International Journal of Recent Development in Engineering and Technology
Website: www.ijrdet.com (ISSN 2347-6435(Online) Volume 14, Issue 12, December 2025)

Explaining *Advaitavedānta* through Computational Models of Non-Duality

Pragya[1], Dinesh Mohan Joshi[2], Pratyush Kumar[3]
[1]*BBA University, Lucknow, India*
[2]*IIT Jodhpur, Jodhpur, Rajasthan, India*
[3]*ETH Zurich, Zurich, Switzerland*

***Abstract*— The thinking process of distinct and unparalleled reality outlined by Śaṅkarācārya in his *Advaitavedānta* philosophy encapsulated a yoga of *jīva* under delusion as an *māyā* (illusion) of a self, captured a dual and non dual reality of the *Brahman* (absolute reality), as well as an illusionary and true *Brahman* reality, in a triumph of philosophy as fusion and synthesis. It is an indisputable fact that these positions remain thesis papers for dharma advocates, since they are far too abstract and contrary for the layman. To interpret this, we can think of the study as constructing various analogies between training examples and partial observations, and underlying data distribution and the unchanging substratum of *Brahman*. Model overfitting is explained in terms of *avidyā* or ignorance where noise and superficial correlations are treated as the truth. Gradient descent in the minimization of a loss function inherits the practice of *Vedānta* of inquiry and *adhyāropa–apavāda* (superimposition and negation) where the error gets eliminated step by step to finally reach the realization. Neural networks' hidden layers describe one over another veils of reality, with adversarial examples exposing the deceitful character of *māyā*, in which appearance veils reality.**

***Keywords-* *āropa–Apavāda*, *Advaitavedānta*, Machine Learning, *Māyā*, Non-Duality, Supervised Learning, *Vedānta* Philosophy**

I. Introduction

Advaitavedānta is a major Indian philosophy that was developed by Śaṅkara in the 8th century CE. Its main idea is that the true self, called *Ātman* and the ultimate reality, called *Brahman* are one and the same. The world we see is considered an illusion called *māyā*, which hides this oneness [1-3]. People get confused because of ignorance, which makes them think they are separated from the real truth. To be free from this confusion, one needs knowledge, not actions, and the knowledge comes from understanding that the self is actually the same as *Brahman*. In other words, Liberation (*mokṣa*) is achieved not by action but by knowledge (*jñāna*), specifically through realizing the identity of self and *Brahman*.

Although this philosophy has deep ideas, it is very abstract and not easy for people from outside philosophy and theology to fully understand.

Advaita is one of the most complex and sophisticated branches of Indian Philosophy. Metaphysical framework and *Advaita* are also recognized By Bauer, the emphasis of *Advaita* on oneness beyond the multiplicity of the empirical world is a challenge to ethical and epistemological structures that are based on difference [1, 4]. Matilal also notes that the *Advaitic* method of *adhyāropa*– a montage of superimposition and negation is, *apavāda* elusive to contemporary analytical frameworks though it is precise in resolving illusion.

Recent scholarship has deepened this concern, recognizing *Advaitavedānta* as more than a philosophical system, but rather, a cultural and textual complex that grew with history, politics, and vernacular traditions. Jessica Frazier, for instance, points out that *Vedānta* is a 'many-stranded fabric' of Indian intellectual thought, rather than a single tradition, as it integrates various metaphysical techniques and socio-cultural transmission modalities [7]. In the same manner, Allen notes the *Advaita* vernacular, regional commentaries, and the social milieu of philosophical production especially in the early modern period [8].

Moreover, as long points out, modern *Advaitavedānta* studies (e.g. Arvind Sharma, neo-*Vedāntins*) have begun to uncover the persistent tension between faithfulness to the texts as opposed to transformation of the experience, particularly in regard to the *Advaitic* frameworks that scholars endorse and critique simultaneously [9].

In contrast, machine learning (ML) has become a central paradigm in engineering and applied sciences. ML algorithms construct predictive models from finite and noisy datasets, seeking to approximate an underlying data-generating process [4–6]. In supervised learning, a model is trained by minimizing a loss function, often using iterative optimization schemes such as stochastic gradient descent.

