Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Ant colony optimization plus Temporal difference and opposition-based learning for semi-supervised multi-label feature selection

Sahar Shams Beyranvand
Lorestan University
Vahid Mehrdad
mehrdad.v@lu.ac.ir

Lorestan University
Mohammad Bagher Dowlatshahi
Lorestan University

Article

Keywords:

Posted Date: August 28th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10304282/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Version of Record: A version of this preprint was published at Scientific Reports on October 12th, 2026.
See the published version at https://doi.org/10.1038/s41598-026-75108-9.

Ant colony optimization plus Temporal difference and opposition-based learning for semi-supervised multi-label feature selection

Sahar Shams Beyranvand[1], Vahid Mehrdad[2,*], and Mohammad Bagher Dowlatshahi[3]

[1]Department of Electrical Engineering, Faculty of Engineering, Lorestan University, Khorramabad, Iran. Shams.s@fe.lu.ac.ir
[2]Department of Electrical Engineering, Faculty of Engineering, Lorestan University, Khorramabad, Iran. mehrdad.v@lu.ac.ir
[3]Department of Computer Engineering, Faculty of Engineering, Lorestan University, Khorramabad, Iran. dowlatshahi.mb@lu.ac.ir
[*]corresponding.author@email.example
[+]these authors contributed equally to this work

ABSTRACT

This paper presents a novel feature selection method for multi-label data, combining Ant Colony Optimization with Reinforcement Learning and Opposition-Based Learning in a semi-supervised framework. Ant colony optimization navigates the feature space, while reinforcement learning dynamically updates heuristic functions to predict rewards for selected features. Opposition-Based Learning enhances exploration by evaluating both original and complementary solutions. The semi-supervised learning aspect allows the algorithm to leverage the structure of the unlabeled data, improving the quality of feature selection even when labeled data is scarce. For the first time, the impact of the proposed algorithms has been systematically investigated. Compared to baseline methods, the approach achieves 7% and 3% higher Accuracy, 6% and 4% lower Hamming loss with 20% and 40% labeled data. Evaluated on 8 datasets using the Multi-Label K-Nearest Neighbors classifier, the method outperforms traditional techniques by refining feature subsets through unlabeled data. The method's higher computational complexity and memory requirements compared to baseline techniques make it better suited for high-accuracy tasks where real-time processing is not critical, such as precision-driven biomedical analysis or batch-oriented industrial data mining. This work establishes a new benchmark for semi-supervised multi-label feature selection, prioritizing accuracy in scenarios where computational latency is tolerable.

Introduction

The rapid growth of data complexity necessitates feature selection to reduce dimensionality, mitigate overfitting, and improve model efficiency[1–3]. This urgency is amplified in multi-label domains like genomics (gene-trait relationships), text mining (multi-topic document tagging), and image classification (scene recognition), where overlapping labels and high-dimensional data demand methods that balance relevance and interdependencies. However, multi-label data, where instances map to multiple overlapping categories[4,5], introduces challenges like label dependencies, imbalances, and sparsity. Traditional filter-based methods, which rank features using statistical metrics[3], often ignore label correlations, leading to redundant or irrelevant feature subsets[4,5]. For example, they may select features useful for individual labels but fail to capture cross-label interactions[6]. Wrapper-based methods, though capable of modeling dependencies via embedded algorithms[7], face prohibitive computational costs and overfitting risks in multi-label settings due to the need to retrain models across all labels[6,8]. Hybrid approaches struggle to adapt to dynamic label relationships and scale with label growth, while imbalance issues skew feature selection toward dominant labels, neglecting underrepresented ones[4–6]. These limitations highlight the need for methods that explicitly address label interdependencies without sacrificing efficiency[6].

Ant Colony Optimization (ACO), inspired by ant foraging behavior, is a metaheuristic algorithm for solving combinatorial problems[9]. Artificial ants explore solutions guided by pheromone trails, enabling efficient exploration and exploitation of promising regions[10]. ACO has been applied to challenges like routing, scheduling, and optimization with notable success[11]. Opposition-Based Learning (OBL) enhances optimization by considering both original and opposite states of solutions[12]. This approach accelerates convergence and promotes diverse solution exploration by reflecting or negating solution components[13]. OBL's versatility improves accuracy and exploration in optimization tasks[14].

Reinforcement Learning (RL) enables agents to learn sequential decision-making by interacting with environments and

maximizing rewards[15]. Unlike supervised learning, RL relies on feedback from actions to guide learning. RL has demonstrated success in robotics, gaming, and autonomous systems, handling complex decision tasks through continuous adaptation[2]. Semi-supervised learning combines limited labeled data with abundant unlabeled data, bridging the gap between supervised and unsupervised techniques[16]. By leveraging both, it enhances model accuracy and generalization while reducing dependency on labeled datasets. Algorithms like Label Propagation and Self-Training illustrate its effectiveness in diverse applications[17].

ACO is a versatile metaheuristic algorithm inspired by ant foraging behavior, widely used for feature selection in multi-label classification tasks. However, traditional ACO faces inherent limitations in this context. Specifically, its reliance on random initial placement of ants within the feature graph structure often traps solutions in local optima, particularly in high-dimensional multi-label data. Furthermore, decision-making processes in ACO depend heavily on pheromone updates and a single heuristic element, which can lead to suboptimal exploration, increased error rates, and redundancy in selected features.

To address these challenges, we integrate TD learning and OBL into the ACO framework. The incorporation of TD learning introduces a dynamic decision-making layer that complements pheromone updates. By enabling ants to interact with the environment and learn optimal feature selection strategies through reward maximization, TD learning reduces reliance on static heuristics, enhances adaptability, and mitigates errors caused by premature convergence. Simultaneously, OBL broadens the algorithm's exploration capabilities by evaluating both candidate solutions and their opposites, effectively diversifying the search space and escaping local optima. This dual mechanism of exploration and adaptive exploitation ensures a balanced traversal of the feature graph, prioritizing globally informative features while minimizing redundancy. In semi-supervised learning scenarios with limited labeled data, this hybrid approach demonstrates significant robustness. By coupling OBL's exploratory efficiency with RL's reward-driven refinement, the algorithm achieves faster convergence and greater stability, even when only 20–40% of training data is labeled. Empirical evaluations confirm that the synergy of these components not only overcomes ACO's traditional weaknesses but also improves multi-label classification accuracy and generalization, particularly in complex, high-dimensional datasets where feature-label interactions are intricate and non-linear.

The proposed method includes the following important points:

- For the first time, the effect of semi-supervised multi-label feature selection is analyzed in an algorithm that uses OBL for pheromone updates in the ACO algorithm.
- The effect of semi-supervised multi-label feature selection is explored in an ACO algorithm that utilizes RL as the heuristic function and incorporates OBL.
- Innovation lies in examining the impact of semi-supervised multi-label feature selection in an algorithm that concurrently uses OBL for both pheromone updates and the heuristic function.
- As far as we know, this is the first time the impact of semi-supervised has been studied in the Ant-TD algorithm.

We evaluated the accuracy and Hamming loss for each scenario. With 20% and 40% labeled data, our semi-supervised multi-label feature selection approach, combined with OBL, demonstrated notable improvements in accuracy and reductions in Hamming loss compared to traditional methods.

Related work

The integration of Ant Colony Optimization (ACO) and related methods has significantly advanced the field of multi-label feature selection. Researchers have developed innovative techniques that leverage the problem-solving abilities of artificial ants and computational strategies to improve feature selection algorithms. For example, MLACO[18] uses ACO to identify the most informative feature subsets, optimizing multi-label classification tasks by thoroughly exploring the feature space. Similarly, MFS-MCDM[19] incorporates multi-criteria decision-making (MCDM) to balance various selection criteria and enhance model performance. Another ensemble approach[20] combines multiple feature selection algorithms within an MCDM framework to improve robustness and efficiency.

Semi-supervised feature selection has also benefited from ACO-based methodologies. SemiACO[21] addresses high-dimensional data challenges by selecting minimally redundant and highly relevant features using a nonlinear heuristic function and TD reinforcement learning (RL). This approach models feature selection as a Markov Decision Process (MDP), enabling dynamic heuristic learning. Another study[22] demonstrates the effectiveness of ACO in classification tasks by utilizing pheromone-based communication to analyze both labeled and unlabeled data, improving classification accuracy. Expanding this concept,[23] proposes an ACO-based semi-supervised learning method for generating classification rules, combining ACO's exploratory capabilities with labeled and unlabeled data for enhanced accuracy.

