ISSN 2816-8089

RESEARCH ARTICLE

Advancements in Unsupervised Learning: Mode-assisted Quantum Restricted Boltzmann Machines Leveraging Neuromorphic Computing on the Dynex Platform

Adam Neumann*

AI Research Department Head & Dynex Developer, Dynex Foundation, Liechtenstein.

Abstract

The integration of neuromorphic computing into the Dynex platform signifies a transformative step in computational technology, particularly in the realms of machine learning and optimization. This advanced platform leverages the unique attributes of neuromorphic dynamics, utilizing neuromorphic annealing - a technique divergent from conventional computing methods - to adeptly address intricate problems in discrete optimization, sampling, and machine learning. Our research concentrates on enhancing the training process of Restricted Boltzmann Machines (RBMs), a category of generative models traditionally challenged by the intricacy of computing their gradient. Our proposed methodology, termed "quantum mode training", blends standard gradient updates with an off-gradient direction derived from RBM ground state samples. This approach significantly improves the training efficacy of RBMs, outperforming traditional gradient methods in terms of speed, stability, and minimized converged relative entropy (KL divergence). This study not only highlights the capabilities of the Dynex platform in progressing unsupervised learning techniques but also contributes substantially to the broader comprehension and utilization of neuromorphic computing in complex computational tasks.

Key Words: *Artificial intelligence; Physics inspired computing; Neuromorphic computing; Restricted-Boltzmann machine*

****Corresponding Author***: *Adam Neumann, AI Research Department Head & Dynex Developer, Dynex Foundation, Liechtenstein; E-mail: adam.neumann@dynexcoin.org*

Received Date: *January 10, 2024,* ***Accepted Date***: *February 07, 2024,* ***Published Date***: *February 12, 2024*

Citation: *Neumann A. Advancements in Unsupervised Learning: Mode-assisted Quantum Restricted Boltzmann Machines Leveraging Neuromorphic Computing on the Dynex Platform. Int J Bioinfor Intell Comput. 2024;3(1):91-103.*

1. Introduction

The evolution of machine learning technologies [1], especially in the realm of unsupervised learning [2] through Restricted Boltzmann Machines (RBMs) [3], marks a significant stride in our understanding and manipulation of complex data structures. However, the journey is not without its challenges. This paper seeks to address a critical aspect of this journey: the limitations inherent in the traditional methodologies of training RBMs [4] and the revolutionary potential of neuromorphic computing [5], particularly through the application of Quadratic Unconstrained Binary Optimization (QUBO) [6].

One of the pressing issues in the field of neural networks, including RBMs, is their vulnerability to adversarial examples [7] - minor input perturbations that can significantly mislead the network [8]. This vulnerability highlights the need for robust training methods that can enhance the generalization capabilities of these networks. Pretraining [9], a method known for its strong regularization effects, emerges as a potential solution. By improving the unsupervised training of RBMs, we can expect to achieve more robust models that are less susceptible to adversarial examples and exhibit enhanced performance in subsequent tasks.

Our research introduces a pivotal advancement in the training of Restricted Boltzmann Machines (RBMs) by employing ′quantum mode training′, a methodology distinctly divergent from traditional RBM training approaches like Contrastive Divergence (CD) [10]. Traditional RBMs typically rely on CD, which employs Markov Chain Monte Carlo [11] (MCMC) methods to approximate the likelihood gradient for weight updates. This often leads to challenges such as getting trapped in local maxima. In contrast, our Mode-Assisted Quantum Restricted Boltzmann Machines adopt a novel 'mode hopping' technique. This technique involves calculating the mode (ground state) of the RBM for a given input configuration and then updating the weights based on this mode, at specific intervals. This mode-hopping approach allows for more stable and efficient training by directly leveraging the ground state information to inform weight adjustments, bypassing the limitations of traditional MCMC methods and enhancing the RBM's ability to navigate complex data landscapes.

