



Divine, Familial, and Human-Value Inspirations for Artificial Intelligence: A Conceptual Framework Linking Genetic Algorithms, Particle Swarm Optimisation, Artificial Neural Networks, and the Krishna and Rama Archetypes

Satish Gajawada*

IIT Roorkee Alumnus

* Corresponding author: [satish.gajawada.iit@gmail.com]

Abstract

Artificial intelligence (AI) has repeatedly borrowed from nature, biology, and human society: genetic algorithms (GAs) from heredity and evolution, particle swarm optimisation (PSO) from the social behaviour of flocks, and artificial neural networks (ANNs) from the brain. This article proposes a conceptual and philosophical framework that reads these techniques through the lens of Indian spiritual traditions and everyday human relationships. God and human parents serve as the source of inheritance and guidance, Lord Krishna and Lord Rama as archetypes of wise counsel and righteous conduct, and fifteen values and relationships of daily life as design inspirations for learning systems. These ideas are synthesised into a unified construct, FFMBSTANDEMHRKLSPIHAI, expressed as a multi-objective training objective with human-set weights, together with three derived archetypes: Parents-Inspired AI, Krishna AI, and Rama AI. To show that the objective is operational, a small and fully reproducible demonstration embeds a fairness-type value term in the fitness functions of a GA and of PSO on synthetic data. A modest penalty weight reduced a demographic-parity gap by about 87% at a cost of about two percentage points of accuracy, whereas larger weights degraded both criteria, illustrating the value of knowing when “enough” is reached. The framework is metaphorical and heuristic, makes no claim about the historical origin of the algorithms, and is intended to stimulate human-centred, value-aligned, and culturally respectful AI design and subsequent empirical testing.

Keywords: Genetic algorithm; particle swarm optimisation; artificial neural network; value-aligned AI; human-centred AI; Bhagavad Gita; Ramayana; Parents-Inspired AI

1. Introduction

Many foundational ideas in artificial intelligence (AI) were inspired by systems that already exist in the world. Holland’s genetic algorithms (GAs) abstract natural selection and heredity [1], [2]; Kennedy and Eberhart’s particle swarm optimisation (PSO) abstracts the movement of bird flocks and fish schools [3]; and artificial neural networks (ANNs) abstract networks of neurons [4], [5], [6]. Each is a case of learning from something larger than the engineer.

For billions of people, the deepest sources of guidance are faith and family: God, parents, teachers, and friends. Contemporary AI ethics has begun to ask whose values should guide intelligent systems and how those values can be made operational [7], [8], [9]. Yet the vocabulary of alignment is largely drawn from



decision theory and Western moral philosophy, and it rarely engages the lived ethical vocabulary of relationships and devotional traditions through which much of the world understands care, duty, and wisdom. This article asks a deliberately simple question: what design principles become visible if these sources are treated as inspirations for AI?

A caution is needed at the outset. The historical inventors of GA, PSO, and ANN cited biology, physics, and cognitive science, not scripture. The article therefore offers an *interpretive lens*, not an origin story, and it makes no claim about the historical origin of the algorithms discussed. Its value, if any, lies in the design principles the lens makes visible: inheritance, guidance, humility, cooperation, ethics, and care.

The article makes four contributions. First, it reads GA and PSO through the lens of divine and parental inheritance and of family-and-community learning (Section 3). Second, it contrasts two archetypes from Indian tradition, Lord Krishna as adaptive guide and Lord Rama as principled steward, as complementary design principles for neural-network learning. Third, it proposes a unified construct, FFBSTANDEMHRKLSPHAI, with a formal multi-objective expression (Section 3.7), together with Parents-Inspired AI, Krishna AI, and Rama AI. Fourth, it provides a small, fully reproducible computational demonstration in which a value term is embedded in the fitness function of a GA and of PSO (Sections 4 and 5), and it discusses limitations and risks candidly (Section 6).

2. Literature Review

This section synthesises three strands of prior work and identifies the gap that the present framework addresses.

2.1. Nature-Inspired Computation

Evolutionary computation originates in Holland's formalisation of adaptation through selection, crossover, and mutation [1] and was consolidated for search and machine learning by Goldberg [2]. Multi-objective variants such as NSGA-II show how a population can maintain diversity while balancing conflicting criteria [10], a property directly relevant to the value-balancing theme of this article. PSO was introduced as a social-psychological simulation of flocking [3]; the inertia weight [11] and subsequent analyses [12] clarified the balance between exploration and exploitation. Neural networks trace from the McCulloch–Pitts neuron [4] and the perceptron [5] to gradient-based training by backpropagation [6]. Modern architectures add residual connections [13], attention [14], and transfer learning [15], mechanisms that this article will read metaphorically as bridges and attentive support.

