

Constraint-Based Regularization of Neural Networks

Benedict Leimkuhler

Timothée Pouchon

Tiffany Vlaar*

Amos Storkey

University of Edinburgh, UK

B.LEIMKUHLER@ED.AC.UK

TIMOTHEE.POUCHON@ED.AC.UK

TIFFANY.VLAAR@ED.AC.UK

A.STORKEY@ED.AC.UK

Abstract

We propose a method for efficiently incorporating constraints into a stochastic gradient Langevin framework for the training of deep neural networks. Constraints allow direct control of the parameter space of the model. Appropriately designed, they reduce the vanishing/exploding gradient problem, control weight magnitudes and stabilize deep neural networks and thus improve the robustness of training algorithms and generalization capabilities of the trained neural network. We present examples of constrained training methods motivated by orthogonality preservation for weight matrices and explicit weight normalizations. We describe the methods in the overdamped formulation of Langevin dynamics and the underdamped form, in which momenta help to improve sampling efficiency. Our methods see performance improvements on image classification tasks.

Keywords: Constrained Optimization, Langevin Dynamics, Orthogonality Preservation.

1. Introduction

In this paper we explore stochastic training methods based on Langevin dynamics combined with algebraic constraints. Our general framework allows for incorporating constraints into standard training schemes and sampling methods for neural networks (NNs). Constraints provide direct control of the parameter space of a model and hence afford a means to improve its generalization performance. Current approaches to enhance the generalization performance of overparameterized NNs consist of both explicit and implicit regularization techniques [31]. Examples of the former are L1 [42, 47] and L2 [14] regularization, which modify the loss by adding a parameter norm penalty term. Batch normalization (BatchNorm) [16] is a technique that causes an implicit regularization effect. BatchNorm can be viewed as tantamount to a constraint imposed on the network’s parameters during training. Although BatchNorm is widely used, explanations for the method’s success remain elusive; claims that it would reduce internal covariance shift [16] or smooth the loss landscape [38] have been disputed [38, 51]. The reliance on increasingly complex strategies does little to enhance the explainability of NNs, so robust simplification of all aspects of training is desirable.

In this paper we highlight the potential of constrained stochastic differential equation (SDE) based algorithms to provide a simpler overall framework for NN training. An example of a constraint that can be easily introduced using our general framework is orthogonality of the weight matrix. We provide a detailed algorithm for this purpose in a Langevin dynamics-based setting.

In NN training one aims to minimize the loss $L_X(\theta)$ for parameters $\theta \in \mathbb{R}^{|n|}$ and data X . A popular training scheme is stochastic gradient descent (SGD). SGD may be improved by incorporating momenta and additive noise, or more generally by embedding the loss gradient in a Langevin

dynamics (LD) framework [9]. Using low temperatures [24, 46], sampling methods have been found to enhance exploration and speed the approach to ‘good’ minima, which enhance their generalization to nearby data sets. Ergodic properties of the idealized SDEs associated with gradient schemes may help these methods to ensure robust exploration of a useful range of parameters.

Constraints can be seen as limiting cases of penalty-based regularization which replaces minimization of the loss $L_X(\theta)$ by that of the augmented loss $L_X^c(\theta) = L_X(\theta) + \frac{1}{\varepsilon^2}g(\theta)^2$, where $g(\cdot)$ is a suitable smooth function of the parameters. In the limit $\varepsilon \rightarrow 0$, these penalty terms introduce an undesirable stiffness and consequent stability restriction in gradient-based training. It is therefore natural to relate the above system to a constrained optimization task subject to $g(\theta) = 0$.

2. Neural networks with constraints

We suggest to use constraints when training NNs. Imposing good priors on NNs is known to improve performance, e.g. CNNs suit image datasets better than overparameterized fully connected NNs, despite being a subset of the latter [10]. Using constraints also arises naturally in the control of vanishing/exploding gradients. Constraints can be used to control the magnitudes of individual weights and/or to limit the growth of gradients in deep NNs. We present various approaches below.

A L -layer NN has parameters $\theta \in \mathbb{R}^{|n|}$, with a weight matrix $W^\ell \in \mathbb{R}^{d^\ell \times d^{\ell-1}}$ and bias vector $b^\ell \in \mathbb{R}^{d^\ell}$ for each layer ℓ . To allow for inequality constraints, we define slack variables vector $\xi \in \mathbb{R}^{n^\xi}$ and consider as variable $q = (\theta, \xi) \in \mathbb{R}^d$, where $d = |n| + n^\xi$. The constraint manifold is

$$\Sigma = \{q \in \mathbb{R}^d \mid g(q) = 0\}, \quad g: \mathbb{R}^d \rightarrow \mathbb{R}^m. \quad (1)$$

We partition $\theta = (\theta^u, \theta^c)$ into unconstrained $\theta^u \in \mathbb{R}^{n^u}$ and constrained $\theta^c \in \mathbb{R}^{n^c}$ parameters.

Circle constraints: In a *circle constraint*, we restrict each parameter in θ^c as $|\theta_i^c| \leq r_i$, where $r_i > 0$ is given. We thus introduce $m = n^c = n^\xi$ slack variables ξ_i and define

$$g_i(q) = |\theta_i^c|^2 + |\xi_i|^2 - r_i^2 \quad 1 \leq i \leq m. \quad (2)$$

Note that if $q \in \Sigma$ then the parameters in θ^c are bounded as desired.

Sphere constraints: In a similar way, we could opt to restrict the sums of squares of weights associated to the input channels of any node. This constraint is analogous to max-norm [40, 41] as used in ad hoc regularization procedures. In our context, introducing such constraints would yield distinctive training methods, although we omit discussion of these here due to space limitations.

Orthogonality constraints: Orthogonal weight matrices can mitigate the vanishing/exploding gradient problem in RNNs [1, 32, 45], and are developing a growing following in the CNN literature [4, 15, 37]. Orthogonal initialization is linked to achieving dynamical isometry [35, 36, 39], which can accelerate training. Xiao et al. [49] were able to train 10,000 layer vanilla CNNs, without learning rate decay, BatchNorm or residual connections, by using initial orthogonal convolution kernels. Methods for enforcing orthogonality during training include the use of ‘soft’ constraints which add a restraint term to the loss [4, 8, 50] and hard constraints based on optimization over Stiefel manifolds [15, 17]. The latter requires repeated singular value decomposition of high-dimensional matrices during training, which is costly. Here we propose a straightforward algorithm to incorporate orthonormality constraints for rectangular matrices within our NN training framework, with manageable additional cost. We make no empirical claims over other manifold optimization methods, but rather provide a framework for network optimization that is theoretically sound, flexible enough

to incorporate new constraints, and demonstrates good properties relative to standard SGD training. We set $\theta^u = b^\ell$, and define the *orthogonality constraint* for layer ℓ (which has n^ℓ parameters) as

$$g(q) = \begin{cases} (W^\ell)^T W^\ell - I_{n^{\ell-1}} & \text{if } n^{\ell-1} \leq n^\ell, \\ W^\ell (W^\ell)^T - I_{n^\ell} & \text{otherwise.} \end{cases} \quad (3)$$

3. Constrained SDEs and their discretization

We now describe SDE-based methods for constrained NN training. An alternative to our approach is constrained Hamiltonian Monte Carlo (HMC) methods [12, 28, 52]. Although HMC schemes have nil sampling bias if fully converged, their acceptance rates depend on stepsize and system size [5, 7]. SDE-based methods are often preferred in high-dimensional sampling calculations as they offer greater overall efficiency for a fixed computational budget. In this section we discuss properties of constrained Langevin Dynamics. For further discussion on (unconstrained) LD see [34]. LD discretizations are studied in [26, Chap. 3], [11] (overdamped) and [21, 27] (underdamped).

Constrained Langevin: ergodicity and central limit theorem. The NN loss function naturally extends to the variable $q = (\theta, \xi) \in \mathbb{R}^d$ as $V(q) = L_X(\theta)$ (note that in particular $\nabla_\xi V = 0$). The first continuous training method we consider is the constrained overdamped Langevin system

$$dq_t = -\nabla V(q_t) dt + \sqrt{2\tau} d\mathcal{W}_t - \nabla_q g(q_t) d\lambda_t, \quad 0 = g(q_t), \quad (4)$$

where \mathcal{W} is a d -dimensional Wiener process, $\tau \geq 0$ is the temperature hyperparameter, and λ_t is an \mathbb{R}^m -valued vector of Lagrange multipliers. Provided the initial configuration q_0 satisfies the constraint, any trajectory q_t of (4) remains on the constraint manifold Σ defined in (1). When $\beta^{-1} = \tau > 0$, (4) is equivalent to an underlying ergodic (unconstrained) SDE (see [26, Chap. 3] and Appx. A.1) with unique invariant measure $d\nu_\Sigma = Z^{-1} e^{-\beta V(q)} d\sigma_\Sigma$, $Z = \int_\Sigma e^{-\beta V(q)} d\sigma_\Sigma$, where σ_Σ is the surface measure on Σ . Ergodicity ensures that averages of observables with respect to ν_Σ can be approximated by time averages of trajectories of (4). To ensure the practical use of (4) as a training method, we need the convergence to occur in a reasonable time. Thanks to the reversibility of the underlying SDE (see Appx. A.1), exponential convergence to equilibrium occurs as a consequence of a Poincaré inequality for ν_Σ , which holds provided the curvature of the manifold is well behaved (see Appx. A.2 and A.3). Poincaré inequalities on manifolds and their use in the analysis of diffusion processes are presented in [3, Chap. 4].

Introducing momenta p leads to constrained underdamped LD, the 2nd order counterpart of (4)

$$dq_t = p_t dt, \quad dp_t = (-\nabla_q V(q_t) - \gamma p_t) dt + \sqrt{2\gamma\tau} d\mathcal{W}_t - \nabla_q g(q_t) d\lambda_t, \quad 0 = g(q_t), \quad (5)$$

where γ is the friction hyperparameter. The constraint induces a cotangency condition: $p \in T_q^* \Sigma$, where $T_q^* \Sigma = \{p \in \mathbb{R}^d \mid \nabla^T g(q)p = 0\}$ is the cotangent space of the manifold Σ . The corresponding phase space is the cotangent bundle $T^* \Sigma = \{(q, p) \mid q \in \Sigma, p \in T_q^* \Sigma\}$. Given an initial pair $(q, p) \in T^* \Sigma$, any trajectory (q_t, p_t) of (5) stays on $T^* \Sigma$ for all time. In case $\tau > 0$, (5) is equivalent to an underlying ergodic SDE, whose invariant measure is $d\mu = e^{-\beta H(q,p)} d\sigma_{T^* \Sigma}$, with Hamiltonian $H(q, p) = V(q) + \frac{1}{2} p^T p$ and $\sigma_{T^* \Sigma}$ the Liouville measure of the cotangent bundle [27]. Exponential convergence also holds here, but the proof is more technical (e.g. based on hypocoercivity [25, 44]).

Discretization of constrained Langevin dynamics. The simplest iteration scheme $q_n \in \Sigma \mapsto q_{n+1} \in \Sigma$ for constrained overdamped Langevin dynamics (4) consists of an Euler–Maruyama step followed by projection onto Σ . The best choice for the projection is constraint-specific. For circle constraints we suggest orthogonal projection, which is both explicit and robust (see Appx. B.3). For orthogonality constraints, we derive an efficient quasi-Newton scheme (Appx. B.5). The latter leads to the following training method (written here for $Q = W^\ell$ if $n^\ell \leq n^{\ell-1}$ and $Q = (W^\ell)^T$ otherwise, $s = \min\{n^\ell, n^{\ell-1}\}$): one training iteration $Q_n \in \Sigma \mapsto Q_{n+1} \in \Sigma$ is given by

$$\text{for } k = 0 \text{ to } K - 1: \quad Q^{(k+1)} = Q^{(k)} - \frac{1}{2}Q_n((Q^{(k)})^T Q^{(k)} - I_s). \quad (6)$$

We initialize $Q^{(0)} = Q_n - h\nabla_Q V(Q) + \sqrt{2\tau h}R_n$, with stepsize h and independent standard random normal matrix R_n of the same size as Q . After K quasi-Newton iterations we set $Q_{n+1} \equiv Q^{(K)}$. For the constrained underdamped Langevin system (5), the ABO splitting strategy from [21] gives

$$\text{A:} \quad dq_t = p_t dt, \quad dp_t = -\nabla_q g(q_t) d\lambda_t, \quad 0 = g(q_t), \quad 0 = \nabla_q g(q_t)p_t, \quad (7)$$

$$\text{B:} \quad dq_t = 0, \quad dp_t = -\nabla_q V(q_t) dt - \nabla_q g(q_t) d\mu_t, \quad 0 = g(q_t), \quad 0 = \nabla_q g(q_t)p_t, \quad (8)$$

$$\text{O:} \quad dq_t = 0, \quad dp_t = -\gamma p_t dt + \sqrt{2\gamma\tau} d\mathcal{W}_t - \nabla_q g(q_t) d\nu_t, \quad 0 = g(q_t), \quad 0 = \nabla_q g(q_t)p_t \quad (9)$$