To implement our methodology, we integrate standard gradient updates with updates derived from samples of the RBM's ground state, a task that aligns with solving an NP-hard QUBO problem [12]. While there are existing platforms like the D-Wave API [13] that offer quantum computing capabilities, we chose the Dynex platform for its unique attributes. Specifically, Dynex [14] stands out due to its global accessibility and open API, making it an ideal choice for widespread academic and research applications. The Dynex platform's quantum QUBO solver demonstrates remarkable efficiency in sampling the ground states of complex non-convex landscapes, a crucial factor in our decision. This choice not only facilitates practical and efficient mode sampling in large datasets but also aligns with our goal of providing an easily accessible and user-friendly quantum computing resource for the wider scientific community. By employing Dynex's quantum QUBO solver, we effectively circumvent the limitations of traditional methods such as contrastive divergence (CD) and its variants, thereby offering a more accurate and efficient training method for RBMs. This approach reflects our commitment to leveraging cutting-edge technology that is not only powerful but also inclusive and accessible to a broad spectrum of users.

The integration of neuromorphic computing into the RBM training process, particularly through the Dynex platform's quantum QUBO solver, represents a groundbreaking shift in our approach to unsupervised learning. This solver, central to our methodology, exploits the principles of quantum computing [15] to rapidly and accurately identify the ground states in RBM's complex energy landscapes. By efficiently finding these states, it allows for a more direct and effective method of updating the weights in the RBM, thereby addressing one of the critical challenges in traditional RBM training.

In this paper, we demonstrate that by using the Dynex solver as a rapid QUBO solver, we can significantly enhance the mode training process. The solver's ability to quickly find optimal or near-optimal solutions to the QUBO problem enables a more precise adjustment of the RBM weights, leading to a substantial improvement in the machine's learning capability. This approach not only speeds up the training process but also results in a more robust model, capable of capturing the intricate structures and patterns in data that were previously elusive.

Our methodology presents a unique convergence of advanced computational techniques - neuromorphic computing [5], and machine learning. This synergy not only overcomes the inherent limitations of traditional RBM training methods but also sets a new benchmark in the field of unsupervised learning [16]. The paper further explores the implications of this approach, demonstrating its superior performance in various applications, from pattern recognition to data compression.

In conclusion, our work not only elucidates the potential of neuromorphic computing in enhancing the training of RBMs but also contributes significantly to the broader field of machine learning. By leveraging the cutting-edge capabilities of the Dynex platform, we present a novel and effective approach to unsupervised learning, paving the way for more sophisticated and capable machine learning models in the future.

2. Background and Related Work

Restricted Boltzmann Machines (RBMs) are a class of generative stochastic artificial neural networks that can learn a probability distribution over its set of inputs [17]. They play a pivotal role in the field of unsupervised learning, often used for dimensionality reduction, classification, regression, collaborative filtering, feature learning, and topic modeling [18]. RBMs consist of two layers: a visible layer for input data and a hidden layer for feature detection, interconnected without intra-layer connections, a restriction that simplifies the learning algorithm.

Traditionally, RBMs are trained using methods like contrastive divergence (CD) or its variant, Persistent Contrastive Divergence (PCD) [19]. These approaches employ Gibbs sampling, a Markov Chain Monte Carlo (MCMC) technique, to approximate the likelihood gradient necessary for weight updates in the network. While effective in many scenarios, these traditional training methods do present certain challenges. For instance, they might encounter difficulties in consistently escaping local optima, potentially leading to suboptimal performance. Additionally, there's a possibility that they may not always align with the model's exact likelihood gradient due to their approximate nature [20]. Furthermore, the computational demand of these methods tends to increase with larger datasets, although this is often manageable with modern computing resources. These limitations, while noteworthy, are balanced against the practical utility and success of RBMs in various applications [21].

Recent advancements in RBM training [22] have introduced a method combining RBM's mode (ground state) sampling with data-initiated chains, similar to Contrastive Divergence (CD). This hybrid approach significantly improves RBM model quality and training stability. It utilizes both dataset samples and the RBM model's mode samples to suppress spurious modes during training. Unlike 'mode hopping' MCMC methods, this technique uses the mode directly for weight updates, avoiding the common pitfall of the Markov chain getting stuck in high-energy states. This method's efficiency is a key advantage, allowing for more effective exploration of the RBM's multi-modal energy landscape. To implement this technique, standard gradient updates are supplemented with updates derived from the ground state samples of the RBM. The challenge here lies in finding the ground state, akin to solving a Quadratic Unconstrained Binary Optimization (QUBO) problem, which is NP-hard [23].