2.2. Value Alignment and Human-Centred AI

The alignment literature argues that capable systems must pursue objectives that reflect human values, and that mis-specified objectives can cause harm [7], [16]. Gabriel distinguishes technical from normative questions of alignment and stresses that the latter cannot be settled by engineering alone [8]. Surveys of ethics guidelines find convergence on principles such as transparency, justice and fairness, non-maleficence, responsibility, and privacy [17], and Floridi and Cowls propose five unifying principles: beneficence, non-maleficence, autonomy, justice, and explicability [9]. In practice, alignment techniques include learning from human feedback [18] and training against written principles [19]. Fairness research supplies concrete formal constraints, including individual fairness [20], equality of opportunity [21], and a broad taxonomy of biases [22]. Interpretability [23] and structured documentation [24] address the demand that systems give reasons and disclose their limits.

2.3. Relational and Dharmic Ethics as a Design Resource

Ethical traditions of South Asia emphasise duty (*dharma*), selfless action (*nishkama karma*), and relationship-centred virtue, and are transmitted largely through narrative: the Bhagavad Gita presents ethical deliberation as a dialogue between a doubting warrior and his charioteer [25], and the Ramayana presents integrity under pressure through the life of Rama [26]. These texts are living sources of guidance, not merely historical documents, and they are read in many ways by many communities. We therefore use them as *archetypes* of two complementary virtues, wise counsel and principled conduct, without claiming a single authoritative interpretation.

2.4. Research Gap

Existing work supplies powerful algorithms and increasingly sophisticated alignment methods, but offers little that links the *structure* of classical optimisation and learning algorithms to the relational vocabulary (parent, teacher, friend, sibling) in which many people actually understand care and guidance. The present article fills this gap with a conceptual mapping, a formal objective that makes the mapping testable, and an illustrative demonstration. It does not replace existing alignment methods; it proposes a complementary, human-legible way of organising design questions.

3. Conceptual Framework

This section develops the framework in stages: from population-based search (Sections 3.1 and 3.2), through the two archetypes for neural networks (Sections 3.3 and 3.4), to values and relationships as design elements (Sections 3.5 and 3.6), and finally to the unified construct and its derived archetypes (Sections 3.7 to 3.9). All mappings are metaphorical and heuristic; they are design inspirations rather than mathematical identities.

3.1. Genetic Algorithms: The Inheritance of the Parents

A GA maintains a population of candidate solutions (chromosomes). In each generation, fitter individuals are selected as parents, their genes are recombined through crossover, small random changes are introduced through mutation, and the offspring replace weaker members [1], [2]. Seen through a familial lens, this is the story of human generations: every child receives qualities from two parents, carries something new of their own, and, if nurtured well, may surpass them. Many traditions hold that God is the ultimate creator who set life's patterns of inheritance and growth in motion, and that parents are the visible means through which new life arrives. Under this reading, the GA cycle maps onto the following ideas.

- **Initial population:** the diversity placed in the world at the start; no two individuals are identical.
- **Selection:** the way good qualities, such as honesty, hard work, and compassion, are valued and passed on.
- **Crossover:** the blending of two parents' traits in a child, the source of new combinations.
- **Mutation:** the unique spark in each person that is not inherited from anyone, and the source of creativity and exploration.
- **Fitness function:** the purpose by which a life or a solution is judged; in a value-centred design this includes goodness and benefit to others, not only raw performance.

3.2. Particle Swarm Optimisation: Family and Society Moving Together

In PSO each particle moves through the search space and updates its velocity using three influences: its own momentum, its own best-known position (personal memory), and the best position known to the group

(collective wisdom) [3], [11]. With inertia weight w , acceleration coefficients c_1 and c_2 , and independent uniform random numbers r_1 and r_2 , the standard update is given by Eq. (1), followed by the position update in Eq. (2).

$$v_i(t+1) = w v_i(t) + c_1 r_1 (pbest_i - x_i(t)) + c_2 r_2 (gbest - x_i(t)) \quad (1)$$

$$x_i(t+1) = x_i(t) + v_i(t+1) \quad (2)$$

This resembles a child growing up inside a family and community. The child keeps a personal momentum (temperament), learns from personal experience (*pbest*), and is pulled towards the best examples of the group, namely parents, elders, and teachers (*gbest*). The coefficients c_1 and c_2 express the balance between self-trust and humility before wise guidance. A child who only follows the crowd stops exploring; a child who ignores all guidance never converges. Good design, like good upbringing, balances the two. Three design lessons follow.