Further innovations include Ant-TD[24], which merges ACO with TD reinforcement learning for multi-label feature selection. Unlike traditional ACO methods with static heuristics, Ant-TD employs a dynamic strategy, refining heuristic functions based on prior experiences and modeling the problem as an MDP.[25] proposes an Enhanced Ant Colony System (ACS), an AI-driven

optimization algorithm, to classify water quality rules in smart aquaculture. The AI model improves decision-making by efficiently analyzing sensor data to optimize water parameters (e.g., pH, temperature) and enhance aquaculture management compared to conventional methods, alongside future research directions like hybrid approaches and deep learning integration. Building on ensemble methods,[26] presents a framework combining multiple heuristics within an MCDM structure to improve feature selection's robustness and effectiveness.

Several enhancements in ACO-related techniques are also noteworthy. Modifications proposed in[27] focus on synchronizing ant paths and incorporating opposite-pheromone updates to improve the solution construction phase of ACS. MOBL[28] employs a multi-operator and neighborhood structure to enhance exploration and exploitation in numerical optimization. Similarly, OOBL[29] introduces orthogonal opposition to refine optimization processes. An advanced Elephant Herding Optimization algorithm with OBL and K-means clustering[30] improves convergence and solution quality, while OPACO[31] applies ACO principles to nonlinear structural analysis, offering potential applications in engineering optimization. These advancements demonstrate the versatility of ACO and its integration with various computational techniques, contributing to robust feature selection, optimization, and classification methodologies across diverse domains.

Despite these developments, no existing literature has combined ACO, TD, and OBL. This article introduces a novel approach that simultaneously employs multi-stage and heuristic opposition along with RL, aiming to push the boundaries of multi-label feature selection

Basic concepts

Ant Colony Optimization (ACO)

Ant Colony Optimization (ACO), introduced by Marco Dorigo in the 1990s, is a metaheuristic inspired by ants' foraging behavior, particularly their ability to find the shortest paths to food sources. It is widely used to solve combinatorial optimization problems like the Traveling Salesman Problem (TSP), Vehicle Routing Problem (VRP), and Quadratic Assignment Problem (QAP)[32]. In ACO, artificial ants construct solutions by probabilistically selecting components based on pheromone trails and heuristic information[33]. Pheromone trails, representing indirect communication, indicate the quality of previous solutions, while heuristic information provides problem-specific insights, such as distances in the TSP. Ants deposit pheromones based on solution quality, reinforcing successful paths, while evaporation prevents stagnation and promotes exploration by reducing the influence of older solutions. This iterative process continues until a termination criterion is met, such as a maximum number of iterations or achieving a desired solution quality. Pheromone updates guide the search towards promising areas and involve two main mechanisms: global update and local update.

Global Pheromone Update: After all ants have constructed solutions in an iteration, the global update is performed to update the pheromone trails based on the quality of the solutions found. The amount of pheromone deposited on an edge is typically proportional to the quality of the solution that includes that edge. Better solutions contribute more pheromone to the edges they traverse. The global pheromone update can be calculated as follows:

$$\tau_{ij}^{(t+1)} = (1-\rho)\,\tau_{ij}^{(t)} + \sum_{k=1}^{N} \Delta\tau_{ij}^{k} \tag{1}$$

where:

- $\tau_{ij}^{(t)}$ is the pheromone level on edge (i, j) at iteration t,
- ρ is the evaporation rate, representing the amount of pheromone that evaporates from the trails,
- τ_{ij}^{k} is the amount of pheromone deposited on edge (i, j) by ant k, and
- N is the total number of ants[34].

Local Pheromone Update: In addition to the global update, a local pheromone update is often applied to intensify the search near the edges that have been recently traversed by ants. The local update involves depositing a small amount of pheromone on the edges immediately after an ant traverses them, regardless of the solution's quality[35]. The local pheromone update can be calculated as follows:

$$\tau_{ij}^{(t+1)} = (1-\alpha)\tau_{ij}^{(t)} + \alpha \cdot \tau_0 \tag{2}$$

where:

- α is the local update factor, determining the amount of pheromone deposited locally,

- τ_0 is a constant representing the initial pheromone level on edges, and
- $\tau_{ij}^{(t)}$ is the pheromone level on edge (i, j) before the local update[36].

Opposition-Based Learning (OBL)

OBL is a computational intelligence technique that enhances optimization algorithms by considering both the current and opposite states of solutions during the search process. Inspired by the concept of opposition in nature, OBL aims to improve the convergence speed and solution quality by exploring a wider solution space more effectively[37]. This technique can be integrated into various optimization algorithms to enhance their performance. The fundamental idea behind OBL is to generate and evaluate opposite solutions alongside the current solutions. The opposition is defined in a way that reflects the current solution around a predefined reference point. This concept leverages the idea that the opposite of a potentially good solution might also be a good or even better solution. By considering both the original and opposite solutions, OBL increases the likelihood of finding optimal solutions more quickly[38].

The absolute value opposition (ABS) operator is a concept used in OBL to generate an opposite solution by taking the absolute value of each component of the current solution. Mathematically, if we have a current solution represented as $X = (x_1, x_2, \ldots, x_n)$, the ABS operator can be defined as:

$$X' = (|x_1|, |x_2|, \ldots, |x_n|). \tag{3}$$

where X' represents the opposite solution generated using the absolute value opposition operator.

- For each component x_i of the current solution X, take the absolute value $|x_i|$.
- Assign the value $|x_i|$ to the corresponding component x_i' of the opposite solution X'.

This equation reflects the current solution around the midpoint of the bounds, ensuring the opposite solution remains within the same search space. Both solutions are evaluated using the optimization problem's objective function, with the better one selected for further processing[39].

Temporal Difference Learning (TD)

TD Learning is a reinforcement learning technique used for predicting future rewards and learning optimal policies based on experience. It estimates the value of states or actions in a sequential decision-making process, with applications in areas such as game playing, robotics, and financial forecasting. TD Learning updates value estimates after each step, enabling real-time learning. The value function represents the expected cumulative reward of being in a specific state or taking a particular action. Two primary value functions are State Value Function ($V(s)$) and Action Value Function ($Q(s,a)$)[40].

The core of TD Learning is the update rule, which adjusts the value function based on the difference between the predicted and actual rewards. This difference, known as the TD error, is used to update the value function incrementally. The TD update rule for the state value function can be expressed as follows:

$$\delta = r + \gamma V(s') - V(s) \tag{4}$$

where:

- δ is the TD error,
- r is the immediate reward obtained after transitioning from state s to state s',
- γ is the discount factor, which determines the importance of future rewards,
- $V(s)$ and $V(s')$ are the estimated values of states s and s', respectively.

TD(0) is a basic form of TD Learning that updates value estimates using the immediate reward and the next state's value.

Semi-Supervised Learning (SSL)

Semi-Supervised Learning (SSL) combines labeled and unlabeled data to enhance model performance, addressing situations where labeled data is limited and costly, while unlabeled data is abundant and inexpensive. By leveraging the structure of unlabeled data, SSL improves classification accuracy, reduces overfitting, and enhances generalization. It builds predictive models using a small labeled dataset and a larger unlabeled dataset, integrating insights from both supervised and unsupervised learning[41].

Key SSL methods include:

- **Self-Training:** Models trained on labeled data generate pseudo-labels for unlabeled data, which are iteratively added to the training set to improve performance.
- **Co-Training:** Multiple models trained on distinct feature subsets label unlabeled data, sharing new labels to enhance each other's learning.
- **Multi-View Learning:** Extends co-training by combining complementary insights from different feature sets to improve predictions.
- **Graph-Based Methods:** Utilize graph representations of data, propagating labels from labeled nodes to similar unlabeled nodes, based on assumptions such as smoothness.
- **Generative Models:** Methods such as VAEs and GANs generate synthetic data to augment labeled datasets, improving learning.

SSL enhances cost-effectiveness by reducing reliance on labeled datasets while improving model accuracy and generalization. It is particularly valuable for imbalanced or limited labeled data, helping balance the learning process and improving performance on underrepresented classes[42].