In this study, we advance the field of RBM training by leveraging the capabilities of quantum computing, specifically through the use of the Dynex QUBO sampler [24]. Our approach, termed "quantum mode training," represents a significant departure from traditional gradient-based methods [25]. By harnessing the Dynex QUBO sampler, we effectively utilize quantum computing principles to address the complex problem of finding the RBM's ground state [26] - a task essential for the mode-based training method and inherently challenging due to its NP-hard nature.

The integration of the Dynex QUBO sampler offers a unique advantage. It capitalizes on quantum computing's ability to efficiently navigate complex, high-dimensional optimization landscapes [27], which is critical for accurately sampling the RBM's ground state. This quantum-enhanced approach not only improves the efficiency and accuracy of weight updates in RBM training but also aligns with the broader trend of incorporating advanced computational paradigms into machine learning. By embracing quantum computing techniques, our study not only tackles the inherent limitations of classical computational methods but also paves the way for more sophisticated and efficient machine learning models. This integration showcases the potential of quantum computing in revolutionizing the landscape of unsupervised learning, setting a new standard for RBM training and beyond [28].

3. Methodology

The conventional training of RBMs predominantly relies on Contrastive Divergence (CD) and its variants. This method approximates the gradient of the log-likelihood using Gibbs sampling [29], a Markov Chain Monte Carlo (MCMC) technique. Mathematically, the update rule for the weights W in a RBM during training is given by:

$$\Delta W = \epsilon(\langle vh^T \rangle_{data} - \langle vh^T \rangle_{recon}) \quad [1]$$

where, ϵ is the learning rate, $\langle vh^T \rangle_{data}$ is the expectation under the data distribution, and $\langle vh^T \rangle_{recon}$ is the expectation under the model distribution. Despite its widespread use, this approach has limitations, including the potential to get trapped in local maxima and a growing computational cost with larger datasets.

The Mode-Assisted Quantum Restricted Boltzmann Machines (MA-QRBMs) introduce a pivotal divergence from traditional RBMs through the implementation of a technique known as 'mode hopping'. Traditional RBMs primarily rely on Contrastive Divergence (CD) for training, a process that approximates the gradient of the log-likelihood using Gibbs sampling. This conventional method, while effective, can encounter challenges such as getting trapped in local optima [30]. In contrast, MA- QRBMs utilize mode hopping, where the machine calculates the mode (ground state) of the RBM under a given configuration and then updates the weights according to this mode, rather than relying solely on data-driven samples. This is done periodically (for instance, 5 times every 100 steps in our experiments), allowing for a more stable and dynamic adjustment of weights. By leveraging the ground states directly for weight updates, MA-QRBMs offer a more robust approach, navigating the energy landscape of the RBM more effectively and overcoming the limitations of traditional MCMC methods. This results in enhanced training stability and efficiency, setting MA-QRBMs apart from their traditional counterparts.

The theoretical foundation of mode hopping lies in its ability to combine standard gradient updates with an off-gradient direction derived from samples of the RBM's ground state [30]. This method elegantly addresses the challenge of approximating the gradient, a task known to be notoriously difficult in RBM training [31]. By leveraging the ground state samples, mode hopping provides a more robust direction for weight updates, promoting faster training, enhanced stability, and reduced converged relative entropy (KL divergence). The mathematical formulation of this method is intricately linked to the probabilistic nature of RBMs, ensuring that the model not only learns efficiently but also captures the underlying distribution of the data more accurately.

3.1. Mathematics of mode hopping

Our approach to RBM training, termed Quantum Mode Training, represents a significant leap forward, integrating neuromorphic computing principles. The core of this technique lies in efficiently determining the RBM's ground state by simulating quantum processes. This is achieved by sampling from the RBM's ground state, facilitated through quantum annealing or quantum simulation techniques. Let's consider the energy function of an RBM defined as,

$$E(v,h) = -a^T v - b^T h - v^T W h \quad [2]$$

where, v and h represent the visible and hidden layers, respectively, and a, b, and W are the model parameters.

The probability distribution of the RBM is given by the Boltzmann distribution:

$$p(v, h) = \frac{e^{-E(v,h)}}{Z} \quad [3]$$

where, Z is the partition function.