International Journal of Recent Development in Engineering and Technology
Website: www.ijrdet.com (ISSN 2347-6435(Online) Volume 14, Issue 12, December 2025)

Overfitting arises when the model adapts too closely to training noise, while regularization strategies aim to improve generalization. Multi-layer neural networks further capture hierarchical structures in data, though they are also sensitive to adversarial perturbations. These concepts — error, approximation, convergence, and robustness — form the working vocabulary of today's engineers and data scientists.

Although *Advaitavedānta* offers profound insights into the nature of reality, its terminology and methods are far removed from the computational reasoning familiar to scientists and engineers. The absence of shared language limits cross-disciplinary dialogue: while *Vedānta* speaks of *Brahman* and *māyā*, the technical community speaks of distributions, noise, and model convergence. Without effective analogical bridges, Advaita's relevance to contemporary discussions on knowledge and reality remains constrained.

The objective of this paper is to reinterpret *Advaitavedānta*'s categories through the conceptual framework of machine learning. By aligning metaphysical abstractions with computational processes — such as the relation between training data and true distribution, the iterative minimization of error, and the pitfalls of overfitting — the study develops a pedagogical model that makes non-dualism intelligible to audiences trained in ML.

The contribution of this paper is threefold. First, it develops a systematic mapping between *Advaitavedānta* and machine learning concepts, equating *Brahman* with the true data distribution, *māyā* with finite training samples, *avidyā* with overfitting, and iterative optimization with the *Advaitic* method of superimposition and negation. Second, it illustrates this mapping through a worked example of handwritten digit recognition using logistic regression in PyTorch, thereby grounding the philosophical analogies in a supervised-learning pipeline familiar to contemporary practitioners. Finally, it proposes a pedagogical framework through which instructors can present *Advaita*'s insights using computational analogies, while also encouraging technical audiences to reflect critically on the epistemic limits of model-based reasoning.

II. CORE CONCEPTS OF *ADVAITAVEDĀNTA*

Advaitavedānta, consolidated by *Śaṅkara* in the eighth century CE, presents a non-dualistic framework of reality that distinguishes between ultimate truth and empirical appearance [1–3].

Its ontology can be understood through four interrelated categories, illustrated schematically in **Fig. 1**. *Brahman*. At the center of *Advaita* metaphysics is Brahman, the absolute, unconditioned reality. The Upaniṣads describe Brahman as "*satyam jñānam anantam*"—truth, knowledge, and infinity (*Taittirīyopaniṣad* 2.1). It is not inert matter but consciousness itself, without beginning or end. Similarly, the *Bṛhadāraṇyakopaniṣad* states: "*vijñānam ānandaṃ brahma*"—*Brahman* is of the nature of knowledge and bliss (3.9.28). In *Advaita* interpretation, *Brahman* is not one object among others but the very ground of all that appears, represented at the core of Fig. 1. *Ātman* and *Jīva*. The inner self (*Ātman*) is, in its essence, identical with *Brahman*. This non-dual identity is explicitly declared in the *Bṛhadāraṇyakopaniṣad*: "*ayam ātmā brahma*"—this self is *Brahman* (2.5.19). However, through limitation and misidentification, the individual self (*jīva*) appears as finite and embodied. As illustrated in Fig. 1, the *jīva* is depicted as a smaller self-enclosed within the veiling layer, signifying that the identity with *Brahman* remains obscured rather than absent [2]. *Māyā* and *Avidyā*. The persistence of misperception is explained through *māyā* (cosmic illusion) and *avidyā* (individual ignorance). *Māyā* projects multiplicity upon the substratum, while *avidyā* sustains the false identification of the self with transient forms. *Śaṅkara* encapsulates this view in the aphorism: "*brahma satyam jagan mithyā jīvo brahmaiva nāparaḥ*"—*Brahman* is the only reality; the world is an appearance; the individual self is no other than *Brahman*. In Fig. 1, this is represented as a semi-transparent layer encircling the self, filtering perception of *Brahman* and sustaining the empirical world of names and forms [1,3]. The classical method of *adhyāropa–apavāda* (superimposition and negation) is intended to remove this distortion layer by layer. *Mokṣa*. Liberation (*mokṣa*) is achieved when ignorance is dispelled, revealing the non-dual identity of *Ātman* and *Brahman*. The *Taittirīyopaniṣad* (3.1) describes *Brahman* as that "from which all beings are born, by which they live, and into which they dissolve." Realization of this truth is not the attainment of a new state but the recognition of an always-present reality. In Fig. 1, *mokṣa* is represented as an arrow cutting through *māyā* and *avidyā*, allowing the *jīva* to realize its unity with *Brahman* at the core.

