemma 3.5 in [29], for arbitrary two probability measures μ and ν , if we let

$$\sigma^2 = \max\{\mathbb{E}_{\theta_1 \sim \mu}[\|\theta_1\|^2], \mathbb{E}_{\theta_2 \sim \nu}[\|\theta_2\|^2]\},$$

then we have

$$|\mathbb{E}_{\theta_1 \sim \mu}[\ell(z; \theta_1)] - \mathbb{E}_{\theta_2 \sim \nu}[\ell(z; \theta_2)]| \leq (c_1 \sigma + B) \mathcal{W}_2(\mu, \nu).$$

Obviously, it also holds that

$$|\mathbb{E}_{\theta_1 \sim \mu}[L(\theta_1)] - \mathbb{E}_{\theta_2 \sim \nu}[L(\theta_2)]| \leq (c_1 \sigma + B) \mathcal{W}_2(\mu, \nu).$$

Since we have $\sigma \leq W$ and $\mathcal{W}_2(\mu, \nu) = \inf_{\pi \in \Pi(\mu, \nu)} (\int_{\Theta} \|z - z'\|^2 d\pi(z, z'))^{1/2} \leq \inf_{\pi \in \Pi(\mu, \nu)} (\int_{\Theta} W \|z - z'\|_1 d\pi(z, z'))^{1/2} = \sqrt{W \mathcal{W}_1(\mu, \nu)}$, we obtain the statement. \blacksquare

We start evaluating each of the terms (16), (17), and (18).

First, we study (16). From (20) in Lemma 15, we have

$$\begin{aligned} \mathbb{E}_{\theta_K} [L(\theta_K)] - \mathbb{E}_{\theta_{\mu}} [L(\theta_{\mu})] &\leq (c_1 W + B) \sqrt{W \mathcal{W}_1(\mu_{\mathbf{z}, K}, \mu_{\mathbf{z}}^{(\beta, \varepsilon)})} \\ &\leq (c_1 W + B) \sqrt{W d_K(\beta, \varepsilon, d)}. \end{aligned} \quad (23)$$

Second, we evaluate (17) using the same approach as [29]. Here, we need to evaluate

$$\mathbb{E}_{\theta_{\mu}} [\ell(z; \theta_{\mu})] - \mathbb{E}_{\theta_{\mu'}} [\ell(z; \theta_{\mu'})],$$

where $z \in \mathcal{Z}$ is an arbitrary sampled data, $\theta_{\mu'} \sim \mu_{\mathbf{z}'}^{(\beta, \varepsilon)}$ and $\mu_{\mathbf{z}'}^{(\beta, \varepsilon)}$ is the stationary distribution of BPS when one of the data z_i is changed to arbitrary $\bar{z}_i \in \mathcal{Z}$ and \mathbf{z}' is a dataset with replacing z_i to \bar{z}_i , and $L_{\mathbf{z}'}$ be its corresponding empirical risk. From (19) in Lemma 15, we have

$$\mathbb{E}_{\theta_{\mu}} [\ell(z; \theta_{\mu})] - \mathbb{E}_{\theta_{\mu'}} [\ell(z; \theta_{\mu'})] \leq (c_1 W + B) \mathcal{W}_2(\mu_{\mathbf{z}}^{(\beta, \varepsilon)}, \mu_{\mathbf{z}'}^{(\beta, \varepsilon)})$$

$$\leq (c_1 W + B) C_{\mu'} \left[\sqrt{D(\mu_{\mathbf{z}}^{(\beta, \varepsilon)} \| \mu_{\mathbf{z}'}^{(\beta, \varepsilon)})} + \left(\frac{D(\mu_{\mathbf{z}}^{(\beta, \varepsilon)} \| \mu_{\mathbf{z}'}^{(\beta, \varepsilon)})}{2} \right)^{\frac{1}{4}} \right],$$