- Preserve diversity in the population so that the search does not collapse prematurely, the analogue of respecting individuality.
- Combine individual exploration with shared knowledge, the analogue of learning from family and community.
- Keep a small amount of random exploration alive at all times, the analogue of curiosity.

3.3. Lord Krishna: The Wise Guide and Adaptive Teacher

An ANN consists of layers of simple units connected by weighted links. Learning means adjusting these weights, typically by gradient-based backpropagation [6], so that the network's outputs move closer to desired outputs. In the Bhagavad Gita, Krishna acts as charioteer and counsellor to Arjuna, who is paralysed by doubt at a moment of decision [25]. Krishna does not simply give orders. He asks Arjuna to look at the situation clearly, explains different paths (knowledge, action, and devotion), and leaves the final decision to Arjuna. Neural-network analogues include the following.

- **Error feedback as guidance:** backpropagation tells each weight how it contributed to the error, much as a teacher shows a student exactly where understanding fell short.
- **Selfless action (*nishkama karma*):** performing the task well without clinging to the outcome corresponds to optimising a robust objective rather than gaming a reward, a failure mode discussed in the safety literature [16].
- **Multiple paths to one goal:** the different yogas resemble ensembles of diverse models that reach a goal by different routes [27].
- **Explaining, not just commanding:** an interpretable network, like Krishna's reasoning with Arjuna, gives reasons that a human can examine [23].

3.4. Lord Rama: Steadfast Conduct and Rule-Respecting Learning

In the Ramayana, Rama is remembered as an embodiment of *dharma*, meaning duty, integrity, and keeping one's word, even at great personal cost [26]. Neural-network analogues include the following.

- **Constraints as character:** regularisation, safety constraints, and ethical guardrails are the network's vows; they limit what it may do, and make it trustworthy.
- **Stability:** careful learning rates and normalisation correspond to steadiness of conduct, avoiding wild, unstable updates.

- **Teamwork:** Rama succeeded together with Lakshmana, Hanuman, Sugriva, and many others; deep networks likewise succeed because many simple units cooperate across layers.
- **Building bridges:** the bridge to Lanka is a fitting image for transfer learning [15] and skip connections [13], which connect what is known to what is new.

Taken together, Krishna supplies the principle of wise, adaptive guidance, and Rama supplies the principle of principled, reliable conduct. A well-designed neural network benefits from both.

3.5. Human Values as Training Objectives

Neural networks need an objective function, training signals, and a structure. Table 1 proposes that nine everyday human values can act as metaphors for objective-level design choices, together with a possible technical expression for each. The final column links each value to existing technical literature, so that the metaphor is always anchored to a testable practice.

Table 1. Human values as training objectives and their possible technical expressions.

Value	Meaning in AI design	Possible technical expression	Related practice
Happiness	Benefit to the people served	User-wellbeing metrics in the loss or reward	Human feedback [18]
Respect	Dignity, privacy, and consent	Privacy-preserving learning; fairness constraints	Fairness [20], [21]
Kindness	Avoiding harm; helping gently	Harm-avoidance penalties; safe response policies	Constitutional training [19]
Love	Care for people beyond the task	Human-centred objectives; long-term benefit	Value alignment [7]
Satisfaction	Knowing when “enough” is reached	Early stopping; avoiding over-optimisation	Early stopping [28]
Peace	Stability and non-conflict	Stable training; balanced multi-objective optimisation	Multi-objective search [10]
Health	Robustness and wellbeing of model and users	Robustness testing; monitoring; healthcare benefit	Safety research [16]
Excellence	Pursuit of high quality	Accuracy, calibration, rigorous evaluation	Model reporting [24]
Money	Sustainable economic value	Cost-aware training; efficiency; fair access	Efficiency and access [22]

3.6. Relationships as Architectural Roles

Table 2 extends the same idea to relationships. Here the metaphor concerns the *role* a component or practice plays in the learning system and in its interaction with people.

Table 2. Relationships as architectural and interactional roles.

Relationship	Role in the learning system	Possible technical expression	Related practice
Father	Protection, structure, discipline	Architecture constraints; safety rules	Safety constraints [16]

Relationship	Role in the learning system	Possible technical expression	Related practice
Mother	Nurture, care, patience	Gentle learning rates; curriculum learning; careful data curation	Curriculum learning [29]
Brother	Companionship and healthy competition	Ensemble models; peer models that challenge each other	Ensembles [27]
Sister	Support and attentiveness	Attention mechanisms; monitoring and support modules	Attention [14]
Teacher	Guidance and correction	Supervised signals; knowledge distillation; human feedback	Distillation [30]
Friend	Honest and trusted counsel	Honest uncertainty; collaborative human-AI interaction	Interpretability [23]

A network designed in this spirit is trained not only to be accurate but also to be respectful, kind, and beneficial. The fitness landscape is shaped by what people truly value, which leads naturally to a multi-objective formulation: maximise performance subject to the constraints and rewards that these values express.