We use an OBA sequence, which in the case $\tau = 0$ and by re-scaling the momentum and step size variables, is equivalent to the standard PyTorch form of SGD with momentum [24, 33]. The B and O components can be solved exactly (in law) while the A component can be approximated using a standard scheme for constrained ODEs (e.g. SHAKE or RATTLE [22, Chap. 7]). Importantly, the A component does not involve the evaluation of the gradient. For circle constraints the A step can be solved explicitly (see Appx. B.4). For orthogonality constraints (Appx. B.6): for $Q \in \Sigma$, the projection onto the cotangent space $T_Q^*\Sigma$ is defined as $\Pi_Q : \mathbb{R}^{r \times s} \rightarrow \mathbb{R}^{r \times s}$, $\bar{P} \mapsto \Pi_Q \bar{P} = \bar{P} - \frac{1}{2}Q(\bar{P}^T Q + Q^T \bar{P})$. Then the ABO steps $(Q_n, P_n) \in T^*\Sigma \mapsto (Q_{n+1}, P_{n+1}) \in T^*\Sigma$ are

$$\begin{aligned} \text{(A)} \quad & \begin{cases} Q^{(0)} = Q_n + hP_n, \text{ for } k = 0:K-1: Q^{(k+1)} = Q^{(k)} - \frac{1}{2}Q_n((Q^{(k)})^T Q^{(k)} - I_s), \\ Q_{n+1} = Q^{(K)}, \bar{P}_{n+1} = P_n + \frac{1}{h}(Q_{n+1} - Q^{(0)}), \quad P_{n+1} = \Pi_{Q_{n+1}} \bar{P}_{n+1}, \end{cases} \\ \text{(B)} \quad & \begin{cases} Q_{n+1} = Q_n, \quad \bar{P}_{n+1} = P_n - h\nabla_Q V(Q_n), \quad P_{n+1} = \Pi_{Q_n} \bar{P}_{n+1}, \end{cases} \\ \text{(O)} \quad & \begin{cases} Q_{n+1} = Q_n, \quad \bar{P}_{n+1} = e^{-\gamma h} P_n + \sqrt{\tau(1 - e^{-2\gamma h})} R_n, \quad P_{n+1} = \Pi_{Q_n} \bar{P}_{n+1}, \end{cases} \end{aligned} \quad (10)$$

4. Numerical Experiments

The use of constraints can enhance generalization performance. We support this claim by comparing the performance of NN architectures trained using our constrained methods vs. unconstrained SGD. We set $\tau = 0$ and use equivalent learning rates to present a fair comparison. We denote our circle and orthogonal Constrained *overdamped* Langevin Algorithms as c-CoLA-*od* and o-CoLA-*od*, respectively. We compare *underdamped* variants (CoLA-*ud*) with SGD with momentum (SGD-m).

Orthogonality Constraints We compare SGD with orthogonality-preserving overdamped Langevin (Fig. 1). The goal is to train a MLP with p hidden layers on a tightly wound spiral binary classification problem (Fig. C5). For SGD we show results for both i) standard PyTorch initialization and ii) orthogonal initialization. A clear advantage imposing orthogonality appears with more than 3 hidden layers. In Fig. 2 we show that the use of a small temperature perturbation $\tau = 1e-6$ speeds up training and slightly increases the test accuracy obtained for MLPs trained on the spiral data set.

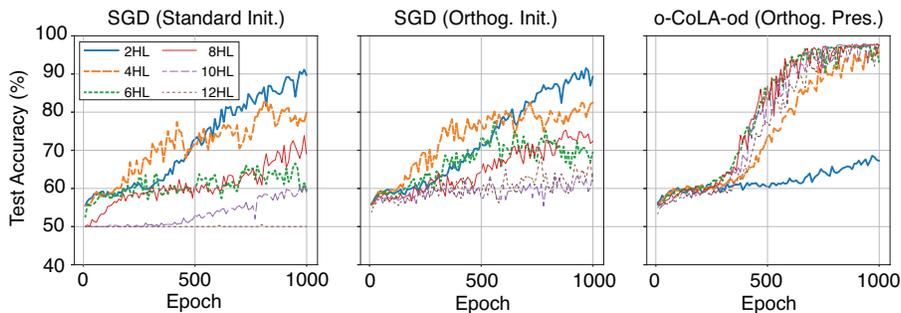


Figure 1: Test acc. of MLPs with p -number of 100-node hidden layers (HL), ReLU activation. The MLPs are trained on a 4-turn spiral dataset (Fig. C5) using SGD with standard initialization (left), SGD with orthogonal initialization (middle) and o-CoLA-od with $\tau = 0$ (right). For o-CoLA-od we constrain weights in all layers, apart from input and output layers. Stepsize $h = 0.1$ for all methods. Results are averaged over 10 runs. o-CoLA-od significantly outperforms unconstrained SGD for MLPs with more than 3 hidden layers.

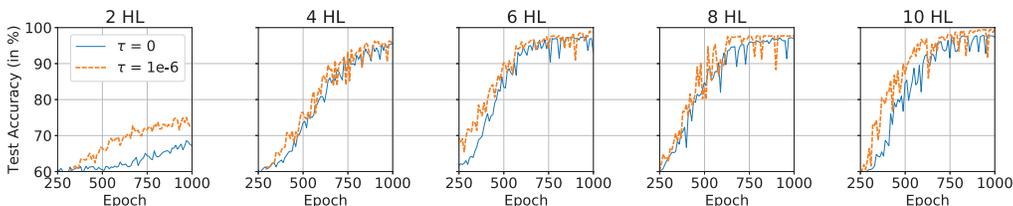


Figure 2: The effect of temperature for the same set-up as for Fig. 1. MLPs with varying numbers of hidden layers (HL) were trained using o-CoLA-od with $h = 0.1$ and either $\tau = 0$ (blue line) or $\tau = 1e-6$ (orange line). Results are averaged over 5 runs. The use of temperature is shown to speed up training and often slightly increases the obtained test accuracies.

For a ResNet-34 architecture with BatchNorm and learning rate (LR) decay on CIFAR-10 [19] data our underdamped orthogonal constrained method, o-CoLA-ud without weight decay (WD) significantly outperforms SGD-m without WD (Fig. 3). In future work we will explore the nuances of combining orthogonality constraints with BatchNorm, residual connections and LR decay. Since o-CoLA outperforms SGD if no LR decay is used, we expect that with more tuning the use of WD can be completely removed by using orthogonality constraints (see also Fig. C6).

Circle Constraints We evaluate our circle constrained c-CoLA-ud method on the Fashion-MNIST data set [48]. We reduce the amount of training data to 10K samples and use the remaining 60K samples as test data. c-CoLA-ud clearly outperforms SGD-m in terms of both test accuracy and test loss for a 1000-node single hidden layer perceptron (see Fig. 4). The lower test loss of c-CoLA-ud is maintained during training and the method shows no signs of overfitting, thus eliminating the need for early stopping. Even with weight decay, SGD-m is outperformed by its constrained counterpart (see Appx. C.2). We also show that a small transformer [43] with 2 encoder layers (each with 2-head self-attention and a 200-node feed-forward network) trained using c-CoLA-ud achieves a lower validation loss on NLP datasets than its unconstrained counterpart, SGD-m (see Table 1).

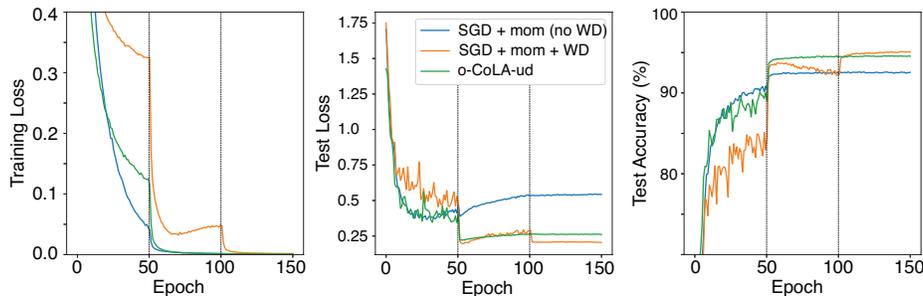


Figure 3: Train (left) & test (middle) loss and test acc. (right) averaged over 5 runs of a ResNet-34 with BatchNorm trained using SGD-m vs. o-CoLA-ud with $\tau = 0$ on CIFAR-10. For SGD we initially use $h = 0.1$ and decay by a factor 10 every 50 epochs (indicated by the vertical black dotted lines). We set momentum = 0.9 and present results with and without WD. o-CoLA-ud (with $\gamma = 0.5$) did not use WD. Its learning rate was re-scaled to match the parameters of SGD-m and used the same LR schedule. The o-CoLA-ud method without weight decay strongly outperforms SGD-m without weight decay.

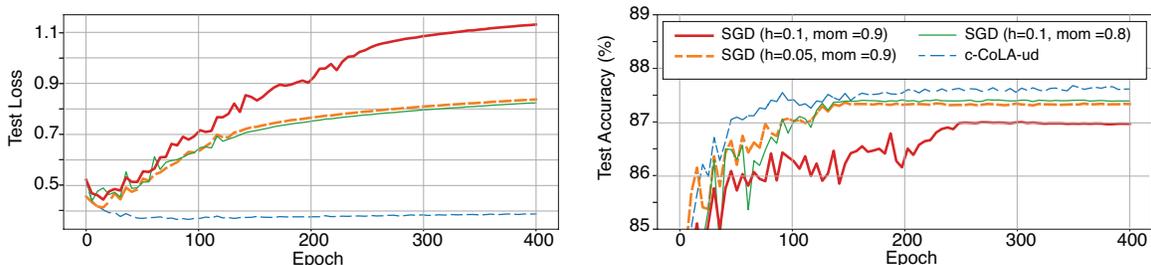


Figure 4: Test loss (left) and test acc. (right) averaged over 5 runs of a 1000-node SHLP trained using SGD-m vs. c-CoLA-ud with $\tau = 0$ on Fashion-MNIST (batchsize 128, # of training data samples reduced to 10K). Hyperparameters of c-CoLA-ud: $h = 0.3, \gamma = 1, r_0 = 0.05, r_1 = 0.1$. Due to the small training data set size both methods quickly reached 100% training accuracy, but c-CoLA-ud is superior in its test loss and test accuracy.

| Data sets | c-CoLA-ud | SGD $h = 0.1$ | | | SGD $h = 0.2$ | | |
|---------------|-------------|---------------|------|------|---------------|------|------|
| | | $mom = 0.7$ | 0.8 | 0.9 | $mom = 0.7$ | 0.8 | 0.9 |
| Penn Treebank | 4.81 | 4.87 | 4.83 | 4.84 | 4.83 | 4.83 | 4.83 |
| Wikitext-2 | 5.09 | 5.13 | 5.13 | 5.13 | 5.13 | 5.14 | 5.13 |

Table 1: Minimum validation loss on Penn Treebank data (batchsize 1024) [29] and Wikitext-2 (batchsize 128) [30] using a transformer trained using c-CoLA-ud (with $\tau = 0$) or SGD-m. Hyperparameters c-CoLA-ud: $h = 0.4, r = 0.5, r_L = 0.1, r_N = 1, r_A = 1, \gamma = 0.5$ (Treebank) and $\gamma = 1$ (Wikitext-2), where the subscripts L, N, A represent the radii belonging to the linear, norm and self-attention layers respectively. The transformer trained using o-CoLA-ud obtains lower validation losses on both datasets.

Acknowledgements

The authors wish to thank Gabriel Stoltz and Tony Lelièvre for helpful discussions on constrained SDEs. Benedict Leimkuhler is a fellow of the Alan Turing Institute which is supported by EPSRC grant EP/N510129/1. Timothée Pouchon is supported by the Swiss National Science Foundation, project P2ELP2_188037. Tiffany Vlaar is supported by The Maxwell Institute Graduate School in Analysis and its Applications, a Centre for Doctoral Training funded by the UK Engineering and Physical Sciences Research Council (grant EP/L016508/01), the Scottish Funding Council, Heriot-Watt University and the University of Edinburgh.