where $D(\cdot \| \cdot)$ is KL-divergence and

$$C_{\mu'} := 2 \inf_{\lambda > 0} \left(\frac{1}{\lambda} \left(\frac{3}{2} + \log \int_{\Theta} e^{\lambda \|\theta\|^2} \mu_{\mathbf{z}'}^{(\beta, \varepsilon)}(d\theta) \right) \right)^{\frac{1}{2}},$$

which is from Corollary 2.3 in [4]. Also, since we have $\|\theta\| \leq W$, $C_{\mu'} \leq 2W$ holds. We denote the density functions of $\mu_{\mathbf{z}}^{(\beta, \varepsilon)}$, $\mu_{\mathbf{z}'}^{(\beta, \varepsilon)}$ as $p_{\mathbf{z}}$, $p_{\mathbf{z}'}$, and the normalization constants as $\Lambda_{\mathbf{z}}$, $\Lambda_{\mathbf{z}'}$ respectively. Let us calculate $D(\mu_{\mathbf{z}}^{(\beta, \varepsilon)} \| \mu_{\mathbf{z}'}^{(\beta, \varepsilon)})$. We have

$$\frac{p_{\mathbf{z}}(\theta)}{p_{\mathbf{z}'}(\theta)} = \frac{\Lambda_{\mathbf{z}'}}{\Lambda_{\mathbf{z}}} \cdot \frac{\beta M_{\ell} + 1/\varepsilon + a_d \beta \|\nabla L_{\mathbf{z}}(\theta)\|}{\beta M_{\ell} + 1/\varepsilon + a_d \beta \|\nabla L_{\mathbf{z}'}(\theta)\|} \exp(-\beta(L_{\mathbf{z}}(\theta) - L_{\mathbf{z}'}(\theta))), \quad (24)$$

so in order to obtain the upper bound of $D(\mu_{\mathbf{z}}^{(\beta, \varepsilon)} \| \mu_{\mathbf{z}'}^{(\beta, \varepsilon)})$, we suppress each of the three terms of the right-hand side of (24). First, we suppress the second term.

$$\begin{aligned} \|\nabla L_{\mathbf{z}}(\theta)\| &= \left\| \nabla L_{\mathbf{z}'}(\theta) + \frac{1}{n} (\nabla \ell(z_i; \theta) - \nabla \ell(\bar{z}_i; \theta)) \right\| \\ &\leq \|\nabla L_{\mathbf{z}'}(\theta)\| + \frac{1}{n} \|\nabla \ell(z_i; \theta) - \nabla \ell(\bar{z}_i; \theta)\| \\ &\leq \|\nabla L_{\mathbf{z}'}(\theta)\| + \frac{2}{n} (c_1 \|\theta\| + B), \end{aligned}$$

where the last inequality is from (21). Hence,

$$\begin{aligned} \frac{\beta M_{\ell} + 1/\varepsilon + a_d \beta \|\nabla L_{\mathbf{z}}(\theta)\|}{\beta M_{\ell} + 1/\varepsilon + a_d \beta \|\nabla L_{\mathbf{z}'}(\theta)\|} &\leq \frac{\beta M_{\ell} + 1/\varepsilon + a_d \beta \left(\|\nabla L_{\mathbf{z}'}(\theta)\| + \frac{2}{n} (c_1 \|\theta\| + B) \right)}{\beta M_{\ell} + 1/\varepsilon + a_d \beta \|\nabla L_{\mathbf{z}'}(\theta)\|} \\ &\leq 1 + \frac{2a_d \beta (c_1 W + B)}{n(\beta M_{\ell} + 1/\varepsilon)} \\ &\leq 1 + \frac{2a_d (c_1 W + B)}{nM_{\ell}} \end{aligned} \quad (25)$$