Taken together, these categories form the foundation for the analogical mapping developed later in this paper, where each *Advaitic* concept is aligned with corresponding structures in machine learning.

International Journal of Recent Development in Engineering and Technology
Website: www.ijrdet.com (ISSN 2347-6435(Online) Volume 14, Issue 12, December 2025)

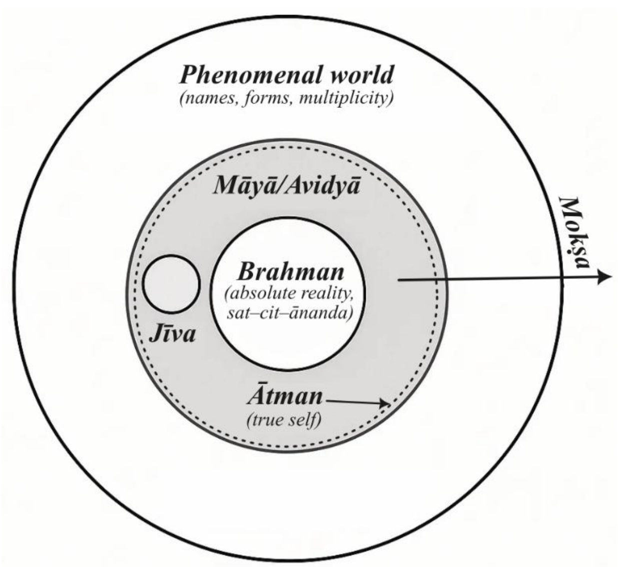

Figure 01: Core categories of *Advaitavedānta*: *Brahman* (absolute substratum), *Ātman* (true self), *Jīva* (conditioned self), *Māyā/Avidyā* (illusion and ignorance), and *Mokṣa* (liberation as realization of non-duality).

III. Mapping *Advaita* Concepts To ML Paradigms

Advaitavedānta's abstract categories can be meaningfully interpreted through the conceptual language of machine learning, a field where engineers and scientists routinely engage with notions of approximation, error, and convergence [4–6]. By aligning these categories with supervised learning workflows, philosophical abstractions are rendered accessible in terms familiar to technical communities. A schematic overview of the mappings is provided in Fig. 2.

3.1 Data and Distribution as Brahman

In machine learning, the true data distribution is the generative substratum from which all finite samples are drawn. This parallels *Advaitavedānta's* understanding of *Brahman* as the absolute reality underlying appearances [1,2]. Training datasets, which consist of finite and noisy examples, correspond to fragmented perceptions shaped by *māyā*. Just as models trained only on limited samples can never exhaustively capture the full distribution, the conditioned self (*jīva*) perceives only partial manifestations of *Brahman*.

3.2 Loss Minimization and Self-Inquiry

In machine learning, optimization is done by the minimization of loss, and so with every step in gradient descent, the prediction is made more certain and aligned with the real outcome [4,5]. This process to achieve the most optimal solution can be compared to *Advaita*'s method of self-inquiry, where ignorance (*avidyā*) is more and more systematically whittled down. Realizing the concepts of non-duality is similar to attaining the global minimum, while being stuck to local minima is akin to gaining partial knowledge that is self-contained and fails to rectify the more fundamental epistemic blunder. Regardless of the paradigm, both require refinement: in ML, there's iteration and with *Vedānta*, there's *adhyāropa–apavāda* (superimposition and negation).