References

- [1] M. Arjovsky, A. Shah, and Y. Bengio. Unitary evolution recurrent neural networks. In *International Conference on Machine Learning*, pages 1120–1128, 2016.
- [2] D. Bakry and M. Émery. Diffusions hypercontractives. In J. Azéma and M. Yor, editors, *Séminaire de Probabilités XIX 1983/84*, pages 177–206, Berlin, Heidelberg, 1985. Springer Berlin Heidelberg. ISBN 978-3-540-39397-9. doi: 10.1007/BFb0075847.
- [3] D. Bakry, I. Gentil, and M. Ledoux. *Analysis and geometry of Markov diffusion operators*, volume 348. Springer Science & Business Media, 2013.
- [4] N. Bansal, X. Chen, and Z. Wang. Can we gain more from orthogonality regularizations in training deep CNNs? In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 4266–4276. Curran Associates Inc., 2018.
- [5] A. Beskos, N. Pillai, G. Roberts, J.-M. Sanz-Serna, and A. Stuart. Optimal tuning of the hybrid Monte Carlo algorithm. *Bernoulli*, 19(5A):1501–1534, 2013.
- [6] R. N. Bhattacharya. On the functional central limit theorem and the law of the iterated logarithm for Markov processes. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 60(2):185–201, 1982. doi: 10.1007/BF00531822.
- [7] N. Bou-Rabee and J.M. Sanz-Serna. Geometric integrators and the Hamiltonian Monte Carlo method. *Acta Numerica*, 27:113–206, 2018. doi: 10.1017/S0962492917000101.
- [8] A. Brock, T. Lim, J. M. Ritchie, and N. J. Weston. Neural photo editing with introspective adversarial networks. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017.
- [9] X. Cheng, N. S. Chatterji, P. L. Bartlett, and M. I. Jordan. Underdamped Langevin MCMC: A non-asymptotic analysis. *arXiv:1707.03663*, 2017.
- [10] S. d’Ascoli, L. Sagun, J. Bruna, and G. Biroli. Finding the needle in the haystack with convolutions: on the benefits of architectural bias. *NeurIPS*, 2019.
- [11] E. Faou and T. Lelièvre. Conservative stochastic differential equations: Mathematical and numerical analysis. *Mathematics of computation*, 78(268):2047–2074, 2009. doi: 10.1090/S0025-5718-09-02220-0.

- [12] M. Graham and A. Storkey. Asymptotically exact inference in differentiable generative models. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54, pages 499–508, 2017.
- [13] K. He, X. Zhang, S. Ren, and J. Sun. Delving deep into rectifiers: Surpassing human-level performance on Imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034, 2015.
- [14] A. Hoerl and R. Kennard. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12:55–67, 1970. doi: 10.1080/00401706.1970.10488634.
- [15] L. Huang, X. Liu, B. Lang, A. Wei Yu, and B. Li. Orthogonal weight normalization: Solution to optimization over multiple dependent stiefel manifolds in deep neural networks. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [16] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International Conference on Machine Learning*, pages 448–456, 2015.
- [17] K. Jia, S. Li, Y. Wen, T. Liu, and D. Tao. Orthogonal deep neural networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019. doi: 10.1109/TPAMI.2019.2948352.
- [18] C. Kipnis and S. R. S. Varadhan. Central limit theorem for additive functionals of reversible Markov processes and applications to simple exclusions. *Communications in Mathematical Physics*, 104(1):1–19, 1986. doi: 10.1007/BF01210789.
- [19] A. Krizhevsky and G. Hinton. Learning multiple layers of features from tiny images. 2009.
- [20] J. M. Lee. *Introduction to Riemannian manifolds*, volume 2. Springer, 2018.
- [21] B. Leimkuhler and C. Matthews. Efficient molecular dynamics using geodesic integration and solvent–solute splitting. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 472(2189):20160138, 2016. doi: 10.1098/rspa.2016.0138.
- [22] B. Leimkuhler and S. Reich. *Simulating Hamiltonian dynamics*, volume 14. Cambridge university press, 2004.
- [23] B. Leimkuhler, C. Matthews, and G. Stoltz. The computation of averages from equilibrium and nonequilibrium Langevin molecular dynamics. *IMA Journal of Numerical Analysis*, 36(1):13–79, 2016. doi: 10.1093/imanum/dru056.
- [24] B. Leimkuhler, C. Matthews, and T. Vlaar. Partitioned integrators for thermodynamic parameterization of neural networks. *Foundations of Data Science*, 1(4):457–489, 2019. doi: 10.3934/fods.2019019.
- [25] T. Lelièvre and G. Stoltz. Partial differential equations and stochastic methods in molecular dynamics. *Acta Numerica*, 25:681–880, 2016. doi: 10.1017/S0962492916000039.
- [26] T. Lelièvre, G. Stoltz, and M. Rousset. *Free energy computations: A mathematical perspective*. Imperial College Press, 2010. ISBN 9781848162488.

- [27] T. Lelièvre, M. Rousset, and G. Stoltz. Langevin dynamics with constraints and computation of free energy differences. *Mathematics of computation*, 81(280):2071–2125, 2012. doi: 10.1090/S0025-5718-2012-02594-4.
- [28] T. Lelièvre, G. Stoltz, and W. Zhang. Multiple projection MCMC algorithms on submanifolds. *arXiv:2003.09402*, 2020.
- [29] M. P. Marcus, B. Santorini, and M. A. Marcinkiewicz. Building a large annotated corpus of English: The Penn Treebank. *Computational Linguistics*, 19(2):313–330, 1993.
- [30] S. Merity, C. Xiong, J. Bradbury, and R. Socher. Pointer sentinel mixture models. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017.
- [31] B. Neyshabur, R. Tomioka, and N. Srebro. In search of the real inductive bias: On the role of implicit regularization in deep learning. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Workshop Track Proceedings*, 2015.
- [32] R. Pascanu, T. Mikolov, and Y. Bengio. On the difficulty of training recurrent neural networks. In *International conference on machine learning*, pages 1310–1318, 2013.
- [33] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer. Automatic differentiation in PyTorch. 2017.
- [34] G. A. Pavliotis. *Stochastic processes and applications: diffusion processes, the Fokker-Planck and Langevin equations*, volume 60. Springer, 2014.
- [35] J. Pennington, S. Schoenholz, and S. Ganguli. Resurrecting the sigmoid in deep learning through dynamical isometry: theory and practice. In *Advances in Neural Information Processing Systems*, pages 4785–4795, 2017.
- [36] J. Pennington, S. Schoenholz, and S. Ganguli. The emergence of spectral universality in deep networks. In *International Conference on Artificial Intelligence and Statistics*, pages 1924–1932, 2018.
- [37] P. Rodríguez, J. González, G. Cucurull, J. M. Gonfaus, and X. Roca. Regularizing cnns with locally constrained decorrelations. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017.
- [38] S. Santurkar, D. Tsipras, A. Ilyas, and A. Madry. How does batch normalization help optimization? In *Advances in Neural Information Processing Systems*, pages 2483–2493, 2018.
- [39] A. M. Saxe, J. L. McClelland, and S. Ganguli. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. *arXiv:1312.6120*, 2013.
- [40] N. Srebro and A. Shraibman. Rank, trace-norm and max-norm. In *International Conference on Computational Learning Theory*, pages 545–560. Springer, 2005. doi: 10.1007/11503415_37.

- [41] N. Srivastava, G.E. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014.
- [42] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996. doi: 10.1111/j.2517-6161.1996.tb02080.x.
- [43] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008, 2017.
- [44] C. Villani. Hypocoercivity. *Memoirs of the American Mathematical Society*, 202(950), 2009.
- [45] E. Vorontsov, C. Trabelsi, S. Kadoury, and C. Pal. On orthogonality and learning recurrent networks with long term dependencies. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 3570–3578. JMLR. org, 2017.
- [46] F. Wenzel, K. Roth, B. S. Veeling, J. Swiatkowski, L. Tran, S. Mandt, J. Snoek, T. Salimans, R. Jenatton, and S. Nowozin. How good is the Bayes posterior in deep neural networks really? *arXiv:2002.02405*, 2020.
- [47] P. Williams. Bayesian regularization and pruning using a laplace prior. *Neural computation*, 7(1):117–143, 1995. doi: 10.1162/neco.1995.7.1.117.
- [48] H. Xiao, K. Rasul, and R. Vollgraf. Fashion-MNIST: a novel image dataset for benchmarking machine learning algorithms. *arXiv:1708.07747*, 2017.
- [49] L. Xiao, Y. Bahri, J. Sohl-Dickstein, S. Schoenholz, and J. Pennington. Dynamical isometry and a mean field theory of CNNs: How to train 10,000-layer vanilla convolutional neural networks. In *International Conference on Machine Learning*, pages 5393–5402, 2018.
- [50] D. Xie, J. Xiong, and S. Pu. All you need is beyond a good init: Exploring better solution for training extremely deep convolutional neural networks with orthonormality and modulation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6176–6185, 2017.
- [51] Z. Yao, A. Gholami, K. Keutzer, and M. Mahoney. PyHessian: Neural networks through the lens of the Hessian. *arXiv:1912.07145*, 2019.
- [52] E. Zappa, M. Holmes-Cerfon, and J. Goodman. Monte Carlo on manifolds: Sampling densities and integrating functions. *Communications on Pure and Applied Mathematics*, 71(12):2609–2647, 2018. doi: 10.1002/cpa.21783.

Overview of the provided supplementary material:

Appendix [A](#): Provides the results necessary to establish exponential convergence to equilibrium of constrained overdamped Langevin dynamics ([4](#)).

Appendix [B](#): Provides discretization schemes and implementation details for our constrained training algorithms. The discretization schemes for a general constraint are described in [Appendix B.1](#) for overdamped Langevin dynamics and in [B.2](#) for underdamped Langevin dynamics. Our c-CoLA circle constrained algorithm is discussed in [Appendix B.3](#) (overdamped) and [B.4](#) (underdamped). [Appendix B.5](#) and [B.6](#) are reserved for our o-CoLA, orthogonality constraint Langevin dynamics, algorithm (overdamped and underdamped, respectively).

Appendix [C](#): Provides further numerical implementation details and results for our constrained methods.

Appendix A. Theory of constrained overdamped Langevin dynamics

We present here the details of the theory summarized in Sec. 3. In particular, we provide the key results and suitable references to establish the exponential convergence to equilibrium of constrained overdamped Langevin dynamics (4).

In the first part (Sec. A.1), we derive the underlying SDE associated with (4), its generator and the invariant measure ν_Σ defined as

$$d\nu_\Sigma = Z^{-1} e^{-\beta V(q)} d\sigma_\Sigma, \quad Z = \int_\Sigma e^{-\beta V(q)} d\sigma_\Sigma, \quad (11)$$

where σ_Σ is the surface measure on Σ . Ergodicity ensures that averages of observables with respect to ν_Σ can be approximated by time averages of trajectories of (4): for all test function $\phi \in \mathcal{C}_c^\infty(\Sigma)$

$$\lim_{T \rightarrow \infty} \langle \phi \rangle_T = \langle \phi \rangle_{\nu_\Sigma} \quad \text{for a.e. } q_0 \in \Sigma, \quad \langle \phi \rangle_T := \frac{1}{T} \int_0^T \phi(q_t) dt, \quad \langle \phi \rangle_{\nu_\Sigma} := \int_\Sigma \phi(q) d\nu_\Sigma(q). \quad (12)$$

Next, in Sec. A.2 we present the Poincaré inequality on a manifold, which holds under a curvature-dimension assumption: there exists $\rho > 0$ such that

$$CD(\rho, \infty) : \quad \text{Ric}_g + \beta \nabla_g^2 V \geq \rho g, \quad (13)$$

in the sense of symmetric matrices. The terms in (13) rely on the structure of Σ as a Riemannian manifold: g is the Riemannian metric, Ric_g is the Ricci curvature tensor and $\nabla_g^2 V$ is the Hessian of V on the manifold. Under (13) the following result holds [3].

Theorem 1 *Assume that there exists $\rho > 0$ and $N > n$ such that $CD(\rho, N)$ holds. Then ν_Σ satisfies a Poincaré inequality: there exists a constant $L > 0$ such that*

$$\int_\Sigma |\phi(q) - \langle \phi \rangle_{\nu_\Sigma}|^2 d\nu_\Sigma(q) \leq \frac{1}{2L} \int_\Sigma |\Pi(q) \nabla_q \phi(q)|^2 d\nu_\Sigma(q) \quad \forall \phi \in H^1(\nu_\Sigma), \quad (14)$$

where $\Pi(q)$ is the projection onto the cotangent space $T_q^* \Sigma$ (18) and $H^1(\nu_\Sigma)$ is the space of functions with square ν_Σ -integrable gradients (17).

Consequences of Theorem 1 are exponential convergence and a central limit theorem (CLT) for the convergence in (12).

Corollary 2 *If (13) holds then*

$$\int_\Sigma |\mathbb{E}(\phi(q_t) | q_0) - \langle \phi \rangle_{\nu_\Sigma}|^2 d\nu_\Sigma(q_0) \leq C(\phi) e^{-2L/\beta t} \quad \forall \phi \in H^1(\nu_\Sigma), \quad (15)$$

where $C(\phi)$ depends only on ϕ . Furthermore we have the following convergence in law:

$$\sqrt{T}(\langle \phi \rangle_T - \langle \phi \rangle_{\nu_\Sigma}) \rightarrow \mathcal{N}(0, \sigma_\phi^2) \quad \text{as } T \rightarrow \infty,$$

where the asymptotic variance σ_ϕ^2 is bounded as $\sigma_\phi^2 \leq \frac{\beta}{L} \int_\Sigma |\phi - \langle \phi \rangle_{\nu_\Sigma}|^2 d\nu_\Sigma$.