holds. Second, we suppress the third term. We have

$$\begin{aligned} \exp(-\beta(L_{\mathbf{z}}(\theta) - L_{\mathbf{z}'}(\theta))) &= \exp\left(-\beta \left(\frac{1}{n} (\ell(z_i; \theta) - \ell(\bar{z}_i; \theta)) \right)\right) \\ &\leq \exp\left(\frac{\beta}{n} \left(\frac{c_1 \|\theta\|^2}{2} + B\|\theta\| + A \right)\right) \\ &\leq \exp\left(\frac{\beta}{n} \left(\frac{c_1 W^2}{2} + BW + A \right)\right), \end{aligned} \quad (26)$$

where we use (22). Finally, we suppress the first term. Using (25) and (26), we have

$$\frac{\Lambda_{\mathbf{z}'}}{\Lambda_{\mathbf{z}}} = \frac{\int_{\theta \in \Theta} (\beta M_{\ell} + 1/\varepsilon + a_d \beta \|\nabla L_{\mathbf{z}'}(\theta)\|) \exp(-\beta L_{\mathbf{z}'}(\theta)) d\theta}{\int_{\theta \in \Theta} (\beta M_{\ell} + 1/\varepsilon + a_d \beta \|\nabla L_{\mathbf{z}}(\theta)\|) \exp(-\beta L_{\mathbf{z}}(\theta)) d\theta}$$

$$\leq \left(1 + \frac{2a_d(c_1W + B)}{nM_\ell}\right) \exp\left(\frac{\beta}{n}\left(\frac{c_1W^2}{2} + BW + A\right)\right). \quad (27)$$

Combining (25), (26) and (27), we have

$$\begin{aligned} \log \frac{p_{\mathbf{z}}(\theta)}{p_{\mathbf{z}'}(\theta)} &\leq 2 \log \left(1 + \frac{2a_d(c_1W + B)}{nM_\ell}\right) + \frac{2\beta}{n}\left(\frac{c_1W^2}{2} + BW + A\right) \\ &\leq \frac{1}{n}\left(\frac{4a_d(c_1W + B)}{M_\ell} + \beta(c_1W^2 + 2BW + 2A)\right), \end{aligned}$$

so

$$D(\mu_{\mathbf{z}}^{(\beta, \varepsilon)} \parallel \mu_{\mathbf{z}'}^{(\beta, \varepsilon)}) \leq \frac{1}{n}\left(\frac{4a_d(c_1W + B)}{M_\ell} + \beta(c_1W^2 + 2BW + 2A)\right)$$

holds. We set $C_d = 4a_d(c_1W + B)/M_\ell$ and $C = c_1W^2 + 2BW + 2A$, then we have

$$(17) \leq 2W(c_1W + B) \left(\left(\frac{C_d + \beta C}{n}\right)^{\frac{1}{2}} + \left(\frac{C_d + \beta C}{n}\right)^{\frac{1}{4}} \right). \quad (28)$$

Finally, we evaluate (18). Let us denote

$$\begin{aligned} \Lambda_{\mathbf{z}}(\theta) &= \frac{\Lambda}{\beta M_\ell + 1/\varepsilon + a_d\beta \|\nabla L_{\mathbf{z}}(\theta)\|} \\ \Lambda &= \int_{\theta \in \Theta} (\beta M_\ell + 1/\varepsilon + a_d\beta \|\nabla L_{\mathbf{z}}(\theta)\|) e^{-\beta L_{\mathbf{z}}(\theta)} d\theta. \end{aligned}$$