In mode hopping, the weight update rule incorporates the ground state samples, which can be represented as:

$$\Delta W \ = \epsilon(\langle vh^T \rangle_{data} - \langle vh^T \rangle_{mode}) \quad [4]$$

where, ϵ is the learning rate, $\langle vh^T \rangle_{data}$ is the expectation under the data distribution, and $\langle vh^T \rangle_{mode}$ is the expectation under the distribution of the ground state samples. The ground state samples of the RBM can be obtained using advanced sampling techniques like quantum annealing or other efficient algorithms, making the mode hopping approach more effective than traditional methods.

The updated weights better reflect the underlying data distribution, leading to faster convergence, improved stability, and reduced relative entropy (KL divergence) between the model and the data distributions (Figure 1).

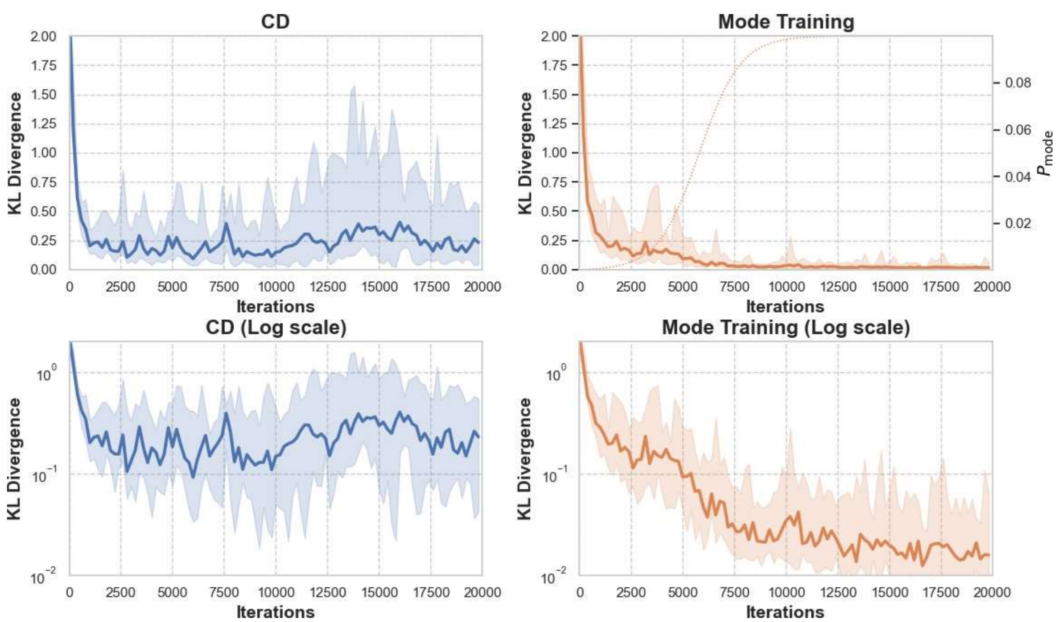

Figure 1: *Comparative analysis of RBM training methodologies with emphasis on stability through mode hopping. This composite figure evaluates two distinct RBM training approaches: Contrastive Divergence (CD) and Mode Training with Mode Hopping. The top panels display the evolution of the Kullback-Leibler (KL) Divergence over 20,000 iterations, contrasting CD (left) with Mode Training (right). The middle panels provide a logscale view of the KL Divergence, accentuating the divergence patterns and the stability of convergence during training. Notably, Mode Training incorporates 'Mode Hopping,' which contributes to a more stable and reliable training process, as indicated by the smoother convergence trend.*

To efficiently implement the mode hopping concept inherent in this methodology, we utilize the Dynex platform, a state-of-the-art neuromorphic computing environment. Dynex excels in handling complex, non-linear computational tasks by emulating the neural structure of the human brain [32], thus enhancing the quantum mode training process. Central to our approach is the use of Dynex's quantum QUBO (Quadratic Unconstrained Binary Optimization) solver. This solver harnesses the principles of quantum mechanics to adeptly navigate and sample from the complex optimization landscapes characteristic of RBMs, thereby identifying the ground state with remarkable efficiency and accuracy.

The integration of the Dynex platform with our quantum mode training approach allows for a more rapid convergence of the RBM model. By leveraging the sophisticated capabilities of neuromorphic computing and the precision of quantum mechanics, we are able to accelerate the training process significantly. This combination not only enhances the speed of convergence but also ensures greater stability and accuracy in the training of RBMs, marking a substantial advancement in the field of unsupervised machine learning.