Appx. A.3 is dedicated to using the Poincaré inequality to proving this.

In \mathbb{R}^n , assumption (13) is equivalent to convexity of V , which is known to be too strong a requirement (a confining assumption is sufficient, see e.g. [25]). Although (13) can certainly be weakened, the above results ensure that, provided the curvature of the manifold is well behaved, sampling on Σ has similar properties as on a flat space.

NOTATION

We collect here additional notation needed for this discussion.

Given a measure μ in a space $E \subset \mathbb{R}^d$, we associate the space of square integrable functions

$$L^2(\mu) = \left\{ \phi : E \rightarrow \mathbb{R} \text{ measurable} : \int_E |\phi|^2 d\mu < \infty \right\}.$$

Equipped with the inner product and associated norm

$$\langle \phi, \psi \rangle_\mu = \int_E \phi \psi d\mu, \quad \|\phi\|_{L^2(\mu)} = \sqrt{\langle \phi, \phi \rangle},$$

$L^2(\mu)$ is a Hilbert space. We further define the subspace $L^2_0(\mu)$ of functions with zero mean by

$$L^2_0(\mu) = \left\{ \phi \in L^2(\mu) : \langle \phi \rangle_\mu = 0 \right\}, \quad \langle \phi \rangle_\mu = \int_E \phi d\mu, \quad (16)$$

as well as the space of functions with square integrable gradient

$$H^1(\mu) = \left\{ \phi \in L^2(\mu) : \partial_i \phi \in L^2(\mu) \quad 1 \leq i \leq d \right\}. \quad (17)$$

For the constraint $g : \mathbb{R}^d \rightarrow \mathbb{R}^m$, we denote the Jacobian matrix as $G(q) = \nabla_q^T g(q)$ and denote its right pseudo-inverse by $G^+ = G^T (GG^T)^{-1}$ (GG^T is invertible if G has full row rank). We verify that the map

$$\Pi : \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}, \quad q \mapsto \Pi(q) = I_d - G^+(q)G(q), \quad (18)$$

defines for each q the orthogonal projection onto the cotangent space $T_q^* \Sigma$.

$$\Pi_q = \Pi(q) : \mathbb{R}^d \rightarrow \mathbb{R}^d, \quad p \mapsto \Pi(q)p.$$

In particular, for all q we have $\Pi_q p \in T_q^* \Sigma$ and the matrix Π_q is symmetric and idempotent: (i.e., $\Pi_q^T = \Pi_q$ and $\Pi_q^2 = \Pi_q$).

A.1. The underlying SDE and the invariant measure

Although presented differently, the results of this section follow closely the treatment of this issue presented in [26, Chap. 3].

We define the mean curvature of the manifold as the vector valued function

$$\mathcal{H} : \mathbb{R}^d \rightarrow \mathbb{R}^d, \quad q \mapsto (\mathcal{H}(q))_i = \Pi_{jk}(q) \partial_j \Pi_{ik}(q) \quad 1 \leq i \leq d, \quad (19)$$

where $\Pi(q) : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is the projection onto the cotangent space defined in (18). We then establish the following result (proved below).

Lemma 3 *The constrained system (4) can be rewritten as the following SDE in \mathbb{R}^d*

$$dq_t = -\Pi(q_t) \nabla V(q_t) dt + \sqrt{2\beta^{-1}} \Pi(q_t) d\mathcal{W}_t + \beta^{-1} \mathcal{H}(q_t) dt. \quad (20)$$

The uniqueness of the invariant measure of (20) and the resulting ergodicity result (12) are proved in [26, Prop. 3.20] (the proof relies on the divergence theorem on manifolds).

The generator associated with (20) is given by

$$\mathcal{L} = -\Pi(q)\nabla V(q) \cdot \nabla + \beta^{-1}\mathcal{H}(q) \cdot \nabla + \beta^{-1}\Pi(q) : \nabla^2.$$

We verify that \mathcal{L} can be written in the following symmetric form

$$\mathcal{L}\psi = \beta^{-1} \operatorname{div}_\Sigma(\nabla_\Sigma\psi) - \nabla_\Sigma V(q) \cdot \nabla_\Sigma\psi = \beta^{-1} e^{\beta V(q)} \operatorname{div}_\Sigma(e^{-\beta V(q)} \nabla_\Sigma\psi), \quad (21)$$

where we denote $\nabla_\Sigma\phi = \Pi\nabla\phi$ and $\operatorname{div}_\Sigma\psi = \nabla_\Sigma \cdot \psi = \sum_{i,j=1}^d \Pi_{ij}\partial_j\psi_i$. This expression directly implies that \mathcal{L} is reversible with respect to ν_Σ :

$$\langle \mathcal{L}\phi, \psi \rangle_{\nu_\Sigma} = -\beta^{-1} \langle \nabla_\Sigma\phi, \nabla_\Sigma\psi \rangle_{\nu_\Sigma} = \langle \phi, \mathcal{L}\psi \rangle_{\nu_\Sigma}. \quad (22)$$

Thanks to this expression, we can prove that the measure ν_Σ is indeed invariant for (4). Let us introduce the forward Kolmogorov equation: given a test function $\phi \in \mathcal{C}_c^\infty(\Sigma)$

$$\partial_t u(t, q) = \mathcal{L}u(t, q) \quad t \geq 0, \quad q \in \Sigma \quad u(0, q) = \phi(q).$$

The solution to this equation is verified to be $u(t, q) = \mathbb{E}(\phi(q_t) \mid q_0 = q)$ (see the Feynmann–Kac formula) and is usually denoted as $u(t, q) = e^{t\mathcal{L}}\phi(q)$. The measure ν_Σ is invariant if for any $t \geq 0$ $\int_\Sigma u(t, q) d\nu_\Sigma(q) = \int_\Sigma u(0, q) d\nu_\Sigma(q) = \langle \phi \rangle_{\nu_\Sigma}$. This is easily verified thanks to (22):

$$\frac{d}{dt} \int_\Sigma u(t, q) d\nu_\Sigma(q) = \frac{d}{dt} \int_\Sigma e^{t\mathcal{L}}\phi(q) d\nu_\Sigma(q) = \int_\Sigma \mathcal{L}e^{t\mathcal{L}}\phi(q) d\nu_\Sigma(q) = \langle \mathcal{L}e^{t\mathcal{L}}\phi, \mathbf{1} \rangle_{\nu_\Sigma} = 0.$$

Proof Let us write λ_t as the Itô process

$$d\lambda_t = \mu(q_t) dt + \sigma(q_t) d\mathcal{W}_t, \quad (23)$$

where $\mu : \mathbb{R}^d \rightarrow \mathbb{R}^m$, $\sigma : \mathbb{R}^d \rightarrow \mathbb{R}^{m \times d}$ and \mathcal{W}_t is the same Wiener process as in (4). Using this expression in (4) brings

$$dq_t = (-\nabla V(q_t) - G(q_t)^T \mu(q_t)) dt + (\sqrt{2\beta^{-1}}I - G(q_t)^T \sigma(q_t)) d\mathcal{W}_t,$$

where we recall the notation for the Jacobian $G = \nabla_q^T g$. Using Itô formula we find

$$0 = dg(q_t) = G(q_t) dq + b_t dt = G(q_t) (-\nabla V(q_t) dt + \sqrt{2\beta^{-1}} d\mathcal{W}_t - G(q_t)^T d\lambda_t) + b_t dt, \quad (24)$$

where b_t is the d -dimensional process defined as (omitting the dependence on q_t)

$$\begin{aligned} (b_t)_i &= \frac{1}{2} (\sqrt{2\beta^{-1}}I - G^T \sigma) (\sqrt{2\beta^{-1}}I - G^T \sigma)^T : \nabla^2 g_i \\ &= \beta^{-1} \Delta g_i - \frac{\sqrt{2\beta^{-1}}}{2} (G^T \sigma + \sigma^T G) : \nabla^2 g_i + \frac{1}{2} G^T \sigma \sigma^T G : \nabla^2 g_i. \end{aligned} \quad (25)$$

From (24) yields

$$d\lambda_t = (G(q_t)G(q_t)^T)^{-1} G(q_t) \left(-\nabla V(q_t) dt + \sqrt{2\beta^{-1}} d\mathcal{W}_t \right) + (G(q_t)G(q_t)^T)^{-1} b_t dt. \quad (26)$$

Identifying with (23) we find $\sigma(q) = \sqrt{2\beta^{-1}}(G^+(q))^T$, which used in (25) yields

$$(b_t)_i = \beta^{-1}(\Delta g_i - (G^T(G^+)^T + G^+G) : \nabla^2 g_i + G^T(G^+)^T G^+G : \nabla^2 g_i).$$

As G^+G is symmetric and $GG^+ = I_m$, we obtain

$$(b_t)_i = \beta^{-1}(\Delta g_i - G^+G : \nabla^2 g_i) = \beta^{-1}\Pi : \nabla^2 g_i. \quad (27)$$

Inserting (26) in (4) brings

$$dq_t = -\Pi(q_t)\nabla V(q_t)dt + \sqrt{2\beta^{-1}}\Pi(q_t)d\mathcal{W}_t - G^+(q_t)b_t dt. \quad (28)$$

To conclude the proof we require the following technical relations on the mean curvature vector ((29a) follows from a direct computation; the proof of (29b) is direct but involved and can be found in [26, Lemma 3.15]).

Lemma 4 *The projection Π and the vector \mathcal{H} defined in (18) and (19) satisfy the following equalities*

$$\mathcal{H} = (I - \Pi)\nabla \cdot \Pi, \quad (29a)$$

$$\Pi : \nabla^2 g_i = -(G\mathcal{H})_i \quad 1 \leq i \leq d, \quad (29b)$$

Equality (29a) ensures that $\Pi\mathcal{H} = 0$. Combining (27) and (29b) we can write $b_t = -\beta^{-1}G\mathcal{H}$. Thanks to these relations and the definition of Π , we obtain

$$-G^+b_t = \beta^{-1}G^+G\mathcal{H} = \beta^{-1}(I - \Pi)\mathcal{H} = \beta^{-1}\mathcal{H}.$$

This equality combined with (28) proves (3) and concludes the proof of Lemma 3. ■

A.2. Poincaré inequality on a manifold

Poincaré inequalities, also called spectral gap inequalities, form an important family of functional inequalities in the theory of Markov diffusion processes. They are the simplest inequalities that provide results on the convergence to equilibrium. Stronger results can be obtained with the family of log-Sobolev inequalities, which are at the center of the Bakry–Émery theory [2]. We follow here closely the book [3] on this subject (more specifically §1.16.2 and sections 4.2, 4.8, C.6). For the necessary terminology of Riemannian manifolds we recommend the introductory textbook [20] (the literature on this topic is vast and contains many works of high quality).

As presented in [3, Chap. 4], a Poincaré inequality can be obtained as a consequence of a curvature-dimension condition. For the sake of presentation, we introduce this result in the setting of a weighted Riemannian manifold. Let $(\mathcal{M}, \mathfrak{g})$ be an n -dimensional Riemannian manifold, where \mathfrak{g} is the Riemannian metric. We consider the diffusion operator

$$\mathcal{L} = \Delta_{\mathfrak{g}} - \langle \nabla_{\mathfrak{g}} W, \nabla_{\mathfrak{g}} \cdot \rangle_{\mathfrak{g}},$$

where $\Delta_{\mathfrak{g}}$ denotes the Laplace–Beltrami operator on the manifold \mathcal{M} , $\nabla_{\mathfrak{g}}$ denotes the Levi–Civita connection (covariant derivative) and $\langle \cdot, \cdot \rangle_{\mathfrak{g}}$ denotes the Riemannian metric ($\langle X, Y \rangle_{\mathfrak{g}} = \mathfrak{g}(X, Y)$)

for all vector fields X, Y). We verify that the associated invariant measure is $d\mu = Z^{-1}e^{-W}d\mu_{\mathfrak{g}}$, where $d\mu_{\mathfrak{g}}$ is the Riemannian measure [3, §1.11.3]. For $N \in [n, \infty]$, we define the 2-tensor

$$\text{Ric}_N(\mathcal{L}) = \text{Ric}_{\mathfrak{g}} + \nabla_{\mathfrak{g}}^2 W - \frac{1}{N-n} dW \otimes dW.$$

where $\text{Ric}_{\mathfrak{g}}$ is the Ricci curvature 2-tensor and $\nabla_{\mathfrak{g}}^2$ denotes the Hessian operator on \mathcal{M} (the case $N = n$ is considered only if W is constant). In this context, a curvature-dimension condition $CD(\rho, N)$ for $\rho \in \mathbb{R}$ and $N \geq n$ holds if and only if (see [3, C.6])

$$CD(\rho, N) : \quad \text{Ric}_N(\mathcal{L}) \geq \rho \mathfrak{g}, \quad (30)$$

in the sense of symmetric $(0, 2)$ -tensors (covariant 2-tensors). In the flat space $\mathcal{M} = \mathbb{R}^n$, the condition $CD(\rho, \infty)$ reads $\nabla^2 W \geq \rho I$, which is nothing but the convexity of the potential W . Under $CD(\rho, N)$, the measure μ is proved to satisfy a Poincaré inequality (in [3], combine Thm 4.8.4 with the discussion in section C.6).