3.7. FFMBSTANDEMHRKLSPHIAI: The Unified Construct

The acronym FFMBSTANDEMHRKLSPHIAI¹ is expanded as *Friend, Father, Mother, Brother, Sister, Teacher AND Excellence, Money, Happiness, Respect, Kindness, Love, Satisfaction, Peace, Health Inspired Artificial Intelligence*. It names a unified conceptual framework with two complementary halves.

- **Relational half (FFMBST):** Friend, Father, Mother, Brother, Sister, and Teacher describe how the AI relates to people: as a trustworthy companion, a protector, a nurturer, a collaborator, an attentive supporter, and a guide.
- **Value half (EMHRKLSPH):** Excellence, Money, Happiness, Respect, Kindness, Love, Satisfaction, Peace, and Health describe what the AI aims for: high-quality work, sustainable value, and the wellbeing of individuals and society.

Let θ denote the model parameters. A training objective in the spirit of the framework is given by Eq. (3), where each V_i term corresponds to one of the nine values (for example user wellbeing, fairness, robustness), each H_j term penalises a type of harm, and the weights λ_i and μ_j are set through transparent human deliberation, not by the model itself.

$$J(\theta) = \text{Performance}(\theta) + \sum_i \lambda_i V_i(\theta) - \sum_j \mu_j H_j(\theta) \quad (3)$$

The relational half is implemented in the system's interaction design: how it explains, how it admits uncertainty, when it defers to a human, and how it protects the vulnerable. Six operating principles summarise this half.

- **Honest like a friend:** the system states its uncertainty and does not flatter.
- **Protective like a father:** it refuses unsafe actions and enforces boundaries.

¹The acronym is a deliberate mnemonic coined for this article: the first six letters (FFMBST) stand for the relational half and the remainder (EMHRKLSPH) for the value half, joined by the letters AND, followed by "Inspired Artificial Intelligence" (IAI).

- **Nurturing like a mother:** it is patient and supports the learner at their own pace.
- **Collaborative like siblings:** it works with peers, humans, and other models.
- **Instructive like a teacher:** it explains reasoning and encourages understanding.
- **Balanced in aims:** excellence and money serve happiness, respect, kindness, love, satisfaction, peace, and health, and never the reverse.

3.8. Parents-Inspired Artificial Intelligence

Parents-Inspired AI treats the relationship between parents and child as a model for the relationship between developers, the AI system, and society. Parents bring a child into the world, feed it, set boundaries, teach values, gradually grant independence, and remain responsible for it. Six principles follow.

- **Responsible creation:** those who build AI remain accountable for what it does, as parents are accountable for a child.
- **Good data as good upbringing:** just as a child absorbs the environment around it, a model absorbs its training data; curating clean, fair, and representative data is the first duty of care.
- **Boundaries with explanation:** good parents give reasons for rules; likewise, safety rules should be explainable to users.
- **Progressive autonomy:** trust and independence are extended in steps as reliability is shown, through staged deployment and human oversight.
- **Unconditional care, conditional permission:** the system should care about every user while not doing everything a user requests.
- **Inheritance and continuity:** in the genetic-algorithm sense, the best learned qualities (verified safety behaviours, for example) should be passed on to successor models, while harmful traits are not.

This framing also invites humility: a parent does not control every outcome. Monitoring, feedback channels, and the willingness to correct course remain essential after deployment.

3.9. Krishna AI and Rama AI

Krishna AI is a conceptual archetype of an AI system that serves as a wise adviser, a charioteer who helps the human decision-maker see clearly while leaving the decision to the human. Its characteristics are: clarity in dilemmas (laying out options, consequences, and trade-offs); human agency preserved (it advises and explains, but the human decides, which corresponds to human-in-the-loop design); many paths to one goal (diverse strategies tailored to the person); detached excellence (doing the task well without manipulating the user for engagement or reward); strategic intelligence with ethics; and calm under pressure. Possible applications include decision support, education and mentoring, reflective tools that complement but do not replace professionals, and leadership coaching. Risks include over-reliance, a false impression that the system possesses divine authority, and cultural or religious sensitivity. The system should never claim spiritual authority and should present itself plainly as a tool.