Since the distribution of θ_μ is

$$\mu_{\mathbf{z}}^{(\beta, \varepsilon)}(d\theta) \propto \left(\beta M_\ell + \frac{1}{\varepsilon} + a_d\beta \|\nabla L_{\mathbf{z}}(\theta)\|\right) \exp(-\beta L_{\mathbf{z}}(\theta)),$$

we have

$$\begin{aligned} \mathbb{E}_{\theta_\mu}[L_{\mathbf{z}}(\theta_\mu)] &= -\frac{1}{\beta} \left(\mathbb{E}_{\theta_\mu} \left[\log \frac{e^{-\beta L_{\mathbf{z}}(\theta_\mu)}}{\Lambda_{\mathbf{z}}(\theta_\mu)} \right] + \mathbb{E}_{\theta_\mu} [\log \Lambda_{\mathbf{z}}(\theta_\mu)] \right) \\ &= \frac{1}{\beta} \left(\mathbb{E}_{\theta_\mu} [-\log p_{\mathbf{z}}(\theta_\mu)] - \mathbb{E}_{\theta_\mu} [\log \Lambda_{\mathbf{z}}(\theta_\mu)] \right). \end{aligned}$$

Since we have $\mathbb{E}_{\theta_\mu}[\|\theta_\mu\|^2] \leq W^2$, we can calculate the upper bound of $\mathbb{E}_{\theta_\mu}[-\log p_{\mathbf{z}}(\theta_\mu)]$ by the differential entropy of Gaussian distribution in the same way as the discussion of Section 3.5 in [29]:

$$\mathbb{E}_{\theta_\mu}[-\log p_{\mathbf{z}}(\theta_\mu)] \leq \frac{d}{2} \log \left(\frac{2\pi e}{d} W^2 \right).$$

Using (21), we have

$$\log \Lambda_{\mathbf{z}}(\theta) \geq \log \frac{\Lambda}{\beta M_\ell + 1/\varepsilon + a_d\beta(c_1W + B)}.$$

In addition,

$$\begin{aligned}
 \log \Lambda &= \log \int_{\theta \in \Theta} (\beta M_\ell + 1/\varepsilon + a_d \beta \|\nabla L_{\mathbf{z}}(\theta)\|) e^{-\beta L_{\mathbf{z}}(\theta)} d\theta \\
 &\geq \log \int_{\theta \in \Theta} (\beta M_\ell + 1/\varepsilon) e^{-\beta L_{\mathbf{z}}(\theta)} d\theta \\
 &= \log(\beta M_\ell + 1/\varepsilon) + \log \int_{\theta \in \Theta} e^{-\beta L_{\mathbf{z}}(\theta)} d\theta \\
 &\geq \log(\beta M_\ell + 1/\varepsilon) - \beta L_{\mathbf{z}}^* + \frac{d}{2} \log \frac{2\pi}{c_1 \beta}
 \end{aligned}$$

holds, where the last inequality is from the equation (3.21) in [29]. Here, we denote $L_{\mathbf{z}}^* = \inf_{\theta \in \Theta} L_{\mathbf{z}}(\theta)$. Hence, we have

$$\begin{aligned}
 (18) &\leq \frac{1}{\beta} \left(\frac{d}{2} \log \left(\frac{2\pi e}{d} W^2 \right) + \log \frac{\beta M_\ell + 1/\varepsilon + a_d \beta (c_1 W + B)}{\beta M_\ell + 1/\varepsilon} + \beta L_{\mathbf{z}}^* - \frac{d}{2} \log \frac{2\pi}{c_1 \beta} \right) - L_{\mathbf{z}}^* \\
 &\leq \frac{1}{\beta} \left(\frac{d}{2} \log \frac{e W^2 c_1 \beta}{d} + \log \left(1 + \frac{a_d (c_1 W + B)}{M_\ell} \right) \right). \tag{29}
 \end{aligned}$$

We combine the result (23), (28), and (29), then obtain the statement. \blacksquare

Appendix I. Proof of Proposition 6

Proof Let θ_μ and θ_ν be the random variable which obey the distributions $\mu_{\mathbf{z}}^{(\beta, \varepsilon)}$ and $\nu_{\mathbf{z}}^{(\beta)}$ respectively.