Theorem 5 [3, Thm 4.8.4] *Under the curvature-dimension condition $CD(\rho, N)$ with $\rho > 0$ and $N \geq n, N > 1$, the measure μ satisfies the Poincaré inequality*

$$\text{Var}_{\mu}(\phi) = \|\phi - \langle \phi \rangle_{\mu}\|_{L^2(\mu)}^2 \leq C_P \|\nabla_{\mathfrak{g}} \phi\|_{L^2(\mu)}^2 \quad \forall \phi \in L^2(\mu) \cap H^1(\mu), \quad (31)$$

with constant $C_P = \frac{N-1}{\rho N}$.

As the tensor $dW \otimes dW$ is positive semi-definite, we verify the monotonicity $\text{Ric}_{N+M}(\mathcal{L}) \geq \text{Ric}_N(\mathcal{L})$ for any $M \geq 0$. This implies in particular that $CD(\rho, N) \Rightarrow CD(\rho, \infty)$ for any $N \in [n, \infty]$. Hence, among all choices of $N \geq n$, $CD(\rho, \infty)$ is the weaker condition.

Let us now consider this result in the context of the constraint manifold Σ in (1). We consider the space \mathbb{R}^d with its Riemannian manifold structure given by the Euclidean metric $\bar{\mathfrak{g}}(v, w) = v \cdot w$ for all $v, w \in \mathbb{R}^d$ (for all $q \in \mathbb{R}^d, p \in T_q \mathbb{R}^d$ is identified with \mathbb{R}^d through a canonical isomorphism). Assuming that g is smooth and that $\nabla_q^T g$ has everywhere full row-rank, Σ is a smooth embedded submanifold of \mathbb{R}^d of dimension $n = d - m$ (see e.g. [20, Cor. A.26]). Furthermore, Σ is equipped with the metric induced by $\bar{\mathfrak{g}}$: for a local parameterization of $\psi : U \subset \Sigma \rightarrow \mathbb{R}^d$, $\bar{\mathfrak{g}}$ is given locally on U by

$$\bar{\mathfrak{g}} = \sum_{i=1}^d \sum_{j,k=1}^n \frac{\partial \psi^i}{\partial x^j} \frac{\partial \psi^i}{\partial x^k} dx^j dx^k = (\nabla_x \psi \nabla_x^T \psi)_{jk} dx^j dx^k. \quad (32)$$

We now define the potential $W = \beta V|_{\Sigma}$, where $V|_{\Sigma}$ denotes the restriction of V to Σ . Assumption 13 corresponds then to condition $CD(\rho, \infty)$ above. Applying Theorem 5 we obtain Poincaré's inequality on the constraint manifold Σ . We note that for a function ϕ defined on \mathbb{R}^d , the covariant derivative in \mathbb{R}^d of $\phi|_{\Sigma}$ on the manifold is the orthogonal projection of the directional derivative of ϕ (in the ambient manifold \mathbb{R}^d) onto the cotangent space: $\nabla_{\bar{\mathfrak{g}}}(\phi|_{\Sigma})(q) = \Pi(q) \nabla_q \phi(q)$. Furthermore, we note that the surface measure σ_{Σ} equals the Riemannian measure on the manifold (compare [26, Rem. 3.4] with [20, Prop. 2.41] and (32)). We thus obtain the result of Theorem 1 with constant $C_P = \frac{1}{\rho} = \frac{1}{2L}$.

A.3. Exponential convergence to equilibrium and central limit theorem

Let us define the norm of a linear operator $\mathcal{A} : L_0^2(\nu_\Sigma) \rightarrow L_0^2(\nu_\Sigma)$ as

$$\|\mathcal{A}\|_{\mathcal{B}(L_0^2(\nu_\Sigma))} = \sup_{\phi \in L_0^2(\nu_\Sigma)} \frac{\|\mathcal{A}\phi\|_{L_0^2(\nu_\Sigma)}}{\|\phi\|_{L_0^2(\nu_\Sigma)}}.$$

Denote $\bar{\phi} = \phi - \langle \phi \rangle_{\nu_\Sigma} \in L_0^2(\nu_\Sigma)$. The Poincaré inequality (14), rewritten on the subspace $L_0^2(\nu_\Sigma)$, is as follows:

$$\|\bar{\phi}\|_{L_0^2(\nu_\Sigma)}^2 \leq \frac{1}{2L} \|\nabla_\Sigma \bar{\phi}\|_{L_0^2(\nu_\Sigma)}^2 \quad \forall \bar{\phi} \in L_0^2(\nu_\Sigma) \cap H^1(\nu_\Sigma). \quad (33)$$

Using the reversibility of the measure (22), we can prove the following result (the proof follows the same lines as [25, Prop. 2.3], see also [3, Thm 4.2.5]).

Lemma 6 *The measure ν_Σ satisfies the Poincaré inequality (33) if and only if*

$$\|e^{t\mathcal{L}}\|_{\mathcal{B}(L_0^2(\nu_\Sigma))} \leq e^{-2\frac{L}{\beta}t}. \quad (34)$$

Exponential convergence to equilibrium is then directly obtained from Lemma 6:

$$\|e^{t\mathcal{L}}\bar{\phi}\|_{L_0^2(\nu_\Sigma)} \leq \|e^{t\mathcal{L}}\|_{\mathcal{B}(L_0^2(\nu_\Sigma))} \|\bar{\phi}\|_{L_0^2(\nu_\Sigma)} \leq e^{-2\frac{L}{\beta}t} \|\bar{\phi}\|_{L_0^2(\nu_\Sigma)}. \quad (35)$$

This inequality implies (15) (note that $e^{t\mathcal{L}}\langle \phi \rangle_{\nu_\Sigma} = \langle \phi \rangle_{\nu_\Sigma}$) and thus proves the first assertion of Corollary 2.

A consequence of the exponential convergence to equilibrium (35) is the following central limit theorem for time averages $\langle \phi \rangle_T = \frac{1}{T} \int_0^T \phi(q_t) dt$ (see also [18]).

Theorem 7 [6] *If (35) holds, then the following convergence in law is satisfied*

$$\sqrt{T}(\langle \phi \rangle_T - \langle \phi \rangle_{\nu_\Sigma}) \rightarrow \mathcal{N}(0, \sigma_\phi^2) \quad \text{as } T \rightarrow \infty,$$

where the asymptotic variance σ_ϕ^2 is given by the formula $\sigma_\phi^2 = 2\langle \bar{\phi}, -\mathcal{L}^{-1}\bar{\phi} \rangle$ with $\bar{\phi} = \phi - \langle \phi \rangle_{\nu_\Sigma}$.

To quantify the asymptotic variance, we use the following classical result.

Lemma 8 (e.g., [25, Prop. 2.1]) *If (34) holds, then the generator \mathcal{L} is invertible and the resolvent can be expressed as $-\mathcal{L}^{-1} = \int_0^\infty e^{t\mathcal{L}} dt$ and satisfies the bound $\|\mathcal{L}^{-1}\|_{\mathcal{B}(L_0^2(\nu_\Sigma))} \leq \frac{\beta}{2L}$.*

Using Lemma 8 and Cauchy–Schwartz inequality, the asymptotic variance in Theorem 7 can thus be bounded as

$$\sigma_\phi^2 = 2 \int_\Sigma \bar{\phi}(-\mathcal{L}^{-1}\bar{\phi}) d\nu_\Sigma \leq 2\|\mathcal{L}^{-1}\|_{\mathcal{B}(L_0^2(\nu_\Sigma))} \|\bar{\phi}\|_{L_0^2(\nu_\Sigma)}^2 \leq \frac{\beta}{L} \|\bar{\phi}\|_{L_0^2(\nu_\Sigma)}^2.$$

This estimate completes the proof of the second assertion of Corollary 2.

Appendix B. Discretization of constrained Langevin dynamics

We present here the details of the constrained training methods considered in this paper. Both the overdamped (4) and underdamped (5) Langevin dynamics are discretized for the constraints presented in Section 2. We emphasize that the initialization of each given method must be done with care: the constrained parameters, the potential slack variable, as well as their momenta in the underdamped case, have to satisfy the constraint initially.

Recall the notation introduced in Section 2: $\theta \in \mathbb{R}^{|n|}$ is the vector of all the parameters of the model, we consider the variable $q = (\theta, \xi) \in \mathbb{R}^d$, $d = |n| + n^\xi$, where $\xi \in \mathbb{R}^{n^\xi}$ is a slack variable to enforce the potential inequality constraints. The loss is extended $q = (\theta, \xi)$ as $V(q) = L_X(\theta)$ (in particular $\nabla_\xi V = 0$) and constraints are given by a map $g : \mathbb{R}^d \rightarrow \mathbb{R}^m$. The parameters are partitioned as $\theta = (\theta^u, \theta^c)$, where $\theta^u \in \mathbb{R}^{n^u}$ are not involved in any constraint while $\theta^c \in \mathbb{R}^{n^c}$ are.

B.1. Discretization of constrained overdamped Langevin (general constraint)

Following [26, Chap. 3] a simple discretization of the constrained overdamped Langevin dynamics (4) is given by the iteration $q_n \in \Sigma \mapsto q_{n+1}$ defined as

$$\begin{aligned} \bar{q}_{n+1} &= q_n - \nabla_q V(q_n)h + \sqrt{2\beta^{-1}h} R_n, & q_{n+1} &= \bar{q}_{n+1} - \nabla_q g(q_n)\lambda_n, \\ \text{where } \lambda_n \in \mathbb{R}^m & \text{ is such that } g(q_{n+1}) = 0, \end{aligned} \quad (36)$$

where $R_n \sim N(0, I)$ is a vector of iid standard normal random variable. The first step of (36), \bar{q}_{n+1} , is an Euler–Maruyama step for standard overdamped Langevin. As \bar{q}_{n+1} in \mathbb{R}^d is generally not on the constrained manifold Σ , the last term is present to project \bar{q}_{n+1} back onto Σ , ensuring $g(q_{n+1}) = 0$. In particular, for the unconstrained parameter we have $\nabla_{\theta^u}^T g = 0_{m \times n^u}$ which implies that $\theta_{n+1}^u = \bar{\theta}_{n+1}^u$ is a standard EM step.

In general, projecting back onto the manifold Σ , i.e., finding λ_n , can be done using root-finding algorithms. Nevertheless, for certain constraints g the roots can be found explicitly. This is the case for the circle constraint (2) (see Section B.3). A potential weakness of method (36) is that the projection process can be guaranteed only for small enough step size h (i.e. \bar{q}_n must be close to Σ). Indeed, even for the circle constraint if h is too large it might not be possible to project \bar{q}_{n+1} back onto the circle following the direction $\nabla_q g(q_n)$. See [28] for some discussion of methods to allow computation to be performed in the large timestep regime.

An alternative method is given by the iteration $q_n \in \Sigma \mapsto q_{n+1} \in \Sigma$ defined as in [26, Chap. 3]

$$\begin{aligned} \bar{q}_{n+1} &= q_n - \nabla_q V(q_n)dt + \sqrt{2\beta^{-1}h} R_n, & q_{n+1} &= \bar{q}_{n+1} - \nabla_q g(q_{n+1})\lambda_n, \\ \text{where } \lambda_n \in \mathbb{R}^m & \text{ is such that } g(q_{n+1}) = 0, \end{aligned} \quad (37)$$

where $R_n \sim N(0, I)$ is a vector of iid standard normal random variable. The projection used in method (37) is in general more robust. The circle constraint is a good illustration of this: while in (36) we project following an oblique direction, in (37) the projection is orthogonal and always exists (see Section B.3).

B.2. Discretization of constrained underdamped Langevin (general constraint)

We next consider the discretization of the constrained underdamped Langevin dynamics (5) where we denote by $p = (p^u, p^c, p^\xi) \in \mathbb{R}^{n^u + n^c + n^\xi}$ the momenta associated with the configuration $q =$

$(\theta^u, \theta^c, \xi)$. Following [21], the system is split into A,B,O components (7)-(9), where B represents a projected impulse defined by the loss gradient (restricted to the cotangent space), O represents a projected stochastic impulse, and A represents evolution along geodesics (i.e., for circle constraints, these are rotations on the circles).