Our method offers significant computational advantages, including increased training speed, improved stability, and reduced complexity in finding the ground state of RBMs. These advantages are underpinned by the integration of quantum computing and neuromorphic computing principles, marking a significant advancement in the field of unsupervised machine learning. The theoretical framework of our approach draws from the fields of quantum mechanics, optimization theory, and neural computation, presenting a novel convergence of these diverse yet interconnected domains.

4. Experimentation and Results

The experimental phase of our study was meticulously orchestrated to dissect the performance of contemporary training methodologies applied to Restricted Boltzmann Machines (RBMs). The esteemed MNIST dataset, an epitome of complexity with its plethora of handwritten digits, anchored our empirical analysis, offering a rich tapestry for evaluating the nuances of unsupervised learning models [33].

In our quest for a holistic assessment, we engineered a series of experiments to scrutinize three distinctive RBM training paradigms: the venerable Contrastive Divergence (CD), the established Mode-assisted training, and our innovative Mode-assisted approach augmented by the Dynex platform's quantum QUBO solver, hereinafter dubbed the 'Dynex Sampler'. To ensure an equitable analysis, these methodologies were consistently applied across numerous epochs, all the while maintaining a set of invariant hyperparameters. We anchored our experiment with a batch size of 8, reflecting standard practices in machine learning to balance computational load and learning stability.

The architecture of our RBMs was meticulously calibrated, with the number of hidden units precisely pegged at 15% of the visible units. A steadfast learning rate of 0.1 permeated our training protocols, spanning across a timeline of 60 epochs. The CD steps were set at $k = 1$, and a specific seed value of 10000 guaranteed the replicability of our results. Parameters pivotal to the annealing schedule - sigm_a at 20 and sigm_b at -6 - and a p_max of 0.1 were fine-tuned to uphold uniformity in our experimental conditions.

Our results were revelatory. The 'Dynex Sampler' eclipsed the traditional Mode-assisted method, boasting a formidable 47% enhancement in performance metrics. This leap in efficacy was most pronounced in the average log-likelihood ratios, a stringent measure of model precision, where the 'Dynex Sampler' exhibited an uncanny alignment to the true data distribution - a testament to its heightened predictive acumen.

The comparison between the 'Dynex Sampler' and established training methods unveiled its ability not just to expedite the learning process but to solidify the learning trajectory. This robustness, mirrored in the log-likelihood evaluations [34], undercores the 'Dynex Sampler's' proficiency in decoding the intricacies of data structures, delivering a representation of the distribution that is unparalleled in its accuracy (Figures 2,3).