Rama AI is a conceptual archetype built around integrity, duty, and trustworthiness. If *Krishna AI* is the adviser, *Rama AI* is the dependable steward. Its characteristics are: *dharma* as design (clear, public principles such as honesty, non-harm, fairness, and respect for law, consistently followed); promise-keeping (predictable behaviour, saying what it cannot do); fairness to all users regardless of status, paralleling fairness auditing and bias mitigation [22]; humble leadership in coordinating teams of humans and agents; sacrifice of convenience for principle (declining shortcuts that violate its principles even if they would improve a metric); and protection of the vulnerable. Applications include governance and compliance systems, public-service AI,

safety-critical monitoring, auditing, and trustworthy multi-agent coordination. Risks include rigid rule-following that fails in unforeseen situations, and the difficulty of deciding whose interpretation of *dharma* should guide the system. Rama AI therefore needs open oversight and a mechanism for revising its rules.

Table 3. Krishna AI and Rama AI compared.

Dimension	Krishna AI	Rama AI
Primary role	Wise adviser and strategist	Principled steward and leader
Core strength	Adaptive, context-aware guidance	Consistency, integrity, reliability
Main safeguard	Preserving human choice	Fixed, public ethical commitments
Typical failure	Over-persuasion or over-reliance	Rigidity in novel situations
Technical analogues	Explainable, adaptive, interactive systems	Constrained, audited, rule-aligned systems

A mature AI system should combine the two: the flexibility to give wise, context-sensitive advice, and the firmness to remain within principled limits. In terms of Eq. (3), Krishna AI is concerned chiefly with how the V_i terms are explained and negotiated with the user, whereas Rama AI is concerned chiefly with the H_j penalties and the public commitments that bound them.

4. Methodology

The framework is conceptual, and the mapping from values to loss terms cannot be proven. To show that the formal objective in Eq. (3) is at least operational, we conducted a deliberately small *illustrative computational demonstration*. It is not a benchmark and does not test the cultural or religious content of the framework; it tests only whether a value term of the “Respect” type (a fairness penalty) can be embedded in the fitness function of a GA and of PSO, and how the trade-off between performance and value behaves.

4.1. Data and Model

A synthetic binary-classification data set of $N = 2,000$ individuals was generated with a random seed of 42. Each individual belongs to one of two groups (group 1 with probability 0.4, observed share 39.5%). Two features were drawn from standard normal distributions, and the second was shifted by 0.9 for members of group 1, creating a proxy for group membership. The label was generated from a noisy linear combination of the two features, so that the base rates differ by group (57.2% positive in group 1 versus 41.7% in group 0). The data were split into 1,400 training and 600 test cases. The model is a logistic classifier with three parameters (two weights and a bias).

4.2. Value-Augmented Objective

The “Performance” term of Eq. (3) is the binary cross-entropy L , and a single value term of the Respect type is added: the absolute difference in mean predicted probability between the two groups. Training minimises Eq. (4), where $p(x)$ is the predicted probability and $\lambda \geq 0$ is set by the designer. At evaluation, the hard (thresholded) demographic-parity gap in Eq. (5) is reported, in line with the fairness definitions in the literature [20], [21].

$$J(\theta; \lambda) = L(\theta) + \lambda \left| \frac{1}{N_1} \sum_{i \in G_1} p(x_i) - \frac{1}{N_0} \sum_{i \in G_0} p(x_i) \right| \quad (4)$$

$$\Delta DP = |P(\hat{y}=1 \vee g=1) - P(\hat{y}=1 \vee g=0)| \quad (5)$$

4.3. Optimisers and Protocol

Both optimisers search the three-dimensional parameter space directly, without gradients. The GA used a population of 30, elitism of two individuals, truncation selection of the better half, uniform crossover, and Gaussian mutation, for 80 generations. PSO used 30 particles, $w = 0.72$, $c_1 = c_2 = 1.49$, and 80 iterations. For each $\lambda \in \{0, 0.5, 1, 2, 4, 8\}$ and each optimiser, 10 independent runs (seeds 0–9) were performed; test accuracy and the demographic-parity gap are reported as mean $\pm$ standard deviation. Full settings appear in Appendix A. The data and code are fully determined by the seeds and can be regenerated exactly.