As in the overdamped case, the equality $\nabla_{\theta^u}^T g = 0_{m \times n^u}$ ensures that the unconstrained parameters and their momenta (θ^u, p^u) evolve following the A,B,O steps for unconstrained underdamped Langevin (see [23]). As the B and O components only involve a variation in the momentum p_t and because the constraint only involves q_t , they can be solved exactly for any constraint. The A component involves a variation of the configuration q_t and thus cannot be solved exactly (in law) for any constraint. However, as this part does not include any force evaluation (which would require back-propagation to compute the gradient), it can be approximated cheaply using a few steps of standard well-known schemes such as SHAKE or RATTLE (see Section B.6 for orthogonal constraints). Furthermore, for simple constraints such as the circle constraint (2) the A component can be solved explicitly (see Section B.4).

Let us present the details of the B and O steps. For convenience, let us introduce the following notation for the variables involved in the constraint $w = (\theta^c, \xi) \in \mathbb{R}^{n^c + n^\xi}$ and associated momentum $p^w = (p^c, p^\xi) \in \mathbb{R}^{n^c + n^\xi}$. The projection onto the cotangent space (18) is then as

$$\Pi(q) = I_d - \begin{pmatrix} 0 & 0 \\ 0 & \Pi_w(q) \end{pmatrix}, \quad \text{with } \Pi_w = \begin{pmatrix} g_{\theta^c}^T H^{-1} g_{\theta^c} & g_\xi^T H^{-1} g_{\theta^c} \\ g_{\theta^c}^T H^{-1} g_\xi & g_\xi^T H^{-1} g_\xi \end{pmatrix}, \quad (38)$$

where we have denoted the partial Jacobians by $g_{\theta^c} = \nabla_{\theta^c}^T g \in \mathbb{R}^{m \times n^c}$, $g_\xi = \nabla_\xi^T g \in \mathbb{R}^{m \times n^\xi}$ and the matrix $H = g_{\theta^c} g_{\theta^c}^T + g_\xi g_\xi^T \in \mathbb{R}^{m \times m}$.

B component. Given $q_0, p_0 \in T^*\Sigma$ and a time $t > 0$, we have

$$q_t = q_0, \quad p_t = p_0 - t \nabla_q V(q_0) - \nabla_q g(q_0) (\mu_t - \mu_0),$$

where μ_t is such that $p_t \in T_{q_t}^*\Sigma$ (i.e., it satisfies the constraint $0 = \nabla_q g(q_t) p_t$). Note that as q_0, p_0 satisfy the constraints we have $\mu_0 = 0$. Projecting onto the cotangent space $T_{q_t}^*\Sigma = T_{q_0}^*\Sigma$ and using $\Pi(q_0) \nabla_q g(q_0) = 0$ and $p_0 = \Pi(q_0) p_0$, we obtain

$$p_t = \Pi(q_t) p_t = \Pi(q_0) (p_0 - t \nabla_q V(q_0) - \nabla_q g(q_0) \mu_t) = p_0 - t \Pi(q_0) \nabla_q V(q_0).$$

The B step is thus obtained for a chosen stepsize $h > 0$ as: given $q_n = (\theta_n^u, \theta_n^c, \xi_n) \in \Sigma$ and $p_n = (p_n^u, p_n^c, p_n^\xi) \in T_{q_n}^*\Sigma$

$$\begin{aligned} & \theta_{n+1}^u = \theta_n^u, \quad \theta_{n+1}^c = \theta_n^c, \quad \xi_{n+1} = \xi_n, \\ \text{(B, gen.)} \quad & p_{n+1}^u = p_n^u - h \nabla_{\theta^u} L_X(\theta_n), \quad \bar{p}_{n+1}^c = p_n^c - h \nabla_{\theta^c} L_X(\theta_n), \quad \bar{p}_{n+1}^\xi = p_n^\xi, \\ & \begin{pmatrix} p_{n+1}^c \\ p_{n+1}^\xi \end{pmatrix} = \Pi_w(w_n) \begin{pmatrix} \bar{p}_{n+1}^c \\ \bar{p}_{n+1}^\xi \end{pmatrix} \quad \text{where } w_n = \begin{pmatrix} \theta_n^c \\ \xi_n \end{pmatrix}. \end{aligned} \quad (39)$$

O component. Similarly as for the B part, the O part can be solved exactly in law for any constraint. Given $q_0, p_0 \in T^*\Sigma$ and a time $t > 0$, we have

$$q_t = q_0, \quad p_t = p_0 - \gamma \int_0^t p_t dt + \sqrt{2\gamma\tau} \int_0^t d\mathcal{W}_t - \nabla_q g(q_0) \nu_t,$$

where ν_t ensures that $p_t \in T_{q_t}^* \Sigma$. Projecting to the cotangent space $T_{q_0}^* \Sigma = T_{q_0}^* \Sigma$ as before, we obtain

$$p_t = \Pi(q_t)p_t = p_0 - \gamma \int_0^t \Pi(q_0)p_t dt + \sqrt{2\gamma\tau}\Pi(q_0) \int_0^t d\mathcal{W}_t.$$

We thus recognize that p_t is an Ornstein–Uhlenbeck process:

$$p_t \stackrel{\text{law}}{=} \Pi(q_0)(e^{-\gamma t}p_0 + \sqrt{\tau(1 - e^{-2\gamma t})}R), \quad \text{with } R \sim N(0, I_d),$$

where the equality holds in law.

The O step is thus obtained for a chosen stepsize $h > 0$ as: given $q_n = (\theta_n^u, \theta_n^c, \xi_n) \in \Sigma$ and $p_n = (p_n^u, p_n^c, p_n^\xi) \in T_{q_n}^* \Sigma$

$$\begin{aligned} \theta_{n+1}^u &= \theta_n^u, & \theta_{n+1}^c &= \theta_n^c, & \xi_{n+1} &= \xi_n, \\ p_{n+1}^u &= e^{-\gamma h}p_n^u + \sqrt{\tau(1 - e^{-2\gamma h})}R^u, \\ \text{(O, gen.) } \bar{p}_{n+1}^c &= e^{-\gamma h}p_n^c + \sqrt{\tau(1 - e^{-2\gamma h})}R^c, \\ \bar{p}_{n+1}^\xi &= e^{-\gamma h}p_n^\xi + \sqrt{\tau(1 - e^{-2\gamma h})}R^\xi, \\ \begin{pmatrix} p_{n+1}^c \\ p_{n+1}^\xi \end{pmatrix} &= \Pi_w(w_n) \begin{pmatrix} \bar{p}_{n+1}^c \\ \bar{p}_{n+1}^\xi \end{pmatrix} \quad \text{where } w_n = \begin{pmatrix} \theta_n^c \\ \xi_n \end{pmatrix}, \end{aligned} \tag{40}$$

and R^u , R^c , and R^ξ are independent standard normal random variables.

B.3. Circle constraint, overdamped Langevin (c-CoLA-od)

We consider here the circle constraint (2), for which the partial Jacobians are computed as

$$\nabla_q^T g = (\nabla_{\theta^u}^T g, \nabla_{\theta^c}^T g, \nabla_{\xi}^T g) \in \mathbb{R}^{m \times (n^u + n^c + m)}, \quad \partial_{\theta_j^u} g_i = 0, \quad \partial_{\theta_j^c} g_i = 2\theta_i^c \delta_{ij}, \quad \partial_{\xi_j} g_i = 2\xi_i \delta_{ij}, \tag{41}$$

where δ_{ij} is the Kronecker delta.

For this constraint, the projection step in (36) can be computed explicitly. Indeed λ_n can be found by solving the m quadratic equations $0 = g_i(\bar{q}_{n+1} - \nabla_{q_i} g(q_n)\lambda_n)$ $1 \leq i \leq m$. The (potential) two roots of each equation corresponds to the (potential) two projections of \bar{q}_{n+1} onto the circle following the direction $\nabla g_i(q_n) = 2(\theta_{n,i}^c, \xi_{n,i})$. When two roots are found, we may select the one closest to the point of origin $(\theta_{n,i}^c, \xi_{n,i})$. However, if the point to project $(\bar{\theta}_{n+1,i}^c, \bar{\xi}_{n+1,i})$ is too far away from the circle, this oblique projection may not be possible (i.e., the quadratic equation has no real root).

For the circle constraint, method (37) thus leads to a more robust projection process. Indeed, as $\nabla g_i(q_{n+1}) = 2(\theta_{n+1,i}^c, \xi_{n+1,i})$, the direction of the projection is now orthogonal to the circle. To find an expression for the orthogonal projection P of a point $(\bar{\theta}_1, \bar{\xi}_1)$ on the circle, it is easier to use a geometrical approach than to find the Lagrange multipliers:

$$(\theta_1, \xi_1) = P(\bar{\theta}_1, \bar{\xi}_1) = (r_i \cos(\alpha), r_i \sin(\alpha)), \quad \text{where } \alpha = \arctan\left(\frac{\bar{\xi}_1}{\bar{\theta}_1}\right).$$

We obtain the following discretization of the overdamped Langevin with circle constraints. We initialize the parameters of the neural network using standard PyTorch initialization [13, 33], i.e., $\mathcal{U}(-1/\sqrt{N_{in}}, 1/\sqrt{N_{in}})$, where N_{in} is the number of inputs to a layer. The auxiliary variables ξ_i corresponding to the constrained parameters θ_i^c are initialized to obey the constraint $(\theta_i^c)^2 + \xi_i^2 = r_i^2$. For a chosen stepsize $h > 0$ and given a configuration $q_n = (\theta_n^u, \theta_n^c, \xi_n) \in \Sigma$, one step of the method is defined by $q_{n+1} = (\theta_{n+1}^u, \theta_{n+1}^c, \xi_{n+1}) \in \Sigma$ as

$$\begin{aligned}\theta_{n+1,i}^u &= \theta_{n,i}^u - h\partial_{\theta_i^u} L_X(\theta_n) + \sqrt{2\beta^{-1}h}R_i^u, \\ \bar{\theta}_{n+1,i}^c &= \theta_{n,i}^c - h\partial_{\theta_i^c} L_X(\theta_n) + \sqrt{2\beta^{-1}h}R_i^c, \\ \bar{\xi}_{n+1,i} &= \xi_{n,i} + \sqrt{2\beta^{-1}h}R_i^\xi, \\ \alpha_{n,i} &= \arctan\left(\frac{\bar{\xi}_{n+1,i}}{\bar{\theta}_{n+1,i}^c}\right), \\ \theta_{n+1,i}^c &= r_i \cos(\alpha_{n,i}), \\ \xi_{n+1,i} &= r_i \sin(\alpha_{n,i}),\end{aligned}\tag{42}$$

where R_i^u, R_i^c, R_i^ξ are independent standard normal random variables.

B.4. Circle constraint, underdamped Langevin (c-CoLA-ud)

We provide here the full discretization of the underdamped Langevin dynamics in the case of the circle constraint (2).

A component. For the circle constraint we can solve the A step explicitly. First recall that as $\nabla_{\theta^u}^T g = 0$, the unconstrained parameters θ^u are obtained with a standard A step of the unconstrained underdamped Langevin. Let us then focus on solving the constrained components: we denote $w = (\theta^c, \xi), p^w = (p^c, p^\xi)$. Then for $1 \leq i \leq m$ the A step in (7) corresponds to the constrained ODEs

$$\begin{aligned}\dot{w}_i &= p_i^w \\ \dot{p}_i^w &= -2\lambda_i w_i \\ |\theta_i^c|^2 + |\xi_i|^2 &= r_i^2, \quad \theta_i^c p_i^c + \xi_i p_i^\xi = 0.\end{aligned}\tag{43}$$

As these constrained ODEs are uncoupled, let us drop the specification of the index i . By assumption, we are given initial conditions that satisfy the constraint $(w_0, p_0^w) \in T^*\Sigma$. Solving the second order ODE $\ddot{w} = -2\lambda w$, we find that any solution has the form $w_t = R_t^{2\lambda} w_0$, where R_t^ω is a rotation matrix with angular speed ω given with its time derivative as

$$R_t^\omega = \begin{pmatrix} \cos(\omega t) & \sin(\omega t) \\ -\sin(\omega t) & \cos(\omega t) \end{pmatrix}, \quad \dot{R}_t^\omega = \omega \begin{pmatrix} -\sin(\omega t) & \cos(\omega t) \\ -\cos(\omega t) & -\sin(\omega t) \end{pmatrix}.$$

Computing the momentum $p_t^w = \dot{w}_t = \dot{R}_t^\omega w_0$, and using the properties of R_t^ω we verify that w_t, p_t^w satisfy the constraints in (43) ($\|\cdot\|$ denotes the Euclidean norm in \mathbb{R}^2 and \cdot the dot product):

$$\|w_t\|^2 = \|R_t^\omega w_0\|^2 = \|w_0\|^2 = r^2, \quad w_t \cdot p_t^w = w_0^T (R_t^\omega)^T \dot{R}_t^\omega w_0 = 0.$$

We still have to find the angular speed $\omega = 2\lambda$ such that the momentum p_t^w is consistent with its initial value p_0^w (we denote $w_0 = (\theta_0^c, \xi_0)$ and $p_0^w = (p_0^c, p_0^\xi)$):

$$p_0^w = \dot{R}_0^\omega w_0 \Leftrightarrow p_0^c = \omega \xi_0 \quad \text{and} \quad p_0^\xi = -\omega \theta_0^c.$$

We thus find that

$$\xi_0 p_0^c - \theta_0^c p_0^\xi = \omega (|\xi_0|^2 + |\theta_0^c|^2) = \omega r^2 \quad \Leftrightarrow \quad \omega = \frac{1}{r^2} (\xi_0 p_0^c - \theta_0^c p_0^\xi).$$

We have thus found an explicit expression for the solution of the A component for circle constraints (43).