3.3 Overfitting and Ignorance (Avidyā)

Fitting a model is said to be overfitted if it becomes too attached to the training data, and in doing so, perceives noise as reality, thus, failing to generalize [4]. This is akin to the notion of *avidyā*, where the self-distracts and identifies with fleeting forms and confuses them with the essence of existence [3]. Regularization approaches are akin to philosophical pondering: the model becomes less attached to the surface level and closer to the bet of the true distribution. This be it the Brahma, transcending the ephemeral distortions (*māyā*).

3.4 Hidden Layers and Levels of Reality

In the realm of neural networks, the hidden layers progressively learn from the raw input to obtain the final decision layers [5, 6]. This stratification provides an analogy of *Advaita*'s hierarchical ontology: the perceptual, the conceptual, and the ultimate. The empirical (*vyāvahārika*) and illusory (*prātibhāsika*) levels of truth is related to the intermediate activations, and knowing the truth non-duality at the output layer reflects *paramārthika* satya (absolute truth). Therefore, the deep hierarchies offer a pedagogical reflection of *Vedānta*'s hierarchy of realities.

3.5 Adversarial Examples and Māyā

One of the striking challenges in ML is adversarial vulnerability, where imperceptible perturbations lead models to misclassify inputs despite the presence of clear ground truth [6]. This phenomenon is evocative of *māyā*, which sustains illusions that appear convincing but conceal reality.

International Journal of Recent Development in Engineering and Technology
Website: www.ijrdet.com (ISSN 2347-6435(Online) Volume 14, Issue 12, December 2025)

Just as adversarial noise manipulates perception without altering the underlying data, *māyā* distorts cognition without modifying Brahman itself. The persistence of such misclassification underscores the importance of deeper awareness—in ML, robust training; in *Vedānta*, realization of non-duality.

Fig. 2. Mapping *Advaitavedānta* to machine learning paradigms. *Brahman* corresponds to the true distribution; finite training data reflects *māyā*; loss minimization parallels self-inquiry; overfitting exemplifies *avidyā*; hidden layers mirror levels of reality; adversarial examples illustrate the deceptive power of *māyā*.

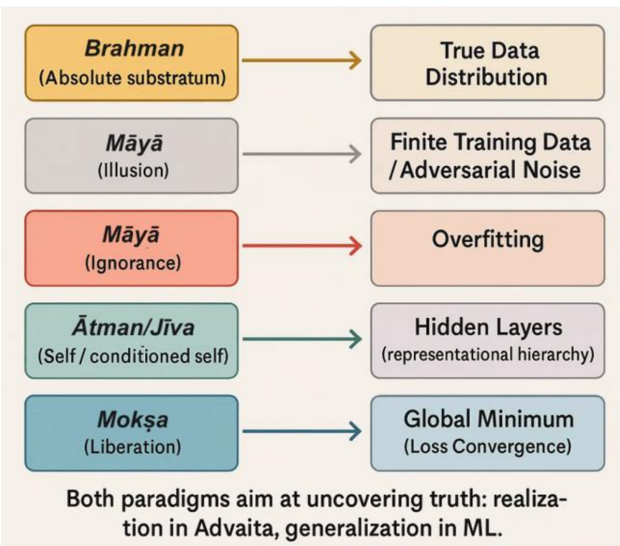

Figure 02: Mapping *Advaitavedānta* to machine learning paradigms. Brahman corresponds to the true distribution; finite training data reflects *māyā*; loss minimization parallels self-inquiry; overfitting exemplifies *avidyā*; hidden layers mirror levels of reality; adversarial examples illustrate the deceptive power of *māyā*.

IV. Case Illustrations From ML Practice

To illustrate the proposed analogy, we adopt the canonical MNIST dataset of handwritten digits, employing logistic regression implemented in PyTorch.

The experimental design follows standard open-source implementations [4–6,10], specifically the tutorial by GeeksforGeeks [10], where the model consists of a single linear layer mapping 784 input pixels to 10 output classes. Training is performed using stochastic gradient descent with cross-entropy loss over five epochs.