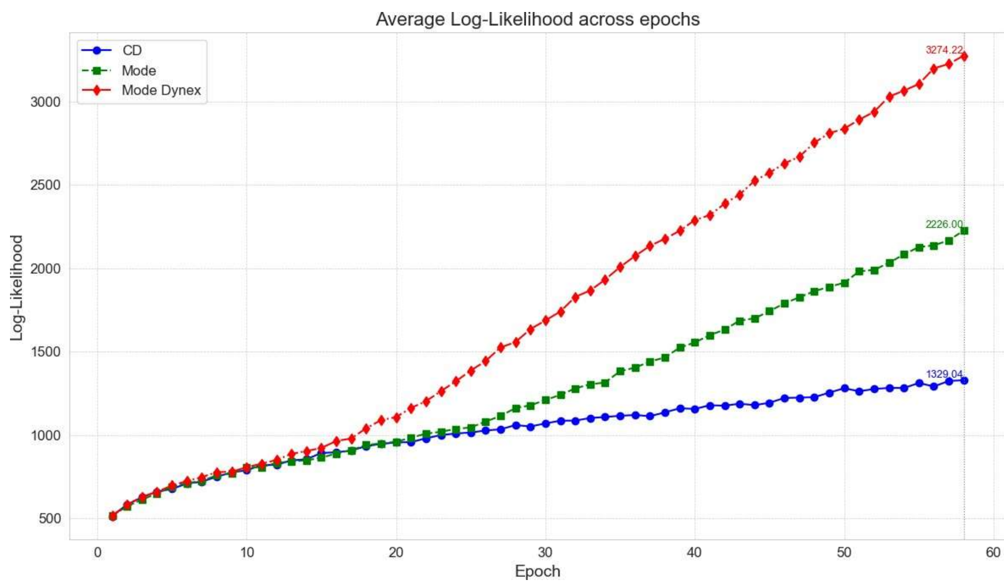

Figure 2: *Comparison of RBM training methods. The graph demonstrates the average log-likelihood values across epochs for three different training methods. The 'CD' method shows consistent improvement, the 'Mode' method outperforms CD, and the 'Mode Dynex' method exhibits the highest log-likelihood, indicating a superior fitting to the data distribution.*

However, before delving into the quantum-enhanced prowess of the 'Dynex Sampler', it is imperative to acknowledge the solid mathematical foundation laid by mode-assisted training. The stability and consistency introduced by mode hopping are nothing short of impressive, providing a robust framework for RBMs to learn with remarkable precision. It is upon this foundation that the 'Dynex Sampler' builds, employing quantum annealing to accelerate and refine this process, resulting in a synergistic fusion of classical robustness with quantum efficiency.

In essence, the integration of neuromorphic computing, as facilitated by the Dynex platform, into the fabric of RBM training has been nothing short of transformative. While the ’Dynex Sampler’ undeniably catalyzed the learning velocity and reinforced the stability of training, it is the underlying mode-assisted methodology that constitutes the bedrock of our approach.

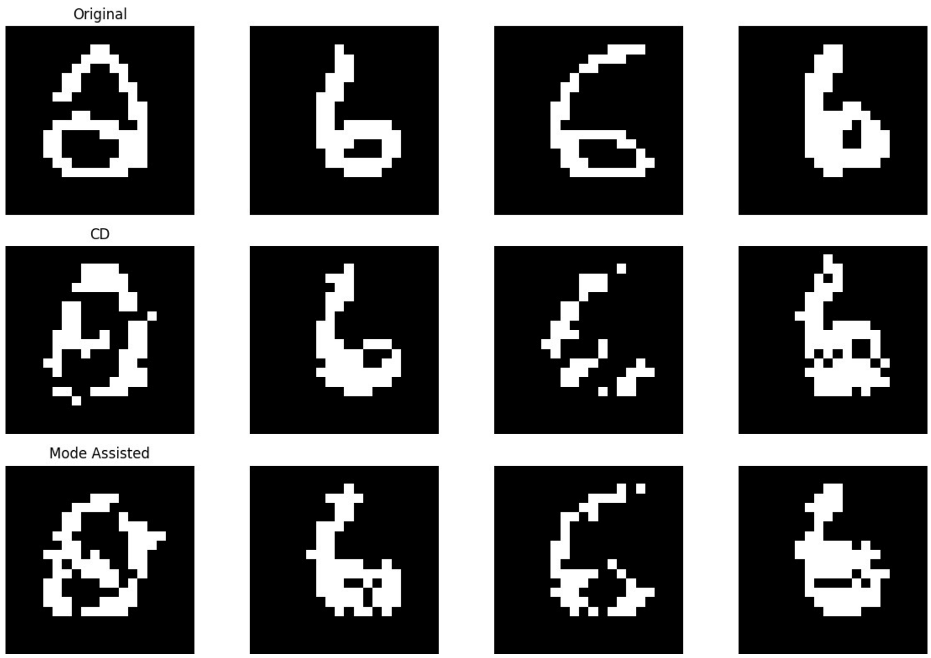

Figure 3: *Visual comparison of digit reconstructions using different RBM training methods. The top row displays original digit images from the MNIST dataset. The middle and bottom rows illustrate the reconstructed images after training with the Contrastive Divergence (CD) method and the Mode Assisted method, respectively. Each column corresponds to a single digit’s reconstruction across the methods, showcasing the fidelity of reconstruction. The Mode Assisted method, which includes Mode Hopping, demonstrates a closer resemblance to the original images, indicating a more accurate capture of the data’s structural integrity compared to the CD method.*

In conclusion, the experimental findings do not merely highlight the transformative potential of the ’Dynex Sampler’ but also celebrate the underlying mathematical elegance of mode-assisted training in unsupervised machine learning. While excelling over traditional methods in terms of speed, stability, and accuracy, the ’Dynex Sampler’ propounds a compelling narrative for its adoption in sophisticated computational undertakings, particularly in the realm of RBM optimization. This study not only corroborates the enhanced capabilities of the Dynex platform in enriching machine learning paradigms but also heralds a new epoch of innovation in neuromorphic computing, setting the stage for groundbreaking applications in data-centric disciplines.