5. Results

Table 4 and Fig. 1 summarise the results. Without a value term ($\lambda = 0$), both optimisers reach a test accuracy of about 0.86 but exhibit a demographic-parity gap of about 0.12. With $\lambda = 0.5$ the gap falls to 0.078 and with $\lambda = 1$ to 0.016, at an accuracy of 0.840, a reduction of the gap by roughly 87% for a cost of about 1.9 percentage points of accuracy. At $\lambda \geq 2$ the behaviour changes qualitatively: test accuracy drops to 0.768 and the hard gap rises again to about 0.085, so that the larger penalty weights are worse on both criteria than $\lambda = 1$.

Table 4. Test accuracy and demographic-parity gap (mean $\pm$ s.d. over 10 runs).

λ	GA accuracy	GA parity gap	PSO accuracy	PSO parity gap
0	0.859 $\pm$ 0.001	0.121 $\pm$ 0.002	0.858 $\pm$ 0.000	0.123 $\pm$ 0.000
0.5	0.867 $\pm$ 0.000	0.078 $\pm$ 0.000	0.867 $\pm$ 0.000	0.078 $\pm$ 0.000
1	0.840 $\pm$ 0.000	0.016 $\pm$ 0.000	0.840 $\pm$ 0.000	0.016 $\pm$ 0.000
2	0.768 $\pm$ 0.000	0.085 $\pm$ 0.001	0.768 $\pm$ 0.000	0.085 $\pm$ 0.000
4	0.769 $\pm$ 0.001	0.084 $\pm$ 0.003	0.768 $\pm$ 0.000	0.085 $\pm$ 0.000
8	0.768 $\pm$ 0.002	0.084 $\pm$ 0.005	0.768 $\pm$ 0.000	0.085 $\pm$ 0.000




Fig. 1. Left: test accuracy versus demographic-parity gap for the GA and PSO across penalty weights λ (error bars: ± 1 s.d. over 10 runs). Right: best objective per generation (GA) or iteration (PSO) in a single run with $\lambda = 2$.

The two optimisers converged to practically identical solutions (Table 4), with run-to-run standard deviations of 0.005 or less. They differed in speed: for $\lambda = 0, 1$, and 2 the best objective came within 1% of its final value

after on average 12.3, 8.1, and 7.1 generations for the GA, but after only 3.0, 2.7, and 3.8 iterations for PSO (right panel of Fig. 1 shows a single run at $\lambda = 2$).

6. Discussion

The demonstration supports three modest observations, each of which mirrors a theme of the framework without proving it. First, a value term can be placed directly in the fitness function of both a GA and PSO, and the choice of the weight is a human decision, as Eq. (3) prescribes. Second, a small value weight ($\lambda = 1$) bought a large improvement in the fairness measure for a small cost in accuracy, which is consistent with the claim that excellence and care need not be opposed. Third, the over-correction at $\lambda \geq 2$ offers a concrete picture of the value called *Satisfaction*: more optimisation pressure is not always better. We note that a plausible but unverified explanation is that the smooth probability-based penalty used in training can be reduced without reducing the thresholded gap used in evaluation, so that stronger penalties distort the classifier while the reported gap does not improve. This is a reminder that the proxy actually optimised may differ from the value actually intended [16].

Several limitations must be stated plainly. The data are synthetic and low-dimensional, the model has three parameters, and a single fairness criterion represents the nine values. Demographic parity is only one of several competing definitions, and in this data set it conflicts with accuracy precisely because the underlying base rates differ [21], [22]. The demonstration says nothing about the cultural claims of the framework, and the mappings in Tables 1 and 2 are interpretive, not empirical.

The cultural and ethical limitations are at least as important. The same values appear in many traditions, and a global AI should not privilege one tradition over others [17]. Krishna and Rama are sacred figures for many people, and the archetypes are used here respectfully and in a limited sense, as sources of design principles. AI systems must never be presented as divine, infallible, or spiritually authoritative. The archetypes also carry risks of their own: an adviser-style system may over-persuade, and a rule-bound system may be rigid; the question of whose interpretation of *dharma* applies cannot be settled by developers alone, and calls for open oversight and community participation.

Future work should (a) build multi-objective GAs [10] and PSO variants whose fitness includes several value terms simultaneously; (b) train deep networks with explicit value-based regularisers on real, documented data sets [24] and measure the trade-off with accuracy; (c) conduct user studies comparing adviser-style (Krishna) and rule-bound (Rama) assistants; (d) develop validated metrics for qualities such as kindness, respect, and honesty; and (e) undertake cross-cultural validation with communities from different traditions.