Mutual Information

Mutual Information (MI) is a fundamental concept in information theory and statistics that measures the amount of information obtained about one random variable through another. It quantifies the degree of dependency between two variables and is used to understand the amount of shared information between them. Mutual Information between two random variables X and Y, denoted as $I(X;Y)$, is defined as the difference between the joint entropy of X and Y and the sum of the marginal entropies of X and Y. Formally, it is given by:

$$I(X;Y) = \sum_{x \in X} \sum_{y \in Y} p(x,y) \log \left(\frac{p(x,y)}{p(x)p(y)} \right) \tag{5}$$

where:

- $p(x,y)$ is the joint probability mass function of X and Y, representing the probability that $X = x$ and $Y = y$,
- $p(x)$ and $p(y)$ are the marginal probability mass functions of X and Y, respectively,
- the summations are taken over all possible values of X and Y.

Mutual Information quantifies how much knowing one variable reduces uncertainty about the other. Unlike correlation, which only measures linear relationships, MI captures both linear and non-linear dependencies. A higher mutual information value indicates a stronger relationship between the variables[43].

Cosine Similarity

Cosine similarity is a metric used to measure the similarity between two non-zero vectors in an inner product space. It is commonly used in text analysis, information retrieval, and various machine learning applications to assess how similar two vectors are, irrespective of their magnitude. Cosine similarity is defined as the cosine of the angle between two vectors. Mathematically, for two vectors $\mathbf{A}$ and $\mathbf{B}$, Cosine(A,B) is given by:

$$\cos(A,B) = \frac{\mathbf{A} \cdot \mathbf{B}}{\|\mathbf{A}\| \|\mathbf{B}\|} \tag{6}$$

where:

- $\cdot$ represents the dot product between two vectors,
- $\|\cdot\|$ represents the Euclidean norm (magnitude) of a vector.

$\mathbf{A} \cdot \mathbf{B}$ measures the alignment between two vectors in space, while the denominator $\|\mathbf{A}\| \|\mathbf{B}\|$ normalizes the result, turning cosine similarity into a bounded similarity metric[44].

Ant-TD Algorithm

The Ant-TD algorithm enhances traditional ACO by integrating it with TD Reinforcement Learning (RL), leveraging RL's dynamic learning capabilities to improve solution quality in complex optimization problems. Unlike static heuristics in standard ACO, Ant-TD employs a TD agent that models the optimization problem as a Markov Decision Process (MDP).

Key contributions of RL in Ant-TD:

- **Dynamic Heuristic Learning:** The TD agent refines heuristic functions by updating state or action values using the TD error, reflecting the difference between predicted and actual rewards.
- **Adaptive Feedback:** During solution construction, the TD agent provides real-time feedback on the quality of ants' decisions, guiding them toward better paths.
- **Influence on Pheromone Update:** RL-adjusted heuristic values modify pheromone trails, reinforcing high-quality solutions and improving exploration–exploitation balance.

This integration enables Ant-TD to dynamically adapt to problem landscapes, enhancing ACO's performance in solving complex and dynamic optimization tasks[45].

Proposed method

In this study, we introduce a novel feature selection approach for multi-label classification that combines ACO with TD RL, OBL, and a semi-supervised learning framework. Prior to applying this method, a preprocessing phase is conducted where rows with missing data or incomplete labels are pruned from the dataset. The cleaned dataset is then partitioned into 60% training and 40% test sets to ensure robust evaluation while maintaining sufficient data for both learning and validation phases. This method is designed to enhance both the efficiency and accuracy of feature selection, particularly in scenarios where labeled data is limited.

To achieve this, we first utilize ACO to navigate the feature space, allowing artificial ants to construct candidate feature subsets based on probabilistic selection guided by pheromone trails and heuristic information. The exploration process is further refined by integrating TD Reinforcement Learning (RL), which dynamically updates the heuristic function. Specifically, the TD agent learns and adjusts the state value V, which estimates the potential reward of selecting specific feature subsets. This adaptive learning process helps the algorithm make more informed decisions during the feature selection.

In addition to these enhancements, we incorporate Opposition-Based Learning (OBL) to expand the search space. By considering both original and opposite solutions, OBL helps the algorithm avoid premature convergence and improves its ability to discover high-quality solutions. The application of OBL is twofold: it is used in updating pheromone levels and in adjusting the heuristic value V, ensuring a thorough exploration of the solution space. The most significant innovation in our method is the integration of semi-supervised learning. To evaluate its impact, we conduct experiments with two different proportions of labeled data—20% and 40%—while the remaining data is left unlabeled. By leveraging the structure of the unlabeled data, the algorithm can enhance its feature selection capabilities even when the amount of labeled data is restricted.

Algorithm 1 outlines the process of integrating OBL into the pheromone update phase of the ACO algorithm. In traditional ACO, ants are randomly placed on the data graph. However, in this modified approach, Mutual Information (MI) between labeled data is leveraged to guide ants to start from more statistically meaningful nodes, ensuring stronger initialization. The implementation employs Equations 5 and (3) to modify pheromone levels based on labeled data. By considering both the original and opposite solutions during pheromone updates, this approach enhances the exploration capability of the algorithm. The use of opposition helps to prevent premature convergence by encouraging diverse exploration paths in the search space, leading to the identification of more optimal feature subsets.

Algorithm 1 Using OBL for pheromone update

```
function OPPOSITION BASED INITIALIZATION STATE VALUE(data)
    Compute MI matrix (Eq. 5) based on all data
    Find min and max values for central point calculation
    Compute opposition matrix (Eq. 3)
    Calculate column-wise mean of opposition matrix
    Initialize pheromone ph_0
end function
```

Algorithm 2 details the application of OBL in the update process of the heuristic vector V. The matrix V is initialized using cosine similarity between data points, leveraging relationships across the entire dataset to establish a robust starting point. This

contrasts with simplistic initialization strategies, ensuring the algorithm begins with a statistically meaningful representation of feature interactions. This vector, which guides the ants' decision-making, is updated using Equations 3 and 6. In addition to leveraging labeled data, the semi-supervised learning component is incorporated into the process by using unlabeled data to update the heuristic values. This integration ensures that the algorithm can utilize both labeled and unlabeled data to refine the feature selection process, improving the overall efficiency and accuracy of the model.

Algorithm 2 Using OBL for update V

function OPPOSITION BASED INITIALIZATION STATE VALUE(data)
    Compute the mutual information matrix using Eq. (5)
    Determine the minimum and maximum values of each feature
    Compute the opposition matrix using Eq. (3)
    Compute the column-wise mean of the opposition matrix
    Initialize the pheromone matrix τ_0
end function

Figure 1 provides a detailed flowchart of the proposed method, visually encapsulating the integration of ACO, OBL, and RL components. The workflow begins with the initialization of the ACO population and pheromone trails, followed by the generation of candidate feature subsets. Key steps include the application of OBL to compute opposition solutions, which diversify the search space and mitigate local optima. Reinforcement Learning (RL) dynamically adjusts pheromone updates and heuristic priorities based on environmental feedback, ensuring adaptive exploration-exploitation balance. The semi-supervised learning module processes both labeled and unlabeled data to refine feature relevance scores iteratively This flowchart serves as a conceptual roadmap, clarifying the synergy between computational techniques and their sequential interactions to achieve robust multi-label feature selection.

Algorithm 3 presents the complete ACO algorithm, integrating the procedures from Algorithms 1 and 2. This comprehensive approach combines the opposition-based pheromone updates and heuristic vector V updates to form a robust feature selection framework. The algorithm capitalizes on the strengths of both OBL and semi-supervised learning, effectively navigating the feature space and improving the selection of relevant features. The simultaneous use of both techniques ensures a thorough exploration and exploitation of the solution space, leading to superior performance in multi-label feature selection tasks.