4.1 Analogy of Vedāntic Analogy through ML Algorithm

The workflow is summarized in Algorithm 1, which parallels Algorithm 2—a *Vedāntic* pathway towards *mokṣa*. Instead of reproducing code in full, we direct readers to [10] for exact implementation details.

Algorithm 1: MNIST Logistic Regression Training

1. Initialize model parameters.
2. Forward propagate batch inputs through the logistic regression model.
3. Compute cross-entropy loss between predicted labels and ground truth.
4. Backpropagate gradients to quantify the direction of error correction.
5. Update parameters using stochastic gradient descent (SGD).
6. Repeat steps 2–5 until training epochs are complete.
7. Evaluate model on test data to assess generalization accuracy.

Algorithm 2. *Vedāntic* Analogy for Liberation

1. Begin with empirical identifications of self as body, mind, and senses.
2. Engage in systematic inquiry through *śravaṇa* (hearing scriptures).
3. Internalize and rationalize through *manana* (reflection).
4. Identify ignorance (*avidyā*) as false associations with transient forms.
5. Apply *neti neti* ("not this, not this") to negate misperceptions.
6. Stabilize in *nididhyāsana* (deep meditation) and cultivate *viveka* (discrimination).
7. Realize the non-dual identity of *Ātman* and *Brahman*, attaining *mokṣa*.

International Journal of Recent Development in Engineering and Technology
Website: www.ijrdet.com (ISSN 2347-6435(Online) Volume 14, Issue 12, December 2025)

TABLE I
CORRESPONDENCE BETWEEN MACHINE LEARNING TRAINING STAGES AND ADWAITIC PROCESS OF LIBERATION

Machine Learning Stage	***Vedāntic* Analogy**	**Description**
Initialization	Empirical identifications	Model parameters start from arbitrary values; the *jīva* begins with mistaken identifications (body, mind, roles).
Input	Listening (*śravaṇa*)	Data batches enter the model; the seeker receives scriptural instruction and teaching.
Error	Ignorance (*avidyā*)	Prediction error reflects misperception; *avidyā* sustains the false sense of separation.
Correction	*Neti neti* (negation)	Gradient correction removes discrepancies; philosophical negation eliminates false identifications.
Update	Meditation/discrimination *(nididhyāsana, viveka)*	Parameters are updated toward convergence; seeker deepens discrimination and stabilizes insight.
Convergence	Liberation *(mokṣa)*	Model achieves optimal generalization; seeker realizes non-duality, perceiving *Ātman = Brahman*.

The parallelism between the machine learning pipeline and the *Advaitic* path to liberation can be systematically illustrated through Algorithm 1 (MNIST Logistic Regression Training) and Algorithm 2 (*Vedāntic* Analogy for Liberation). In Algorithm 1, the training process begins with the arbitrary initialization of model parameters, followed by successive cycles of forward propagation, error computation, gradient-based correction, and parameter updates until convergence is achieved and generalization accuracy is established. Algorithm 2 mirrors this structure: the seeker begins with empirical identifications of self as body and mind, listens to scriptural teachings (*śravaṇa*), reflects on them (*manana*), identifies ignorance (*avidyā*), and applies the method of *neti neti* (negation of false identifications).

This culminates in stabilization through meditation (*nididhyāsana*) and discrimination (*viveka*), ultimately converging in the realization of non-duality (*mokṣa*). The correspondence is further distilled in Table 1, where each stage of ML training finds its analogue in *Vedāntic* inquiry: initialization aligns with mistaken identifications, error with ignorance, correction with *neti neti*, update with meditative discrimination, and convergence with liberation. Taken together, these mappings demonstrate how iterative reduction of loss in ML parallels the systematic removal of ignorance in *Advaitavedānta*, thereby providing an intuitive pedagogical bridge for computational audiences.

4.2 Layered Representations and Advaitic Ontology

The architecture of machine learning models naturally lends itself to comparison with the stratified reality described in *Advaitavedānta*. Fig. 3 illustrates this layered correspondence.