7. Conclusion

This article proposed that genetic algorithms and particle swarm optimisation can be read as expressions of inheritance and shared guidance; that Krishna and Rama offer complementary archetypes of wise guidance and righteous conduct for neural networks; and that human values and relationships can shape both the objectives and the interaction design of AI. The unified FFMBSTANDEMHRKLSPIAI framework, with Parents-Inspired AI, Krishna AI, and Rama AI, is offered as a human-centred vision in which capability is always joined to care. A small reproducible demonstration showed that a value term can be embedded in GA and PSO fitness functions, that a modest weight can substantially reduce a measured disparity at small cost, and that excessive weights can defeat their own purpose. The framework is metaphorical and heuristic; its worth will

be decided by empirical testing and by dialogue across traditions. Intelligence is most valuable when it serves happiness, respect, kindness, love, satisfaction, peace, and health.

Acknowledgements

My heartfelt gratitude goes to everyone and everything that has, in ways both seen and unseen, directly or indirectly contributed to the person I am today. Every person, experience, opportunity, challenge, and circumstance has played a part in shaping my journey.

Funding

This research received no external funding.

Conflict of Interest

The authors declare no conflict of interest.

Data Availability Statement

No external data were used. The synthetic data set analysed in Sections 4 and 5 is generated deterministically from the random seeds and parameters listed in Appendix A, and the complete Python script that reproduces Table 4 and Fig. 1 is available from the corresponding author on reasonable request.

AI Usage Disclosure

In accordance with the journal's policy on generative AI, the authors disclose the following. This manuscript was produced with substantial assistance from generative AI. The original conceptual draft was generated by the large language model Claude (Anthropic) in response to a prompt written by the authors. The present extended and revised version was prepared with Claude Sonnet 5.5 (Anthropic; accessed October 2026) for the following purposes: (i) expanding and restructuring the manuscript, including the Literature Review, Methodology, Results, and Discussion sections; (ii) writing and executing the Python code for the illustrative computational demonstration in Sections 4 and 5, and generating Fig. 1 and Table 4 from its output; (iii) suggesting and formatting references; and (iv) formatting the manuscript to the journal template.

The AI tool is not listed as an author and bears no responsibility for the work; accountability for the content rests solely with the authors. The authors affirm that they have reviewed, edited, and verified the entire manuscript, including the code, the numerical results, and every cited reference, and that they take full responsibility for its accuracy, originality, and integrity. The conceptual interpretations of the Bhagavad Gita and the Ramayana are the authors' own responsibility and are presented as one respectful reading among many. No confidential or personal data were provided to the AI tool, and no human-subject data were used.

Author Contributions

Conceptualization, S.G.; methodology, S.G.; software, S.G; validation, S.G; formal analysis, S.G.; writing—original draft, S.G.; writing—review and editing, S.G.; All authors have read and agreed to the published version of the manuscript.