To complete the B and O steps given in (39) and (40), we need an explicit expression for the projection Π_w in (38) (using (41), recall that $m = n^c = n^\xi$):

$$\Pi_w(w) = \begin{pmatrix} I_m - D^{11} & -D^{12} \\ -D^{12} & I_m - D^{22} \end{pmatrix},$$

where $D^{kl} \in \mathbb{R}^{m \times m}$ are the diagonal matrices defined as

$$D_{ii}^{11} = \frac{|\theta_i^c|^2}{|\theta_i^c|^2 + |\xi_i|^2}, \quad D_{ii}^{12} = \frac{\theta_i^c \xi_i}{|\theta_i^c|^2 + |\xi_i|^2}, \quad D_{ii}^{22} = \frac{|\xi_i|^2}{|\theta_i^c|^2 + |\xi_i|^2}.$$

Assuming that $w = (\theta^c, \xi)$ satisfies the constraint, the projection of (\bar{p}^c, \bar{p}^ξ) is thus computed as

$$\begin{pmatrix} p^c \\ p^\xi \end{pmatrix} = \Pi_w(w) \begin{pmatrix} \bar{p}^c \\ \bar{p}^\xi \end{pmatrix}, \quad \text{where} \quad \begin{aligned} p_i^c &= \bar{p}_i^c - \frac{\theta_i^c}{r_i^2} (\theta_i^c \bar{p}_i^c + \xi_i \bar{p}_i^\xi) & 1 \leq i \leq m, \\ p_i^\xi &= \bar{p}_i^\xi - \frac{\xi_i}{r_i^2} (\theta_i^c \bar{p}_i^c + \xi_i \bar{p}_i^\xi) & 1 \leq i \leq m. \end{aligned}$$

Note that in the B step (39), the above expressions can be simplified by combining the simple definition of $(\bar{p}_n^c, \bar{p}_n^\xi)$ with the constraint

$$0 = (\nabla^T g(q)p)_i = 2(\theta_i^c p_i^c + \xi_i p_i^\xi).$$

We provide below the explicit updates for the A, B and O components for circle constraints. We initialize the parameters of the net using standard PyTorch initialization [13, 33]. The auxiliary variables ξ corresponding to the constrained parameters θ^c are initialized to obey the constraint $(\theta^c)^2 + \xi^2 = r^2$, so that $q_0 = (\theta_0^u, \theta_0^c, \xi_0) \in \Sigma$. The momenta, p^u , p^c , and p^ξ , are generated in the same manner as for standard SGD with momentum in PyTorch, i.e., as equal to the initial gradients. Subsequently, the momenta belonging to the constrained variables p^c and to the auxiliary variables p^ξ are projected using Π_w , so that $p_0 = (p_0^u, p_0^c, p_0^\xi) \in T_{q_0}^* \Sigma$. For a stepsize $h > 0$ we obtain

$$\text{(A step, circle)} \quad \begin{cases} \theta_{n+1,i}^u = \theta_{n,i}^u + h p_{n,i}^u, \\ \omega_i = \frac{1}{r_i^2} (\xi_{n,i} p_{n,i}^c - \theta_{n,i}^c p_{n,i}^\xi), \\ \theta_{n+1,i}^c = \cos(\omega_i h) \theta_{n,i}^c + \sin(\omega_i h) \xi_{n,i}, \\ \xi_{n+1,i} = -\sin(\omega_i h) \theta_{n,i}^c + \cos(\omega_i h) \xi_{n,i}, \\ p_{n+1,i}^u = p_{n,i}^u, \\ p_{n+1,i}^c = \omega_i (-\sin(\omega_i h) \theta_{n,i}^c + \cos(\omega_i h) \xi_{n,i}), \\ p_{n+1,i}^\xi = -\omega_i (\cos(\omega_i h) \theta_{n,i}^c + \sin(\omega_i h) \xi_{n,i}), \end{cases}$$

$$\begin{aligned}
 \text{(B step, circle)} \quad & \begin{cases} \theta_{n+1}^u = \theta_n^u, & \theta_{n+1}^c = \theta_n^c, & \xi_{n+1} = \xi_n, \\ p_{n+1}^u = p_n^u - h \nabla_{\theta^u} L_X(\theta_n), \\ \bar{p}_{n+1,i}^c = p_{n,i}^c - h \left(1 - \frac{1}{r_i^2} |\theta_{n,i}^c|^2\right) \partial_{\theta_i^c} L_X(\theta_n), \\ \bar{p}_{n+1,i}^\xi = p_{n,i}^\xi + h \frac{1}{r_i^2} \theta_{n,i}^c \xi_{n,i} \partial_{\theta_i^c} L_X(\theta_n), \end{cases} \\
 \text{(O step, circle)} \quad & \begin{cases} \theta_{n+1}^u = \theta_n^u, & \theta_{n+1}^c = \theta_n^c, & \xi_{n+1} = \xi_n, \\ p_{n+1}^u = e^{-\gamma h} p_n^u + \sqrt{\beta^{-1}(1 - e^{-2\gamma h})} R^u, \\ \bar{p}_{n+1}^c = e^{-\gamma h} p_n^c + \sqrt{\beta^{-1}(1 - e^{-2\gamma h})} R^c, \\ \bar{p}_{n+1}^\xi = e^{-\gamma h} p_n^\xi + \sqrt{\beta^{-1}(1 - e^{-2\gamma h})} R^\xi, \\ p_{n+1,i}^c = \left(1 - \frac{1}{r_i^2} |\theta_{n,i}^c|^2\right) \bar{p}_{n+1,i}^c - \frac{1}{r_i^2} \theta_{n,i}^c \xi_{n,i} \bar{p}_{n+1,i}^\xi, \\ p_{n+1,i}^\xi = -\frac{1}{r_i^2} \theta_{n,i}^c \xi_{n,i} \bar{p}_{n+1,i}^c + \left(1 - \frac{1}{r_i^2} |\xi_{n,i}|^2\right) \bar{p}_{n+1,i}^\xi, \end{cases}
 \end{aligned}$$

where R^u , R^c , and R^ξ are vectors of independent standard normal random variables.

B.5. Orthogonality constraint, overdamped Langevin dynamics (o-CoLA-od)

We present here a particular discretization of the constrained overdamped Langevin dynamics (4) for the orthogonality constraint (3).

For notational convenience, we present the updates for the weight matrix W^ℓ of a given layer ℓ . The updates for the biases are standard Euler–Maruyama steps such as given for θ^u in (42).

Referring to (3), we denote

$$\begin{aligned}
 Q &= W^\ell, & r &= n^\ell, & s &= n^{\ell-1} & \text{if } n^{\ell-1} \leq n^\ell, \\
 Q &= (W^\ell)^T, & r &= n^{\ell-1}, & s &= n^\ell & \text{otherwise.}
 \end{aligned} \tag{44}$$

so that $Q \in \mathbb{R}^{r \times s}$. With this notation, the constraint (3) is $g(Q) = 0$ where

$$g : \mathbb{R}^{r \times s} \rightarrow \mathbb{R}^{s \times s}, \quad g(Q) = Q^T Q - I_s. \tag{45}$$

Recall that due to symmetry, the matrix equality $g(Q) = 0_s$ corresponds to $s(s+1)/2$ constraints. We compute the partial derivative

$$\partial_{Q_{kl}} g_{ij}(Q) = \delta_{li} Q_{kj} + \delta_{lj} Q_{ki} \quad 1 \leq i, j, k \leq s, \quad 1 \leq l \leq r. \tag{46}$$

In particular, if Λ is an $s \times s$ symmetric matrix, we verify that

$$\sum_{i,j=1}^s \partial_{Q_{kl}} g_{ij}(Q) \Lambda_{ij} = 2(Q\Lambda)_{kl}.$$

We thus obtain the natural matrix form of the constrained dynamics (4): $Q_t : (0, \infty) \rightarrow \mathbb{R}^{r \times s}$ solves

$$\begin{aligned} dQ_t &= -\nabla_Q V(Q_t) dt + \sqrt{2\beta^{-1}} dW_t - Q_t d\Lambda_t, \\ g(Q_t) &= 0, \end{aligned} \quad (47)$$

where $(\nabla_Q V)_{ij} = \partial_{Q_{ij}} V = \partial_{W_{ij}^\ell} L_X$ (or $\partial_{W_{ji}^\ell} L_X$) and W_t is a Wiener process in $\mathbb{R}^{r \times s}$. Furthermore the process Λ_t has values in the $s \times s$ symmetric matrices and is the Lagrange multiplier corresponding to the $s(s+1)/2$ constraints.

Applying discretization scheme (36) to (47), we obtain the iteration step $Q_n \in \Sigma \mapsto Q_{n+1} \in \Sigma$ given by

$$\begin{aligned} \bar{Q}_{n+1} &= Q_n - h\nabla_Q V(Q) + \sqrt{2\beta^{-1}h}R_n, & Q_{n+1} &= \bar{Q}_{n+1} - Q_n\Lambda_n, \\ \text{where } \Lambda_n &\text{ is a symmetric } s \times s \text{ matrix s.t. } g(Q_{n+1}) = 0, \end{aligned} \quad (48)$$

and $R_n \in \mathbb{R}^{r \times s}$ is a matrix of independent standard normal random variables.

Note that the projection step in (48) requires to solve a non-linear system. Following a similar technique as described in [22, Chap. 8], we derive a quasi-Newton scheme for that task. Using the fact that Q_n satisfies the constraint we verify that

$$\bar{Q}_{n+1}^T Q_n = I_s - h\nabla_Q V(Q_n)^T Q_n + \sqrt{2\beta^{-1}h}R_n^T Q_n.$$

The constraint $g(Q_{n+1}) = 0$ thus reads

$$0 = (\bar{Q}_{n+1} - Q_n\Lambda_n)^T (\bar{Q}_{n+1} - Q_n\Lambda_n) - I_s = (\bar{Q}_{n+1}^T \bar{Q}_{n+1} - I_s) - 2\Lambda_n + \mathcal{O}(\sqrt{h}), \quad (49)$$

where $\mathcal{O}(\sqrt{h})$ denotes a matrix whose 2-norm has order \sqrt{h} . Solving for Λ_n , we find

$$\Lambda_n = \frac{1}{2}(\bar{Q}_{n+1}^T \bar{Q}_{n+1} - I_s) + \mathcal{O}(\sqrt{h}).$$

Neglecting the terms of order \sqrt{h} and higher, we obtain the following quasi-Newton scheme: setting $Q^{(0)} = \bar{Q}_{n+1}$, repeat the iteration

$$Q^{(k+1)} = Q^{(k)} - Q_n \Lambda^{(k)}, \quad \text{where } \Lambda^{(k)} = \frac{1}{2}((Q^{(k)})^T Q^{(k)} - I_s), \quad (50)$$

until the process reaches convergence and set $Q_{n+1} = Q^{(k+1)}$. To assess whether convergence has been reached, a tolerance on the 2-norm of $\Lambda^{(k)}$ can be assigned: $\|\Lambda^{(k)}\| \leq \text{TOL}$. However in practice, to ensure that the process ends and to avoid undesirable overhead we typically prefer to either combine this stopping criterion with a limit for the number K of iterations, or use a fixed number of iterations K . Note that estimate (49) ensures that a small number of iterations K is sufficient for the constraint to be satisfied up to a small error.

The initialization for the constrained weights is performed following [39], which is an built-in option in PyTorch. Other parameters are initialized using the standard PyTorch initialization [13, 33] unless otherwise indicated. Constraints are applied layer-wise, where for convolutional layers with weight tensors of the size $n_l \times n_{l-1} \times n_h \times n_w$ (where n_h and n_w are the height and width of the kernel) the weight matrices are reshaped as $n_l \times n_{l-1}n_hn_w$. For CNNs these reshaped matrices are typically rectangular. If they are thin, but long (i.e., $n_l > n_{l-1}n_hn_w$) we apply the constraint $W^T W = I$, but if they have more columns than rows we apply the constraint $W W^T = I$.

B.6. Orthogonality constraint, underdamped Langevin (o-CoLA-ud)

To discretize the underdamped Langevin constrained dynamics, we need the orthogonal projection Π onto the cotangent space $T_Q^*\Sigma$. As the constraint (45) is given in a matrix form, using the formula (18) is not very convenient so we will rather derive Π from its projection property.