Algorithm 3 Using OBL & RL in ACO for semi-supervised data

while termination condition is not met **do**
    for each ant in the population **do**
        Initialize the ant's path
        while the ant has not completed its tour **do**
            if $\lambda > \lambda_0$ **then**
                Perform opposition-based initialization for pheromone (Algorithm 1)
                Perform opposition-based initialization for state value vector (Algorithm 2)
            else
                Select the next feature to visit based on pheromone levels and heuristic information (Eq. 2)
                Update the ant's path and pheromone trail
            end if
        end while
        if ant's solution is better than global best solution **then**
            global best solution $\leftarrow$ ant's solution (Eq. 1)
        end if
    end for
    for each edge (i, j) in the solution paths **do**
        pheromoneMatrix$[i][j] \leftarrow (1 - \text{evaporationRate}) \cdot \text{pheromoneMatrix}[i][j] + \text{quality of ant solution}$
    end for
end while

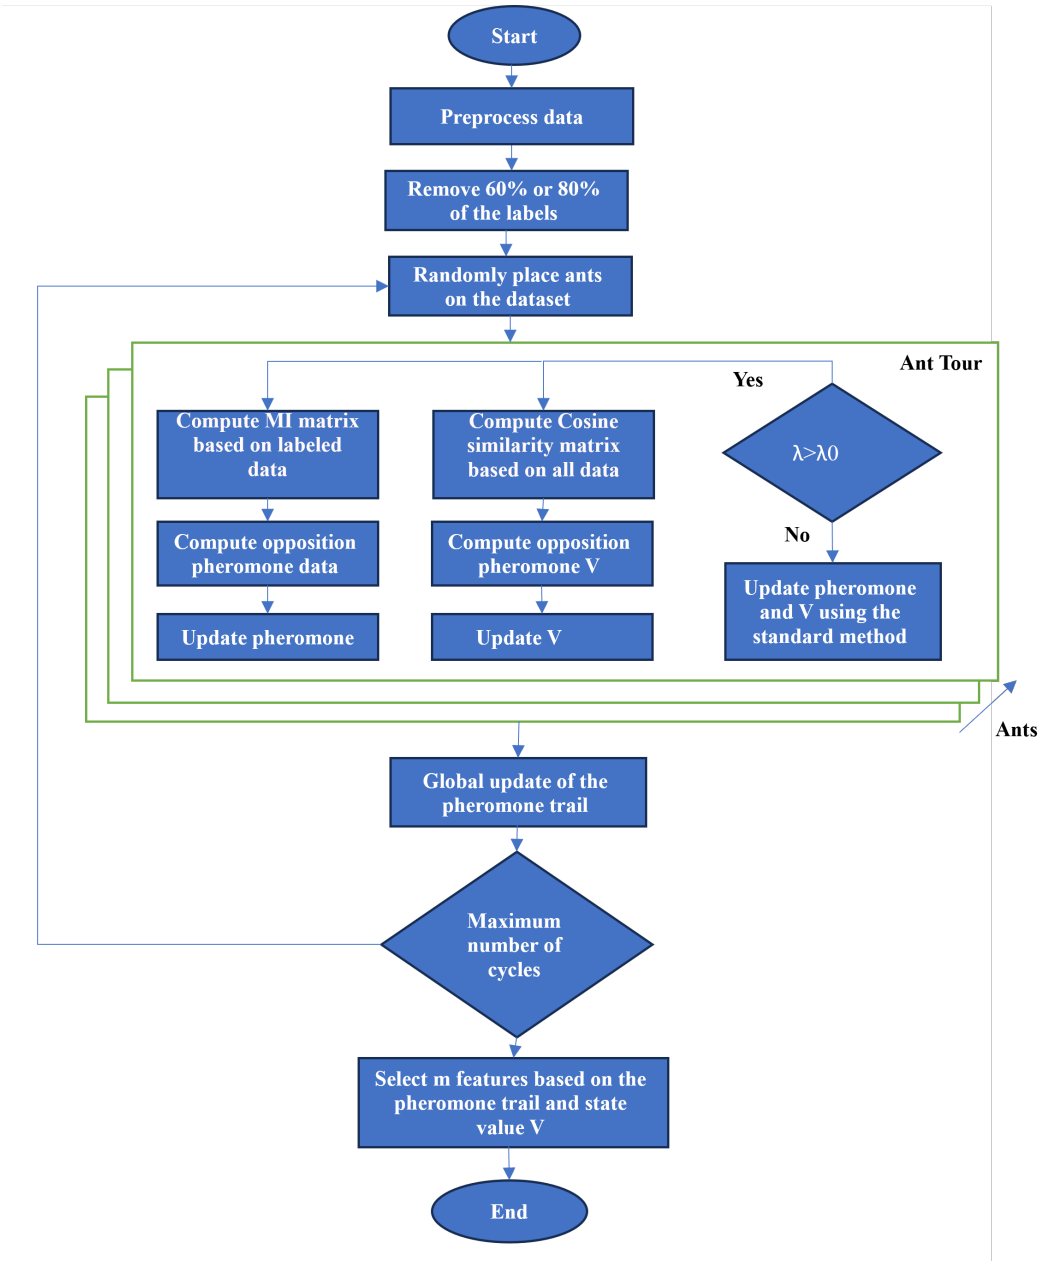


Figure 1. The general flowchart.

Experiment Setup

Datasets

The hypothesis was tested on eight multi-label datasets: Birds, Scene, Arts, Science, Chess, Cooking, Image, and Language. These datasets [1] cover various types of multi-label data[46,47]. The Language dataset is text-based, the Image dataset consists of images, and the other datasets belong to different categories.

Table 1 summarizes the characteristics of the datasets used in this study.

Table 1. Datasets and their characteristics

Name	Features	Labels	Instances	Label Cardinality	Sparsity/Density	Domain
Birds	258	21	645	166	1.0000	Text
Scene	294	6	2407	15	1.0000	Text
Arts	262	26	5000	462	1.0000	Text
Science	243	40	5000	398	1.0000	Text
Chess	812	227	1675	1078	0.0106	Text
Cooking	977	400	10491	6386	0.0056	Text
Image	294	5	2000	20	1.0000	Image
Language	1004	75	1459	304	1.0000	Text

To maintain reproducibility and fairness, the preprocessed datasets were partitioned using stratified random sampling into 40% test and 60% training subsets, preserving the original label distribution across splits. Within the training subset, semi-supervised learning conditions were simulated by systematically masking labels: 80% of labels were removed in the first scenario, and 60% were masked in the second scenario. This design allowed for a rigorous evaluation of the algorithm's robustness under varying levels of label scarcity.

All datasets were sourced from established benchmarks in references[46,47], ensuring alignment with prior research and facilitating direct comparisons. The structured approach to preprocessing, partitioning, and label masking underscores the methodological rigor of the evaluation, enabling precise insights into the algorithm's performance in multi-label feature selection tasks.

Parameter Setting

To predict the relevant labels using the implicit patterns among the data, we employed the well-known and widely used multi-label classifier, Multi-Label K-Nearest Neighbors (MLKNN)[48]. This classifier utilizes the K-nearest neighbors method specifically designed for multi-label data, predicting the label set of an unseen instance based on information obtained from the neighborhood of the instance. Key elements used in this algorithm include membership count statistics and the application of the Maximum A Posteriori (MAP) principle. In our study, we set the neighborhood size to $K = 10$, providing a balance between local and global information for label prediction.

Our approach incorporates semi-supervised learning techniques to enhance the performance of the ACO and RL algorithms, particularly in scenarios with different proportions of labeled data. Semi-supervised learning leverages both labeled and unlabeled data, making it especially useful when the available labeled data is limited. In our experiments, we tested two scenarios with varying proportions of labeled data: 20% and 40%. This approach allows the algorithm to utilize the structure of the unlabeled data to improve feature selection and ultimately achieve better classification performance.

The optimized values for the base algorithm parameters were selected based on the methodological choices established in[24], ensuring alignment with proven ACO and RL frameworks. The RL and ACO algorithms include several parameters that need to be fine-tuned to achieve optimal performance. Through extensive experimentation, we determined the most appropriate values for these parameters, which are listed in Table 2 for the datasets used in this study. For the ACO algorithm, critical parameters include the number of ants ($nAnt$), the number of graph traversal cycles each ant performs ($nCycle$), the number of features each ant can observe in each cycle (NF), and the pheromone evaporation rate per move (ρ). These parameters influence the exploration and exploitation balance within the ACO framework.

In the heuristic function and RL component, essential parameters include the learning rate (α), the discount rate (γ), and the parameter β. These values govern the learning dynamics and the update process within the RL framework. The parameter m represents the value of the final top features selected, reflecting the most relevant features identified through the semi-supervised learning process. Additionally, for the OBL function, we used the value of λ_0, which determines the degree of opposition incorporated into the learning process.