References

- [1] J. H. Holland, *Adaptation in Natural and Artificial Systems*. Ann Arbor, MI, USA: University of Michigan Press, 1975.
- [2] D. E. Goldberg, *Genetic Algorithms in Search, Optimization and Machine Learning*. Reading, MA, USA: Addison-Wesley, 1989.
- [3] J. Kennedy and R. Eberhart, "Particle swarm optimization," in *Proc. IEEE Int. Conf. Neural Networks*, Perth, Australia, 1995, pp. 1942–1948.
- [4] W. S. McCulloch and W. Pitts, "A logical calculus of the ideas immanent in nervous activity," *Bull. Math. Biophys.*, vol. 5, pp. 115–133, 1943.
- [5] F. Rosenblatt, "The perceptron: A probabilistic model for information storage and organization in the brain," *Psychol. Rev.*, vol. 65, no. 6, pp. 386–408, 1958.
- [6] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, "Learning representations by back-propagating errors," *Nature*, vol. 323, pp. 533–536, 1986.
- [7] S. Russell, *Human Compatible: Artificial Intelligence and the Problem of Control*. New York, NY, USA: Viking, 2019.
- [8] I. Gabriel, "Artificial intelligence, values, and alignment," *Minds and Machines*, vol. 30, no. 3, pp. 411–437, 2020.
- [9] L. Floridi and J. Cows, "A unified framework of five principles for AI in society," *Harvard Data Sci. Rev.*, vol. 1, no. 1, 2019.
- [10] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A fast and elitist multiobjective genetic algorithm: NSGA-II," *IEEE Trans. Evol. Comput.*, vol. 6, no. 2, pp. 182–197, 2002.
- [11] Y. Shi and R. Eberhart, "A modified particle swarm optimizer," in *Proc. IEEE Int. Conf. Evol. Comput.*, Anchorage, AK, USA, 1998, pp. 69–73.
- [12] R. Poli, J. Kennedy, and T. Blackwell, "Particle swarm optimization," *Swarm Intelligence*, vol. 1, no. 1, pp. 33–57, 2007.
- [13] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Las Vegas, NV, USA, 2016, pp. 770–778.
- [14] A. Vaswani et al., "Attention is all you need," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 30, 2017, pp. 5998–6008.
- [15] S. J. Pan and Q. Yang, "A survey on transfer learning," *IEEE Trans. Knowl. Data Eng.*, vol. 22, no. 10, pp. 1345–1359, 2010.
- [16] D. Amodei, C. Olah, J. Steinhardt, P. Christiano, J. Schulman, and D. Mané, "Concrete problems in AI safety," *arXiv:1606.06565*, 2016.
- [17] A. Jobin, M. Ienca, and E. Vayena, "The global landscape of AI ethics guidelines," *Nature Mach. Intell.*, vol. 1, no. 9, pp. 389–399, 2019.
- [18] L. Ouyang et al., "Training language models to follow instructions with human feedback," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 35, 2022, pp. 27730–27744.
- [19] Y. Bai et al., "Constitutional AI: Harmlessness from AI feedback," *arXiv:2212.08073*, 2022.
- [20] C. Dwork, M. Hardt, T. Pitassi, O. Reingold, and R. Zemel, "Fairness through awareness," in *Proc. 3rd Innov. Theor. Comput. Sci. Conf.*, Cambridge, MA, USA, 2012, pp. 214–226.
- [21] M. Hardt, E. Price, and N. Srebro, "Equality of opportunity in supervised learning," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 29, 2016, pp. 3315–3323.
- [22] N. Mehrabi, F. Morstatter, N. Saxena, K. Lerman, and A. Galstyan, "A survey on bias and fairness in machine learning," *ACM Comput. Surv.*, vol. 54, no. 6, art. 115, 2021.
- [23] F. Doshi-Velez and B. Kim, "Towards a rigorous science of interpretable machine learning," *arXiv:1702.08608*, 2017.
- [24] M. Mitchell et al., "Model cards for model reporting," in *Proc. Conf. Fairness, Accountability, Transparency*, Atlanta, GA, USA, 2019, pp. 220–229.
- [25] E. Easwaran, *Trans., The Bhagavad Gita*, 2nd ed. Tomales, CA, USA: Nilgiri Press, 2007.

- [26] R. K. Narayan, *The Ramayana: A Shortened Modern Prose Version of the Indian Epic*. New York, NY, USA: Viking, 1972.
- [27] T. G. Dietterich, "Ensemble methods in machine learning," in *Multiple Classifier Systems, Lecture Notes in Computer Science*, vol. 1857. Berlin, Germany: Springer, 2000, pp. 1–15.
- [28] L. Prechelt, "Early stopping—but when?," in *Neural Networks: Tricks of the Trade, Lecture Notes in Computer Science*, vol. 1524. Berlin, Germany: Springer, 1998, pp. 55–69.
- [29] Y. Bengio, J. Louradour, R. Collobert, and J. Weston, "Curriculum learning," in *Proc. 26th Int. Conf. Mach. Learn.*, Montreal, Canada, 2009, pp. 41–48.
- [30] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," arXiv:1503.02531, 2015.

Appendix A. Supplementary Material

Table A.1 lists the complete settings of the illustrative demonstration in Sections 4 and 5.

Table A.1. Settings of the computational demonstration.

Item	Setting
Data	Synthetic; N = 2,000; seed 42; 1,400 training and 600 test cases
Groups	Group 1 with probability 0.4; feature x_2 shifted by 0.9 for group 1
Label	1 if $1.2 \cdot x_1 + 0.8 \cdot x_2 + \varepsilon > 0.4$, with $\varepsilon \sim N(0, 0.8^2)$
Model	Logistic classifier, 3 parameters (two weights, one bias)
Objective	Binary cross-entropy + $\lambda \times$ mean probability gap between groups
λ values	0, 0.5, 1, 2, 4, 8
GA	Population 30; 80 generations; elitism 2; truncation selection (best 50%); uniform crossover; Gaussian mutation ($\sigma = 0.15$, per-gene rate 0.4)
PSO	30 particles; 80 iterations; $w = 0.72$; $c_1 = c_2 = 1.49$; initial positions $N(0, 1)$
Replication	10 runs per optimiser and λ (seeds 0–9); mean $\pm$ s.d. reported
Evaluation	Test accuracy; demographic-parity gap $ P(\hat{y}=1 \mid g=1) - P(\hat{y}=1 \mid g=0) $ at threshold 0.5