- Input Layer → Sensory Perception (*pratyakṣa*): Raw pixel intensities of MNIST digits are analogous to the immediate sense-data through which the individual (*jīva*) perceives the empirical world. These impressions are fragmented and context-dependent, offering no direct access to the substratum.
- Hidden Representations → Conceptual Constructions (*vyāvahārika satya*): Through linear transformations, the model organizes inputs into meaningful intermediate features. Similarly, the mind constructs categories, language, and causal narratives, which, while pragmatically useful, are products of *māyā*.
- Decision Layer → Absolute Truth (*paramārthika satya*): The final classification into digits represents the convergence toward recognition of a single underlying identity. This mirrors the *Vedāntic* claim that beyond appearances lies the realization that *Ātman = Brahman*.

Thus, the training dynamics of even a shallow logistic regression model exemplify how cognition progresses from multiplicity to unity.

4.3 MNIST Training Dynamics

The logistic regression classifier was trained on the MNIST dataset using stochastic gradient descent with a learning rate of 0.001 across five epochs. Training curves (loss vs. epoch, Fig. 4a) demonstrated steady reduction in cross-entropy error, indicating progressive alignment between predictions and true labels.

International Journal of Recent Development in Engineering and Technology
Website: www.ijrdet.com (ISSN 2347-6435(Online) Volume 14, Issue 12, December 2025)

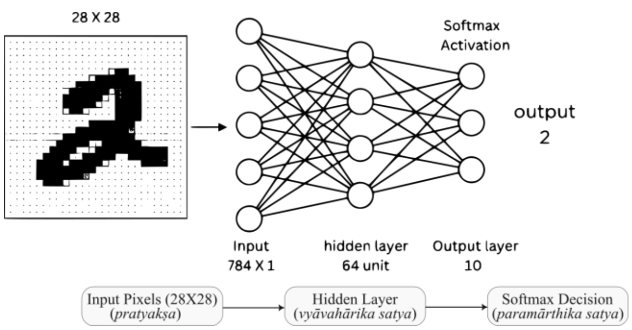

Figure 03: Layered analogy diagram [10] — input pixels (*pratyakṣa*), hidden layers (*vyāvahārika satya*), final classification (*paramārthika satya*)

On the held-out test set of 10,000 images, the model achieved an accuracy of 92–94%, consistent with canonical benchmarks. Misclassifications largely occurred between visually similar digits (e.g., "4" vs. "9" or "3" vs. "5"), which points to the limitations of linear models in capturing deeper invariant structures (Figure 4b).

These results emphasize a key pedagogical insight: the model learns to generalize beyond its finite samples toward the underlying digit distribution. This generalization ability, limited but real, is directly analogous to the seeker transcending superficial identifications to intuit the substratum of *Brahman*.

4.4 Vedāntic Interpretation of Results

The MNIST outcomes acquire philosophical depth when interpreted through *Advaita*:

- Loss Reduction ↔ Dissolution of Ignorance (*avidyā*): Each epoch of training reduces prediction error, paralleling the way *śravaṇa* (listening), *manana* (reflection), and *nididhyāsana* (deep meditation) progressively eliminate epistemic errors sustained by *avidyā*.
- Avoiding Overfitting ↔ Detachment from *māyā*: When models cling too rigidly to training data, they mistake noise for truth, failing to generalize. This mirrors the plight of the conditioned self that mistakes transient phenomena for reality. Techniques like regularization resemble *Vedāntic* discrimination (*viveka*), which prevents attachment to appearances.
- Convergence ↔ Liberation (*mokṣa*): The point at which training stabilizes and the model reliably recognizes unseen digits is akin to *mokṣa*: the recognition that what seemed manifold (digits/forms) is underpinned by a single truth.

Even the residual error has interpretive value. Just as a linear model cannot capture the full complexity of digit variations, intellectual reasoning alone cannot fully grasp *Brahman*. At best, models and methods point toward the reality that transcends them.