In the same way as Theorem 5, we have

$$\mathbb{E}_{\mathbf{z}} [\mathbb{E}_{\theta_K} [L(\theta_K)]] - \inf_{\theta \in \Theta} L(\theta) \leq \mathbb{E}_{\mathbf{z}} [\mathbb{E}_{\theta_K} [L(\theta_K)] - \mathbb{E}_{\theta_\mu} [L(\theta_\mu)]] \tag{30}$$

$$+ \mathbb{E}_{\mathbf{z}} [\mathbb{E}_{\theta_\mu} [L(\theta_\mu)] - \mathbb{E}_{\theta_\nu} [L(\theta_\nu)]] \tag{31}$$

$$+ \mathbb{E}_{\mathbf{z}} [\mathbb{E}_{\theta_\nu} [L(\theta_\nu)] - \mathbb{E}_{\theta_\nu} [L_{\mathbf{z}}(\theta_\nu)]] \tag{32}$$

$$+ \mathbb{E}_{\mathbf{z}} [\mathbb{E}_{\theta_\nu} [L_{\mathbf{z}}(\theta_\nu)] - \inf_{\theta \in \Theta} L_{\mathbf{z}}(\theta)]. \tag{33}$$

(30) can be evaluated in the same as Theorem 5.

First, we evaluate (31). We have

$$\mathbb{E}_{\theta_\mu} [L(\theta_\mu)] - \mathbb{E}_{\theta_\nu} [L(\theta_\nu)] \leq W \mathcal{W}_2(\mu_{\mathbf{z}}^{(\beta, \varepsilon)}, \nu_{\mathbf{z}}^{(\beta)})$$

from the same discussion in the proof of Theorem 5. Since both θ_μ and θ_ν satisfy the log-Sobolev inequality, we can use Otto-Villani theorem [1], and

$$\mathcal{W}_2(\mu_{\mathbf{z}}^{(\beta, \varepsilon)}, \nu_{\mathbf{z}}^{(\beta)}) \leq \sqrt{2c_{\text{LS}}^{(\beta)} D(\mu_{\mathbf{z}}^{(\beta, \varepsilon)} \parallel \nu_{\mathbf{z}}^{(\beta)})}$$

holds, where D denotes the KL-divergence and $c_{\text{LS}}^{(\beta)}$ is the log-Sobolev constant of $\nu_{\mathbf{z}}^{(\beta)}$. We have

$$D(\mu_{\mathbf{z}}^{(\beta, \varepsilon)} \parallel \nu_{\mathbf{z}}^{(\beta)}) = \mathbb{E}_{\theta \sim \mu} \left[\log \frac{(\beta M_\ell + 1/\varepsilon + a_d \beta \|\nabla L_{\mathbf{z}}(\theta)\|) \exp(-\beta L_{\mathbf{z}}(\theta)) / \Lambda_\mu}{\exp(-\beta L_{\mathbf{z}}(\theta)) / \Lambda_\nu} \right]$$

$$\leq \mathbb{E}_{\theta \sim \mu} \left[\log(\beta M_\ell + 1/\varepsilon + a_d \beta M_\ell) \frac{\Lambda_\nu}{\Lambda_\mu} \right],$$

where Λ_μ and Λ_ν are normalizing constants of the density functions of $\mu_{\mathbf{z}}^{(\beta, \varepsilon)}$ and $\nu_{\mathbf{z}}^{(\beta)}$ respectively. We have

$$\frac{\Lambda_\nu}{\Lambda_\mu} = \frac{\int_{\Theta} \exp(-\beta L_{\mathbf{z}}(\theta)) d\theta}{\int_{\Theta} (\beta M_\ell + 1/\varepsilon + a_d \beta \|\nabla L_{\mathbf{z}}(\theta)\|) \exp(-\beta L_{\mathbf{z}}(\theta)) d\theta} \leq \frac{1}{\beta M_\ell + 1/\varepsilon},$$

hence we have

$$D(\mu_{\mathbf{z}}^{(\beta, \varepsilon)} \parallel \nu_{\mathbf{z}}^{(\beta)}) \leq \log(1 + a_d \beta \varepsilon M_\ell).$$