[1] https://github.com/mrapp-ke/Boomer-Datasets

To evaluate the proposed algorithm against other feature selection methods for multi-label data, we employed several evaluation metrics, including accuracy and Hamming loss. These metrics are example-based, meaning they assess each test instance individually before calculating the average for the entire test set. Accuracy measures the proportion of correctly predicted labels out of the total predicted and actual labels for each instance, following the specified relationship.

Table 2. Parameter settings

Parameter	Description	Value
α	Learning rate (TD(0))	0.5
ρ	Pheromone evaporation rate	0.2
γ	Discount factor (TD(0))	0.8
$nCycle$	Number of iterations	25
$nAnt$	Number of ants	5
NF	Features per ant	$\frac{1}{8}d \leq NF \leq \frac{1}{6}d$
q	Exploration-exploitation coefficient	0.7
β	Pheromone-heuristic trade-off	1
m	Number of selected features	$10 \leq m \leq 100$
λ_0	OBL activation threshold	0.001

Hamming loss, on the other hand, evaluates the fraction of incorrectly assigned labels, offering insight into the algorithm's ability to reduce prediction errors across multiple labels. By combining these metrics, we can comprehensively compare the performance of different feature selection methods, highlighting their effectiveness in handling multi-label data and their ability to identify relevant features accurately.

Accuracy in multi-label classification measures the proportion of correctly predicted labels to the total number of labels. It considers both true positives (correctly predicted labels) and true negatives (labels correctly not predicted). The formula for accuracy is given by:

$$Accuracy = \frac{1}{n}\sum_{i=1}^{N}\frac{|\hat{Y}_i \cap Y_i|}{|\hat{Y}_i \cup Y_i|} \tag{7}$$

where:

- $|\hat{Y}_i \cap Y_i|$ is the number of correctly predicted labels for the i-th instance.
- $|\hat{Y}_i \cup Y_i|$ is the total number of unique labels in both the predicted and true sets for the i-th instance[49].

Hamming loss in multi-label classification quantifies the fraction of labels that are incorrectly predicted. It takes into account both false positives (labels incorrectly predicted) and false negatives (labels not predicted but present in the true set). The formula for Hamming loss is given by:

$$Hamming\ Loss = \frac{1}{N \cdot L}\sum_{i=1}^{N}\sum_{j=1}^{L}\mathbf{1}(x_{ij} \neq y_{ij}) \tag{8}$$

where:

- N is the number of instances.
- L is the total number of labels.
- x_{ij} is the predicted label for the j-th label of the i-th instance.
- y_{ij} is the true label for the j-th label of the i-th instance.
- $\mathbf{1}$ is the indicator function, which equals 1 if $x_{ij} \neq y_{ij}$ and 0 otherwise[50].

These formulas provide a comprehensive evaluation of the algorithm's performance, balancing both correct predictions and errors in multi-label classification[50].

Results

This section presents a comparative analysis of the proposed semi-supervised method against benchmark feature selection algorithms, including Cluster[51], PPT-MI[52], LRFS[53], MCLS[54], MDFS[55], MSSL[56], and Ant-TD[24]. All baseline methods are implemented according to their original papers with recommended parameter settings.

Each method is evaluated by selecting the top-100 features and measuring classification performance using Accuracy and Hamming Loss.

The proposed method is tested under two labeling conditions (20% and 40%) to evaluate robustness under limited supervision. Results demonstrate that the proposed approach consistently outperforms competing methods, particularly under low-label conditions.

For clarity, the following abbreviations are used:

- **op-ph**: Improved ACO with RL using opposition-based pheromone update
- **op-v**: Improved ACO with RL using RE in V updates
- **op-ph-v**: Improved ACO with RL using both pheromone and V opposition updates

Figure 2 illustrates a detailed comparison of accuracy between our proposed semi-supervised multi-label feature selection algorithms and a range of established methods when using 20% labeled data. The results highlight the superior performance of our algorithms, which effectively utilize both OBL and semi-supervised techniques. With only 20% of the data labeled, our approach leverages the semi-supervised context to enhance feature selection by providing a more nuanced understanding of the data.

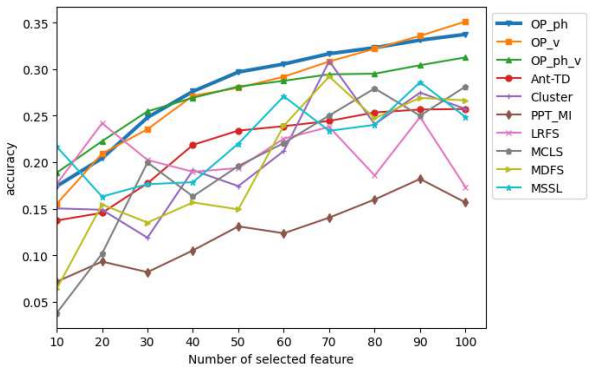
In this scenario, the algorithm strategically incorporates less prominent data points, which might initially be overlooked due to the ant's movement choices or pheromone evaporation. These points are re-evaluated through semi-supervised learning, allowing the algorithm to potentially identify and reinforce relevant features. If these paths prove useful, they are reinforced by pheromone deposition, increasing their chances of selection in future iterations. Conversely, less effective paths are excluded, refining the search focus and improving the overall accuracy. This approach helps reduce Hamming loss and enhances the feature selection process compared to traditional methods.

Figure 3 provides a comparative analysis of Hamming loss for our semi-supervised multi-label feature selection algorithms against established methods, using 20% labeled data. The results reveal that our proposed algorithms significantly reduce Hamming loss compared to traditional approaches. With only 20% of the data labeled, our algorithms effectively use semi-supervised learning to improve the accuracy of feature selection by refining how irrelevant features are handled.

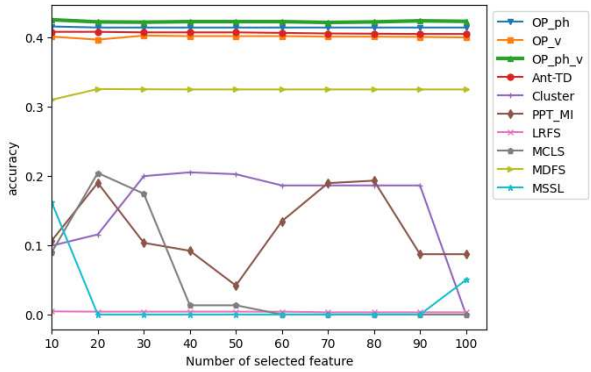
In this scenario, semi-supervised learning allows the algorithm to better utilize unlabeled data, making more informed decisions about which features to include. By incorporating OBL, our algorithms are able to strategically focus on paths that have the potential to select relevant features, while minimizing the impact of less useful data points. This leads to a reduction in Hamming loss, as the algorithm is more adept at avoiding incorrect label assignments and improving the overall quality of feature selection.

Figure 4 presents a comparative analysis of accuracy for our semi-supervised multi-label feature selection algorithms using 40% labeled data, demonstrating their enhanced performance relative to several established methods. The higher proportion of labeled data provides a more robust framework for our algorithms, enabling them to leverage OBL more effectively.

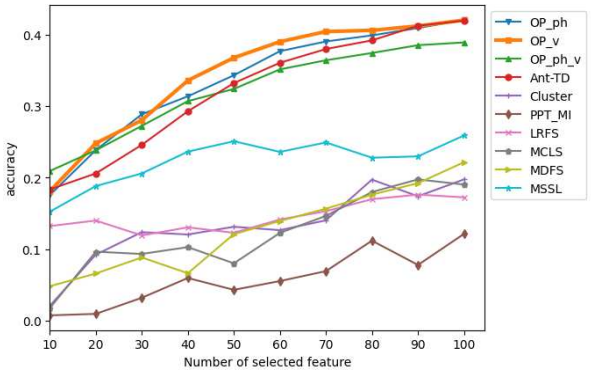
In this setup, the algorithm makes use of a greater amount of labeled data to make more informed feature selection decisions. Points that might initially appear less significant due to ant movement or pheromone evaporation are more thoroughly evaluated through semi-supervised learning. If these points lead to the identification of relevant features, they are reinforced through pheromone deposition, thereby improving their likelihood of selection. Paths that do not contribute positively are excluded, ensuring a more focused and effective search process. This approach significantly reduces Hamming loss and improves accuracy, showing a clear advantage over traditional feature selection methods.