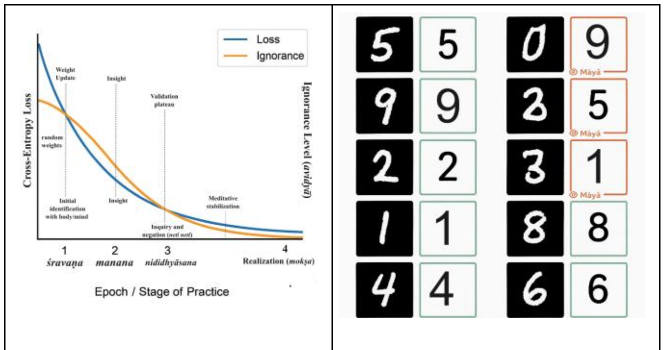

Fig. 4: (a) Loss vs. epoch curve (showing monotonic decline), paralleled with a conceptual "ignorance vs. inquiry" curve in *Vedānta;* (b) Sample misclassifications. Although the true digit (left) remains unchanged, the model predicts an incorrect label (right). These errors, highlighted and annotated as *māyā*, symbolize how illusion can obscure but not alter the underlying reality.

V. Conclusion

This paper developed a novel pedagogical analogy between *Advaitavedānta* and machine learning, aligning core philosophical categories with computational processes. *Brahman* was interpreted as the underlying data distribution, *māyā* as the world of finite samples, *avidyā* as overfitting to noise, and *mokṣa* as convergence to truth. The iterative refinement of a model through gradient descent was compared to the *Advaitic* method of *adhyāropa–apavāda* (superimposition and negation), while hidden layers and adversarial examples were shown to reflect the veils of perception and illusion.

The case study on handwritten digit recognition using logistic regression in PyTorch illustrated these ideas in practice. The monotonic decline of loss across epochs mirrored the reduction of ignorance, while misclassifications demonstrated how *māyā* can distort recognition without altering the substratum. These parallels do not equate machine learning and philosophy but make complex metaphysics accessible to engineers and data scientists by using familiar technical language.

The analogy also invites a broader reflection: just as no model can perfectly capture a data distribution, conceptual reasoning cannot fully capture Brahman. Recognizing these limits encourages epistemic humility in both fields.

International Journal of Recent Development in Engineering and Technology
Website: www.ijrdet.com (ISSN 2347-6435(Online) Volume 14, Issue 12, December 2025)

Future work could extend this framework to deeper models or more formal mappings, enriching interdisciplinary teaching and dialogue.

REFERENCES

[1] N.F. Bauer, Advaita Vedānta and contemporary Western ethics, Philos. East West 37 (1) (1987) 1–15.

[2] P. Williams, Review: Encyclopedia of Indian Philosophies, Vol. III: Advaita Vedānta up to Śaṅkara and His Pupils (K.H. Potter, ed.), J. R. Asiatic Soc. Gt. Brit. Irel. (1) (1983) 166–168.

[3] B.K. Matilal, Epistemology, logic and grammar in Indian philosophical analysis, Philos. East West 21 (3) (1971) 321–336.

[4] S. Badillo, A. Banfai, E. Birzele, et al., An introduction to machine learning, Clin. Pharmacol. Ther. 107 (4) (2020) 871–885.

[5] A.J. Smola, An introduction to machine learning: L1 basics and probability theory, MLSS, Tata Institute of Fundamental Research, Pune, 2007.

[6] G. Rätsch, A brief introduction into machine learning, Friedrich Miescher Laboratory, Max Planck Society, Tübingen, 2007.

[7] J. Frazier, Vedāntic Traditions in Hindu Studies, J. Hindu Stud. 8 (1) (2015) 1–15. https://doi.org/10.1093/jhs/hiv013

[8] M.S. Allen, Vedānta: A Survey of Recent Scholarship (I), J. Indian Philos. 51 (2023) 731–759. https://doi.org/10.1007/s10781-023-09551-7

[9] J.D. Long, The Study of Indian Philosophy in the Work of Arvind Sharma (with Particular Reference to Advaita Vedānta), Int. J. Hindu Stud. 28 (2024) 91–104. https://doi.org/10.1007/s11407-024-09366-5

[10] GeeksforGeeks, "Identifying Handwritten Digits using Logistic Regression in PyTorch," GeeksforGeeks, 2024. Available at: https://www.geeksforgeeks.org/machine-learning/identifying-handwritten-digits-using-logistic-regression-pytorch/