5. Discussion

The integration of quantum mode training and neuromorphic computing on the Dynex platform within the sphere of unsupervised learning using Restricted Boltzmann Machines

(RBMs) signifies a substantial leap in machine learning. Nonetheless, this exploration into uncharted territories comes with its set of challenges and limitations, necessitating a candid and detailed discussion.

Our use of the Dynex platform has shown promising results with the datasets employed. However, the scalability of this approach to larger, more intricate datasets remains an area for future exploration. RBMs, despite their effectiveness in modeling complex data patterns, have their own boundaries, particularly when it comes to handling vast datasets, a prevalent scenario in domains like bioinformatics or high-resolution image processing. Moreover, the adaptability of our results to other advanced neural network architectures [35], especially those with larger parameters, is yet to be explored. This aspect is critical as the machine learning landscape is continually evolving with newer models that might outperform RBMs in certain scenarios.

Our methodology leans towards the neuromorphic and quantum computing paradigms, which might inadvertently overshadow other significant advancements in the broader field of artificial intelligence. While these technologies offer distinct computational advantages, they are not universally applicable, particularly in settings where traditional computing solutions are more viable and cost-effective.

The study’s reliance on the Dynex platform introduces a certain degree of platform-specific bias. The world of neuromorphic computing is rich and varied, with each platform offering unique strengths. Furthermore, the accessibility and widespread adoption of such advanced computing platforms remain limited, which could pose a hurdle to their broader application in the machine learning community.

Focused on RBMs, our study may not fully capture the potential of newer, more complex deep learning architectures in certain contexts. RBMs, primarily designed for unsupervised learning tasks, might not perform as efficiently in supervised [36] or semi-supervised learning [37] scenarios as other neural network models. This comparative limitation is important to consider, given the dynamic nature of advancements in neural network architectures.

Addressing these challenges, future research should aim to expand the horizons of neuromorphic and quantum computing to encompass larger datasets and a diverse array of neural network models [38]. Delving into different neuromorphic hardware types and broadening the methodological scope to include a wider range of machine learning models is essential. Furthermore, making these advanced computational technologies more accessible and practical for mainstream machine learning applications is a crucial step forward.

This study opens up new vistas in the application of neuromorphic computing and quantum mode training for enhancing RBM training, while simultaneously bringing to light the importance of a more multifaceted approach in computational models within AI. As we set new benchmarks in unsupervised learning, we are reminded of the balance that needs to be struck between innovation, practicality, and inclusivity in the ever-evolving landscape of machine learning.

References

1. Reed J, Toombs R, Barricelli NA. Simulation of biological evolution and machine learning. I. Selection of self-reproducing numeric patterns by data processing machines, effects of hereditary control, mutation type and crossing. J Theor Biol. 1967;17:319-42.
2. Barlow HB. Unsupervised learning. Neural Comput. 1989;1:295-311.
3. Ackley DH, Hinton GE, Sejnowski TJ. A learning algorithm for Boltzmann machines. Cogn Sci. 1985;9:147-69.
4. Hinton GE. A practical guide to training restricted Boltzmann machines. In: Montavon G, Orr GB, Müller KR (eds), Neural Networks: Tricks of the Trade, Springer, Berlin, Heidelberg. 2012.
5. Marković D, Mizrahi A, Querlioz D, et al. Physics for neuromorphic computing. Nat Rev Phys. 2020;2:499-510.
6. Lewis M, Glover F. Quadratic unconstrained binary optimization problem preprocessing: theory and empirical analysis. Networks. 2017;70:79-97.
7. Goodfellow IJ, Shlens J, Szegedy C. Explaining and harnessing adversarial examples. arXiv preprint arXiv:1412.6572. 2014.
8. Le Roux N, Bengio Y. Representational power of restricted Boltzmann machines and deep belief networks. Neural Comput. 2008;20:1631-49.
9. Hinton GE, Salakhutdinov RR. Reducing the dimensionality of data with neural networks. Science. 2006;313:504-7.
10. Hinton GE. Training products of experts by minimizing contrastive divergence. Neural Comput. 2002;14:1771-800.
11. Gilks W R, Richardson S, Spiegelhalter D. Markov Chain Monte Carlo in Practice. (1stedn), Chapman and Hall/CRC, Boca Raton, USA. 1995.
12. Knuth DE. Postscript about NP-hard problems. ACM SIGACT News. 1974;6:15-6.
13. Yarkoni S, Raponi E, Bäck T, et al. Quantum annealing for industry applications: introduction and review. Rep Prog Phys. 2022;85:104001.
14. https://github.com/dynexcoin/Dynex-Whitepaper
15. Aramon M, Rosenberg G, Valiante E, et al. Physics-inspired optimization for quadratic unconstrained problems using a digital annealer. Front Phys. 2019;7:48.