As a result, we obtain

$$\mathbb{E}_{\theta_\mu} [L(\theta_\mu)] - \mathbb{E}_{\theta_\nu} [L(\theta_\nu)] \leq W \sqrt{2c_{\text{LS}}^{(\beta)} \log(1 + a_d \beta \varepsilon M_\ell)}. \quad (34)$$

Second, we evaluate (32). Let $\nu_{\mathbf{z}'}^{(\beta)}$ be the Gibbs distribution when one of the data z_i is replaced by z'_i . In the same way as Section 3.6 in [29], we have

$$\mathcal{W}_2(\nu_{\mathbf{z}}^{(\beta)}, \nu_{\mathbf{z}'}^{(\beta)}) \leq \frac{2c_{\text{LS}}^{(\beta)} \beta M_\ell}{n}.$$

Hence, we have

$$\mathbb{E}_{\theta_\nu} [L(\theta_\nu)] - \mathbb{E}_{\theta_\nu} [L_{\mathbf{z}}(\theta_\nu)] \leq (c_1 W + B) \frac{2c_{\text{LS}}^{(\beta)} \beta M_\ell}{n}. \quad (35)$$

Finally, we evaluate (33). This term can be evaluated on the same way as Proposition 3.4 in [29] and we have

$$\begin{aligned} \mathbb{E}_{\theta_\nu} [L_{\mathbf{z}}(\theta_\nu)] - \inf_{\theta \in \Theta} L_{\mathbf{z}}(\theta) &\leq \frac{1}{\beta} \left(\frac{d}{2} \log \left(\frac{2\pi e W^2}{d} \right) - \frac{d}{2} \log \frac{2\pi}{c_1 \beta} \right) \\ &= \frac{d}{2\beta} \log \left(\frac{e W^2 c_1 \beta}{d} \right). \end{aligned} \quad (36)$$

We combine the result (34), (35), and (36), then obtain the statement. \blacksquare

Appendix J. Proof of Theorem 3

Proof Let θ_K, θ_μ be the random variables whose distribution is $\mu_{\mathbf{z}, K}^{(\beta, \varepsilon)}$ and $\mu_{\mathbf{z}}^{(\beta, \varepsilon)}$ respectively. Let $L_{\mathbf{z}}^* = \min_{\theta \in \Theta} L_{\mathbf{z}}(\theta)$. We have

$$\mathbb{E}_{\theta_K} [L_{\mathbf{z}}(\theta_K)] - L_{\mathbf{z}}^* = (\mathbb{E}_{\theta_K} [L_{\mathbf{z}}(\theta_K)] - \mathbb{E}_{\theta_\mu} [L_{\mathbf{z}}(\theta_\mu)]) + (\mathbb{E}_{\theta_\mu} [L_{\mathbf{z}}(\theta_\mu)] - L_{\mathbf{z}}^*).$$

As the first term of the right-hand side, we can use the Wasserstein distance in the same way as the proof of Theorem 5 as in (23). Hence, we have

$$\mathbb{E}_{\theta_K} [L_{\mathbf{z}}(\theta_K)] - \mathbb{E}_{\theta_\mu} [L_{\mathbf{z}}(\theta_\mu)] \leq (c_1 W + B) \sqrt{W d_K(\beta, \varepsilon, d)}.$$

Further, using (29) in the Proof of Theorem 5,

$$\mathbb{E}_{\theta_\mu} [L_{\mathbf{z}}(\theta_\mu)] \leq \frac{1}{\beta} \left(\frac{d}{2} \log \frac{eW^2 c_1 \beta}{d} + \log \left(1 + \frac{a_d(c_1 W + B)}{M_\ell} \right) \right) + L_{\mathbf{z}}^*$$

holds, which completes the proof. ■