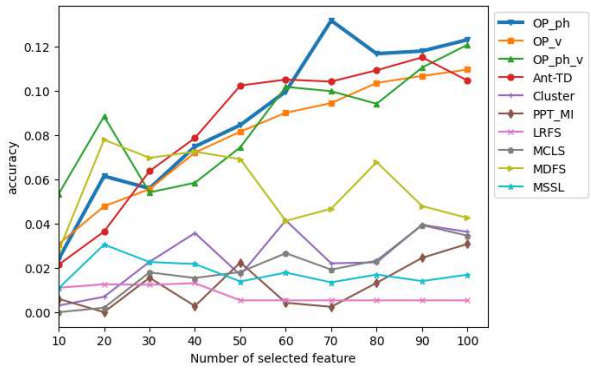
(a) Image



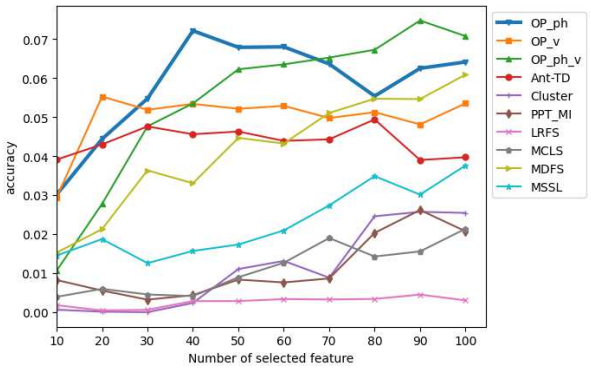
(b) Birds



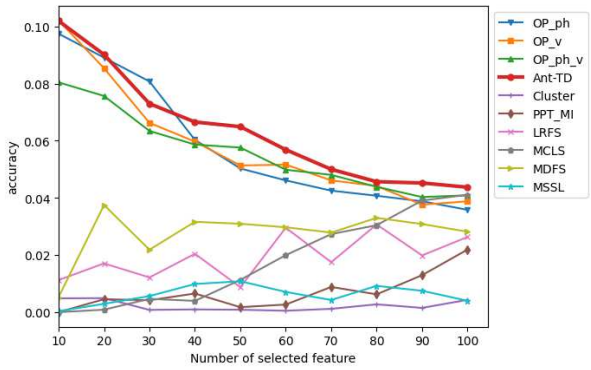
(c) Scene



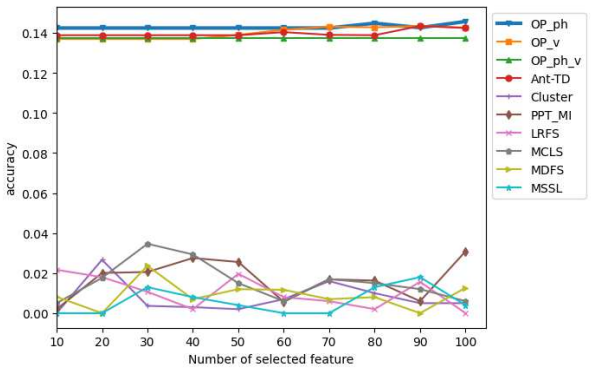
(d) Arts



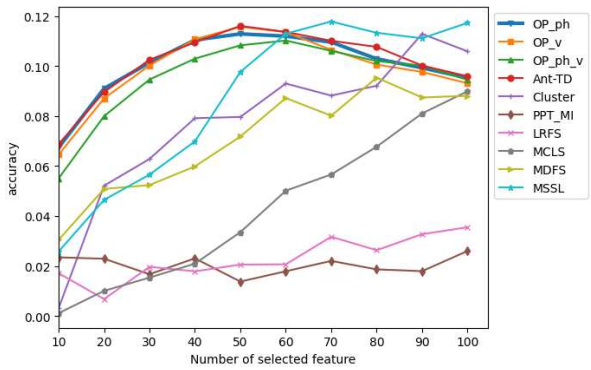
(e) Science



(f) Chess

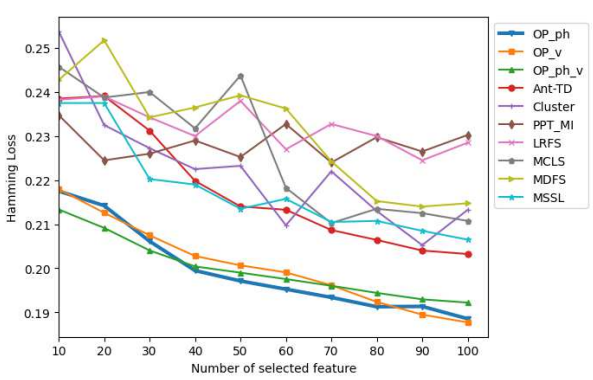


(g) Language

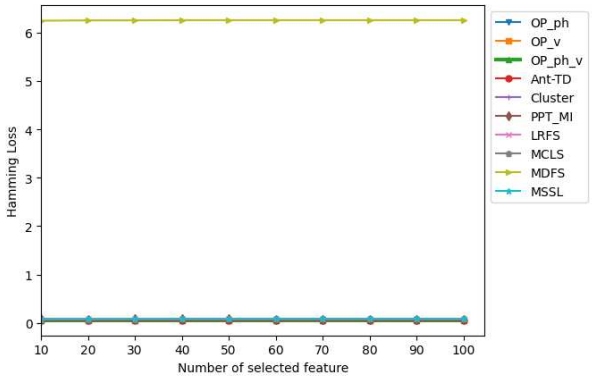


(h) Cooking

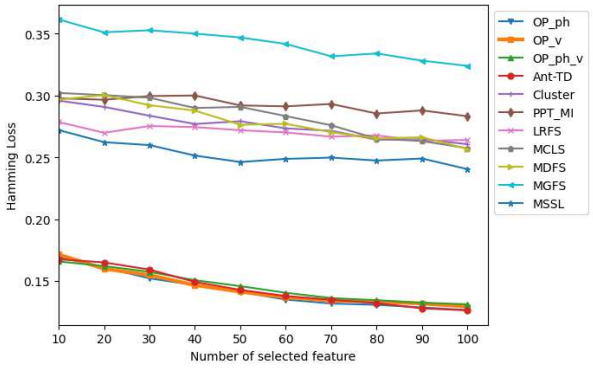
Figure 2. Average performance in terms of Accuracy for 20% of labeled data ↑