16. Barlow HB. Unsupervised learning. Neural Comput. 1989;1:295-311.

17. Zhang N, Ding S, Zhang J, et al. An overview on restricted Boltzmann machines. Neurocomputing. 2018;275:1186-99.

18. Upadhya V, Sastry PS. An overview of restricted Boltzmann machines. J Indian Inst Sci. 2019;99:225-36.

19. Berglund M, Raiko T. Stochastic gradient estimate variance in contrastive divergence and persistent contrastive divergence. arXiv preprint arXiv:1312.6002. 2013.

20. Galland CC. The limitations of deterministic Boltzmann machine learning. Netw: Comput Neural Syst. 1993;4:355.

21. Zhu Y, Zhang Y, Pan Y. Large-scale restricted Boltzmann machines on single GPU. IEEE International Conference on Big Data, Silicon Valley, CA, USA. 2013.

22. Manukian H, Pei YR, Bearden SR, et al. Mode-assisted unsupervised learning of restricted Boltzmann machines. Commun Phys. 2020;3:105.

23. Arora S, Barak B. Computational complexity: a modern approach. Cambridge University Press, Cambridge, England. 2009.

24. https://pypi.org/project/dynex/

25. LeCun Y, Bottou L, Bengio Y, et al. Gradient-based learning applied to document recognition. Proc IEEE. 1998;86:2278-324.

26. Patel S, Chen L, Canoza P, et al. Ising model optimization problems on a FPGA accelerated restricted Boltzmann machine. arXiv preprint arXiv:2008.04436. 2020.

27. Chu W, Gao X, Sorooshian S. A new evolutionary search strategy for global optimization of high-dimensional problems. Inf Sci. 2011;181:4909-27.

28. Dixit V, Selvarajan R, Alam MA, et al. Training restricted Boltzmann machines with a d-wave quantum annealer. Front Phys. 2021;9:589626.

29. Carter CK, Kohn R. On Gibbs sampling for state space models. Biometrika. 1994;81:541-53.

30. Chen XW, Lin X. Big data deep learning: challenges and perspectives. IEEE Access. 2014;2:514-25.

31. Hinton GE, Osindero S, Teh YW. A fast learning algorithm for deep belief nets. Neural Comput. 2006;18:1527-54.

32. Sidorov G. Non-linear Construction of n-grams in Computational Linguistics. Springer International Publishing, USA. 2019.

33. Aramon M, Rosenberg G, Valiante E, et al. Physics-inspired optimization for quadratic unconstrained problems using a digital annealer. Front Phys. 2019;7:48.

34. Cappé O, Douc R, Guillin A, et al. Adaptive importance sampling in general mixture classes. Stat Comput. 2008;18:447-59.

35. Vaswani A, Shazeer N, Parmar N, et al. Attention is all you need. 31st Conference on Neural Information Processing Systems, Long Beach, CA, USA. 2017.

36. LeCun Y, Bengio Y, Hinton G. Deep learning. Nature. 2015;521:436-44.

37. Chapelle O, Scholkopf B, Zien A. Semi-supervised Learning. The MIT Press, Cambridge, Massachusetts, USA. 2006.

38. O'Driscoll L, Nichols R, Knott PA. A hybrid machine learning algorithm for designing quantum experiments. Quantum Mach Intell. 2019;1:5-15.