Using (46), we find that for $1 \leq i, j \leq s$

$$0 = \sum_{k=1}^s \sum_{l=1}^r \partial_{Q_{kl}} g_{ij}(Q) P_{kl} = (P^T Q + Q^T P)_{ij},$$

which leads to the following convenient expression for the cotangent space

$$T_Q^*\Sigma = \{P \in \mathbb{R}^{r \times s} \mid P^T Q + Q^T P = 0_s\}.$$

Now, given $\bar{P} \in \mathbb{R}^{r \times s}$ we want to find a symmetric $s \times s$ matrix Λ such that $P = \bar{P} - Q\Lambda$ belongs to $T_Q^*\Sigma$, i.e.,

$$0_s = P^T Q - Q^T P = \bar{P}^T Q + Q^T \bar{P} - \Lambda Q^T Q - Q^T Q \Lambda.$$

This equation is easily solved for $Q \in \Sigma$ and we find $\Lambda = \frac{1}{2}(\bar{P}^T Q + Q^T \bar{P})$. We obtain the following expression for the projection onto the cotangent space:

$$\Pi_Q : \mathbb{R}^{r \times s} \rightarrow \mathbb{R}^{r \times s}, \quad \bar{P} \mapsto \Pi_Q \bar{P} = \bar{P} - \frac{1}{2}Q(\bar{P}^T Q + Q^T \bar{P}).$$

We then verify that Π_Q is indeed a projection onto the cotangent space $T_Q^*\Sigma$ (i.e., $\Pi_Q \bar{P} \in T_Q^*\Sigma \forall \bar{P} \in \mathbb{R}^{r \times s}$ and $\Pi_Q^2 = \Pi_Q$) and that this projection is orthogonal with respect to the Frobenius inner product on $\mathbb{R}^{r \times s}$ (i.e., $\langle \bar{P} - \Pi_Q \bar{P}, P \rangle = 0$, where $\langle A, B \rangle = \text{tr}(A^T B)$).

A component. For the orthogonal constraint, the A component in (7) can only be solved approximately. A simple yet efficient discretization of A is the RATTLE scheme (see e.g. [22, Chap. 8]):

$$\begin{aligned} Q_{n+1} &= Q_n + hP_{n+1/2}, \\ P_{n+1/2} &= P_n - Q_n \Lambda_{n+1/2} \quad \text{where } \Lambda_{n+1/2} \text{ is s.t. } Q_{n+1}^T Q_{n+1} = I_s, \\ P_{n+1} &= P_{n+1/2} - Q_{n+1} \Lambda_{n+1} \quad \text{where } \Lambda_{n+1} \text{ is s.t. } Q_{n+1}^T P_{n+1} + P_{n+1}^T Q_{n+1} = 0_s. \end{aligned} \tag{51}$$

Denoting $\bar{\Lambda}_{n+1/2} = h\Lambda_{n+1/2}$, $\bar{P}_{n+1} = P_{n+1/2}$ and using the projection operator Π_Q , (51) can be rewritten as

$$\begin{aligned} \bar{Q}_{n+1} &= Q_n + hP_n, \\ Q_{n+1} &= \bar{Q}_{n+1} - Q_n \bar{\Lambda}_{n+1/2} \quad \text{where } \bar{\Lambda}_{n+1/2} \text{ is s.t. } Q_{n+1}^T Q_{n+1} = I_s \quad (\text{use (50)}), \\ \bar{P}_{n+1} &= P_n - \frac{1}{h} Q_n \bar{\Lambda}_{n+1/2}, \quad P_{n+1} = \Pi_{Q_{n+1}} \bar{P}_{n+1}. \end{aligned} \tag{52}$$

As in the overdamped case, we may now use the quasi-Newton scheme (50) for the projection step (to approximate $\bar{\Lambda}_{n+1/2}$). Using K iterations of the quasi-Newton scheme (50) (i.e., $Q_{n+1} = Q^{(K)}$), we verify that $-Q_n \bar{\Lambda}_{n+1/2}$ satisfies

$$-Q_n \bar{\Lambda}_{n+1/2} = \sum_{k=0}^{K-1} Q_n \Lambda^{(k)} = \sum_{k=0}^{K-1} Q^{(k+1)} - Q^{(k)} = Q^{(K)} - Q^{(0)} = Q_{n+1} - \bar{Q}_{n+1},$$

so that $\bar{P}_{n+1} = P_n + \frac{1}{h}(Q_{n+1} - \bar{Q}_{n+1})$.

We obtain the following full discretization of the underdamped Langevin dynamics with orthogonality constraint. The initialization for the constrained weights is performed following [39]. Corresponding momenta are initialized as the initial gradients (equivalently to standard PyTorch initialization) and subsequently projected using $P_0 = \bar{P}_0 - \frac{1}{2}Q_0(\bar{P}_0^T Q_0 + Q_0^T \bar{P}_0)$. The A,B,O steps are then given as:

$$\begin{aligned}
 \text{(A, OG)} \quad & \left\{ \begin{array}{l} \bar{Q}_{n+1} = Q_n + hP_n, \quad Q^{(0)} = \bar{Q}_{n+1}, \\ \text{for } k = 0 : K - 1 \quad Q^{(k+1)} = Q^{(k)} - Q_n \Lambda^{(k)}, \quad \text{where } \Lambda^{(k)} = \frac{1}{2} \left((Q^{(k)})^T Q^{(k)} - I_s \right), \\ Q_{n+1} = Q^{(K)}, \\ \bar{P}_{n+1} = P_n + \frac{1}{h}(Q_{n+1} - \bar{Q}_{n+1}), \\ P_{n+1} = \Pi_{Q_{n+1}} \bar{P}_{n+1} = \bar{P}_{n+1} - \frac{1}{2}Q_{n+1} \left(\bar{P}_{n+1}^T Q_{n+1} + (Q_{n+1})^T \bar{P}_{n+1} \right). \end{array} \right. \\
 \text{(B, OG)} \quad & \left\{ \begin{array}{l} Q_{n+1} = Q_n, \\ \bar{P}_{n+1} = P_n - h \nabla_Q V(Q_n), \\ P_{n+1} = \Pi_{Q_n} P_{n+1} = \bar{P}_{n+1} - \frac{1}{2}Q_n \left(\bar{P}_{n+1}^T Q_n + (Q_n)^T \bar{P}_{n+1} \right), \end{array} \right. \\
 \text{(O, OG)} \quad & \left\{ \begin{array}{l} Q_{n+1} = Q_n, \\ \bar{P}_{n+1} = e^{-\gamma h} P_n + \sqrt{\beta^{-1}(1 - e^{-2\gamma h})} R_n, \\ P_{n+1} = \Pi_{Q_n} \bar{P}_{n+1} = \bar{P}_{n+1} - \frac{1}{2}Q_n \left(\bar{P}_{n+1}^T Q_n + (Q_n)^T \bar{P}_{n+1} \right), \end{array} \right.
 \end{aligned}$$

where R_n is a matrix of independent standard normal random variables.

Appendix C. Additional Numerical Details and Results

We compare our constrained methods with PyTorch’s [33] SGD with momentum optimiser. Unless otherwise indicated, we use for SGD $h = 0.1$ and $mom = 0$ (to compare with our constrained overdamped Langevin method) or $mom = 0.9$ (to compare with our constrained underdamped Langevin method). We use standard PyTorch initialization for all unconstrained parameters [13, 33].

C.1. Orthogonality Constraints

For our experiments on the spiral data set (see Fig. C5) we use multi-layer perceptrons with ReLU activation and binary cross entropy loss. In our experiments we vary the number of 100-node hidden layers of the multi-layer perceptrons. To compare the performance of our o-CoLA-od constrained method with standard SGD we set the temperature $\tau = 0$ to generate Figure 1. We used a small temperature perturbation $\tau = 1e-6$ to generate Figure 2. The size of the temperature parameter was chosen to approximately match observed fluctuations in the loss function. A more precise parameterization is left for a subsequent work.

A plot of the planar spiral data set binary classification problem as used to produce Figure 1 and Figure 2 is provided in Figure C5. The first class of the data set is generated using

$$\begin{aligned} x &= 2\sqrt{t} \cos(8\sqrt{t}\pi) + 0.02\mathcal{N}(0, 1), \\ y &= 2\sqrt{t} \sin(8\sqrt{t}\pi) + 0.02\mathcal{N}(0, 1), \end{aligned} \quad (53)$$

where t is drawn repeatedly from the uniform distribution $\mathcal{U}(0, 1)$ to generate data points. The other class of this dataset is obtained by shifting the argument of the trigonometric functions by π . For our experiments we used 500 training data points, 1000 test data points and 5% subsampling.

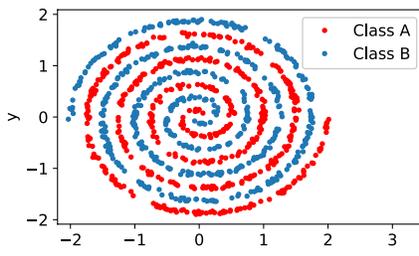


Figure C5: 4-turn spiral data set.

We also applied our orthogonality-constrained methods to the ResNet-34 architecture on CIFAR-10 image classification data [19]. In this setting, running SGD with orthogonal initialization worsened the generalization performance of the resulting net and hence the standard PyTorch initialization was used for SGD. We train for 150 epochs and use a batchsize of 128. In Figure C6 we compare the overdamped variant o-CoLA-od (with $\tau = 0$) to its unconstrained counterpart. We observe that constraining orthogonality gives lower test loss throughout training.

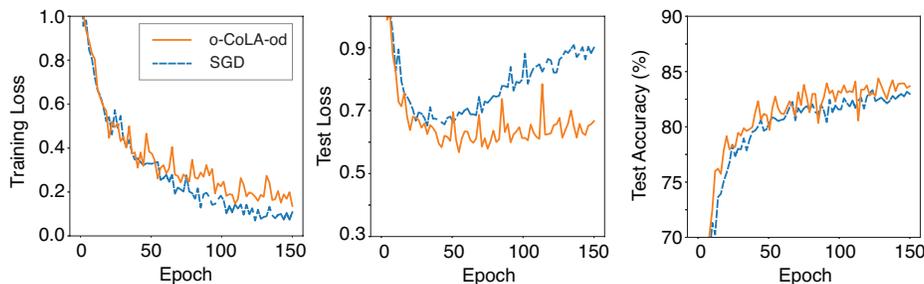


Figure C6: Training loss (left), test loss (middle) and test accuracy (right) of a ResNet-34 architecture trained using SGD vs. o-CoLA-od on CIFAR-10 data, $h = 0.1$ (averaged over 5 runs). The orthogonality constraint provides modestly higher test accuracy and inhibits overfitting.

C.2. Circle constraints

For our Fashion-MNIST [48] example we reduce the number of training data samples to 10,000 and we increase the number of test data samples to 60,000. We use a 1000-node single hidden layer

perceptron (SHLP) with ReLU activation, cross entropy loss and batchsize 128. Our main result with our circle constrained approach is presented in Figure 4. In this section we present extensive hyperparameter tests for the test accuracy and test loss obtained after 400 epochs (averaged over 5 runs) using SGD-m with and without weight decay (WD).

| SGD with mom | | no WD | | with WD | |
|--------------|------------|-----------|-----------|-----------|-----------|
| | | Test Acc. | Test Loss | Test acc. | Test Loss |
| h = 0.2 | mom = 0.8 | 87.18% | 1.06 | 84.05% | 0.696 |
| | mom = 0.7 | 87.38% | 0.890 | 87.0% | 0.547 |
| h = 0.1 | mom = 0.9 | 86.97% | 1.133 | 85.35% | 0.634 |
| | mom = 0.8 | 87.39% | 0.824 | 87.47% | 0.531 |
| | mom = 0.7 | 87.39% | 0.750 | 87.25% | 0.517 |
| h = 0.05 | mom = 0.95 | 86.67% | 1.226 | 85.63% | 0.623 |
| | mom = 0.9 | 87.33% | 0.837 | 86.24% | 0.569 |
| | mom = 0.8 | 87.27% | 0.719 | 87.33% | 0.511 |

The results presented in the two right-hand columns are all obtained with weight decay set to $1e-4$. We found this value to give the best results for SGD-m during a hyperparameter search.

In comparison our circle constrained net reaches test accuracy **87.61%**, with test loss **0.386** without using weight decay (see Figure 4). Hence it outperforms standard SGD with momentum both with and without weight decay.

C.3. NLP

We evaluate the performance of a small transformer model [43] on the Penn Treebank data set [29] and Wikitext-2 data set [30]. The transformer has 2 encoder layers. Each encoder layer consists of self-attention with 2 heads and a feedforward network with 200 nodes followed by layer norms. We use batch size 1024 for the Penn Treebank data set and batchsize 128 for the Wikitext-2 dataset. We present the lowest validation loss obtained in 200 epochs by SGD-m and our circle constrained method c-CoLA-ud in Table 1 of the main paper.