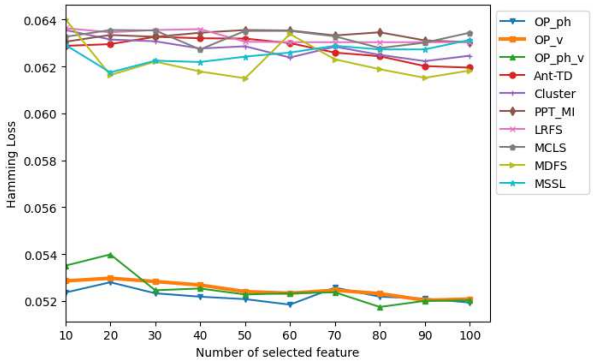
(a) Image



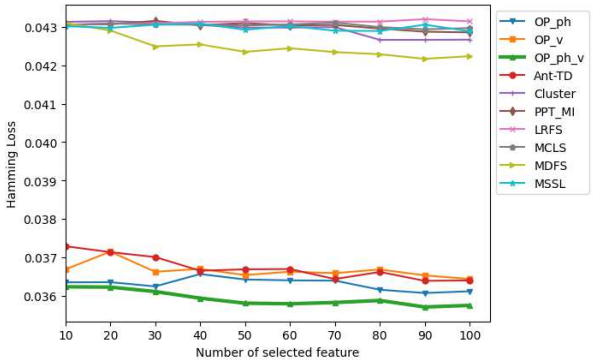
(b) Birds



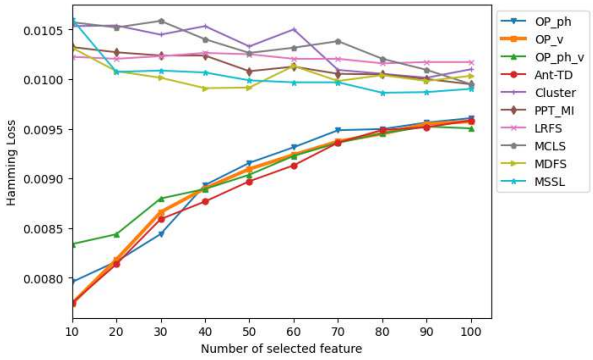
(c) Scene



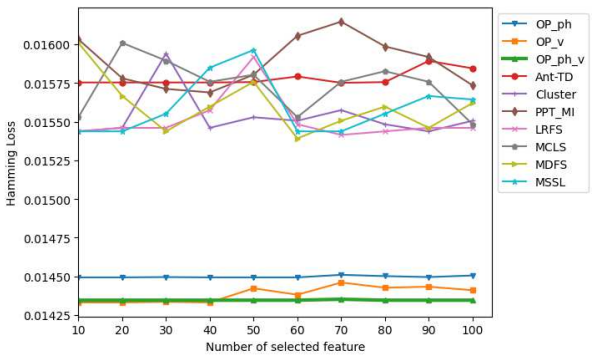
(d) Arts



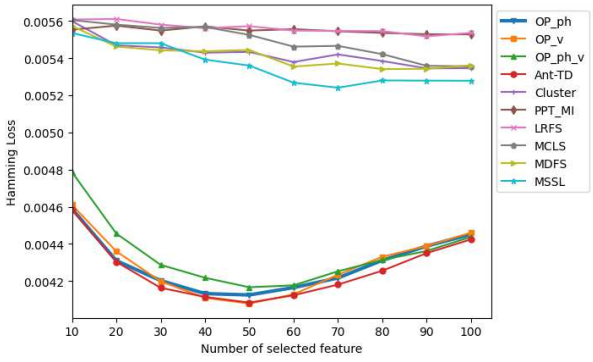
(e) Science



(f) Chess

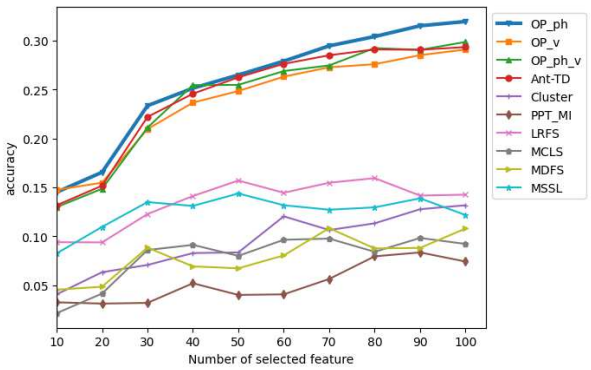


(g) Language

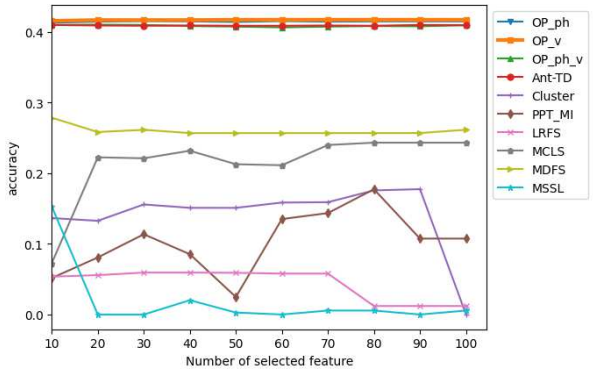


(h) Cooking

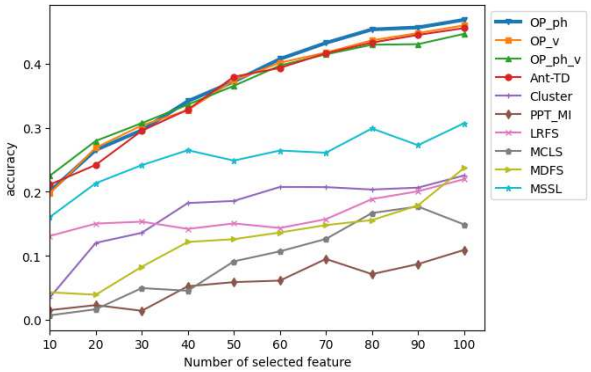
Figure 3. Average performance in terms of Hamming Loss for 20% of labeled data ↓



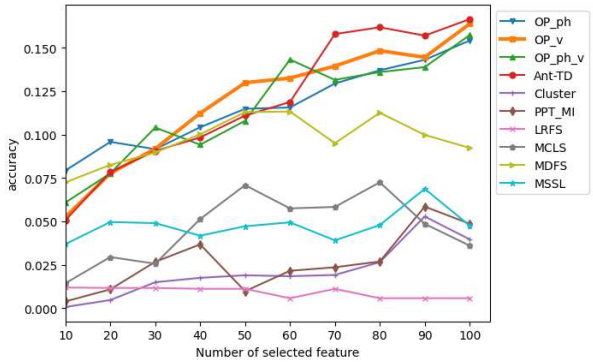
(a) Image



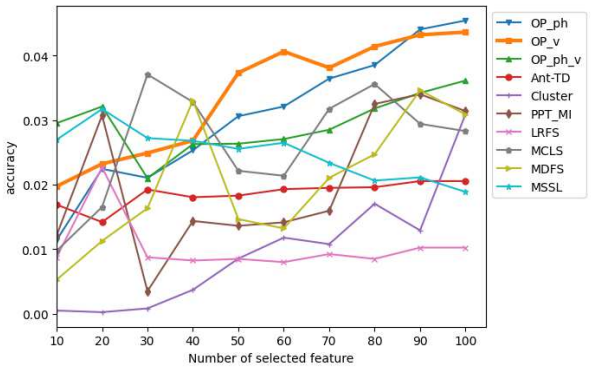
(b) Birds



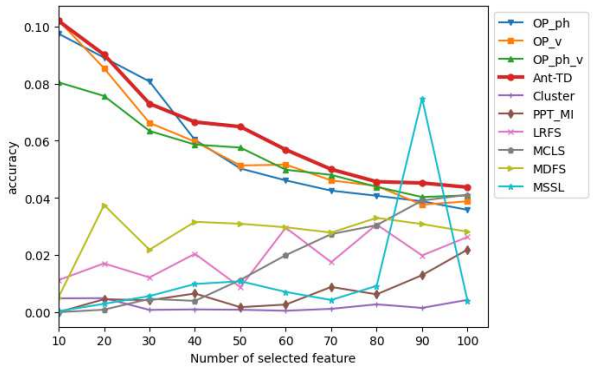
(c) Scene



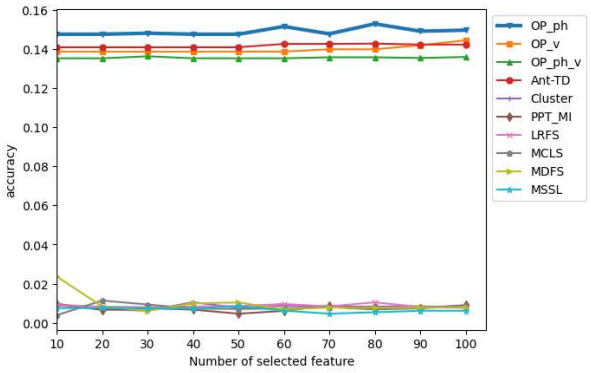
(d) Arts



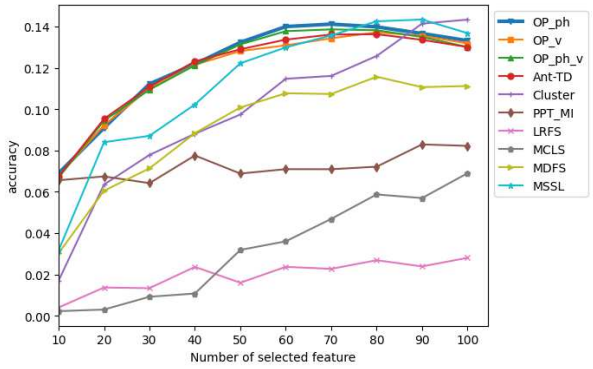
(e) Science



(f) Chess

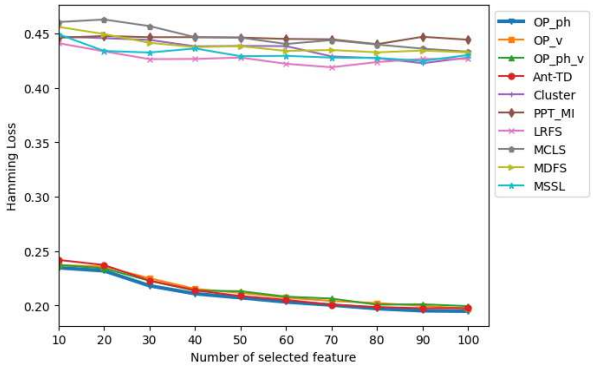


(g) Language

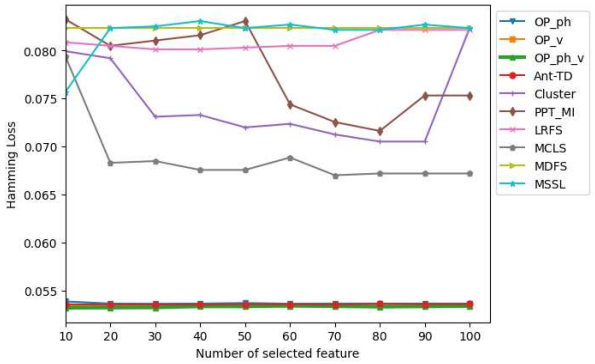


(h) Cooking

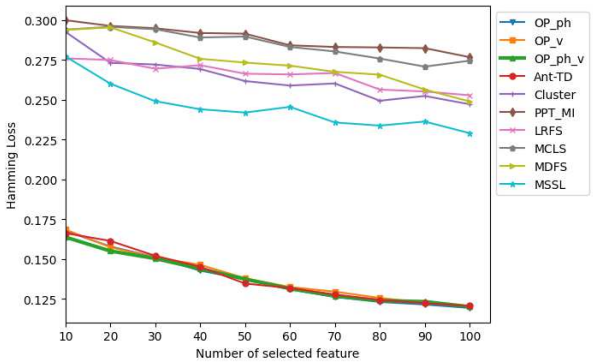
Figure 4. Average performance in terms of Accuracy for 40% of labeled data ↑



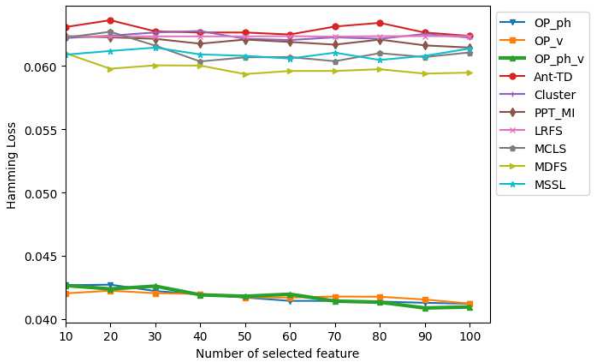
(a) Image



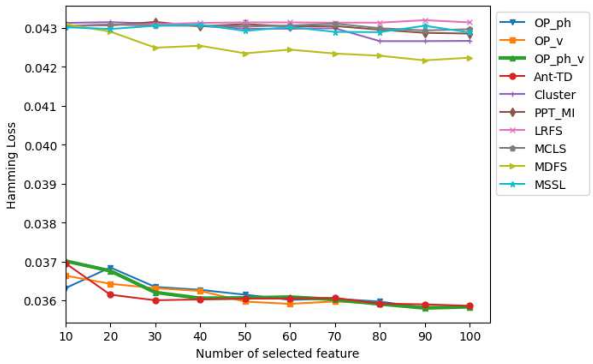
(b) Birds



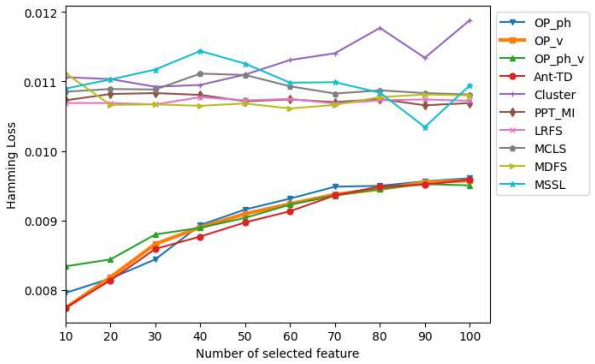
(c) Scene



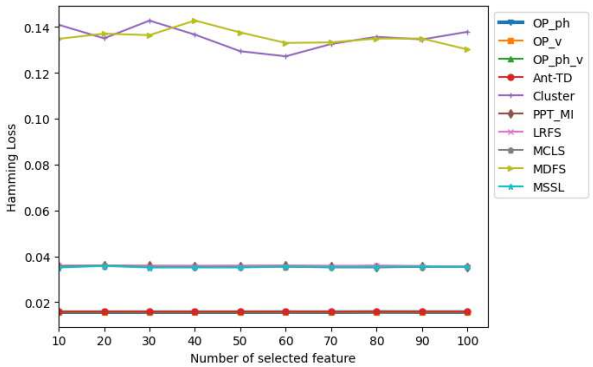
(d) Arts



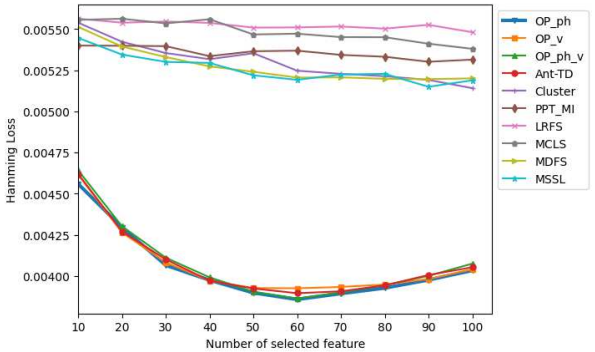
(e) Science



(f) Chess



(g) Language



(h) Cooking

Figure 5. Average performance in terms of Hamming Loss for 40% of labeled data ↓

Figure 5 shows a detailed comparison of Hamming loss for our semi-supervised algorithms using 40% labeled data. The results demonstrate that our algorithms achieve a substantial reduction in Hamming loss compared to several other methods. With a higher proportion of labeled data, our algorithms are able to leverage both semi-supervised learning and OBL techniques more effectively.

In this context, the increased amount of labeled data enhances the algorithm's ability to differentiate between relevant and irrelevant features. The semi-supervised approach allows for more precise updates to the feature selection process, while OBL helps in reinforcing effective paths and excluding less useful ones. As a result, the overall Hamming loss is significantly reduced, reflecting a more accurate and reliable feature selection process that minimizes incorrect label assignments. This improved performance highlights the advantage of using a higher proportion of labeled data in conjunction with advanced learning techniques.

Table 3 displays the average accuracy of the selected feature subsets across 10 different groupings, evaluated with 20% labeled data. The results illustrate that our three proposed semi-supervised methods consistently outperform other established techniques. Specifically, when using opposition in the pheromone update, the accuracy is often optimal. In other scenarios, employing opposition in the vector V yields better results. Furthermore, the combination of both opposition techniques usually provides the highest accuracy. This demonstrates that our proposed methods are more effective in leveraging semi-supervised learning to enhance feature selection accuracy compared to traditional methods. The improvements in accuracy can be attributed to the refined feature selection process enabled by the integration of semi-supervised data and OBL.

Table 3. Average classification performance in terms of Accuracy on 10 feature subsets for 20% of labeled data (the higher the better). ↑

	Op-ph-v	Op-ph	Op-v	Ant-TD	PPT-MI	MSSL	Cluster	MDFS	LRFS	MCLS
Cooking	0.0953	**0.1013**	0.0989	0.1002	0.0202	0.0868	0.0796	0.0703	0.0229	0.0426
Chess	0.0558	0.0581	0.0583	**0.0638**	0.0069	0.0061	0.0022	0.0276	0.0193	0.0178
Arts	0.0856	**0.0890**	0.0792	0.0841	0.0122	0.0179	0.0470	0.0563	0.0081	0.0196
Language	0.1375	**0.1430**	0.1400	0.1398	0.0171	0.0060	0.0078	0.0090	0.0103	0.0157
Science	0.0543	**0.0582**	0.0497	0.0437	0.0113	0.0229	0.0111	0.0415	0.0026	0.0110
Scene	0.3215	0.3355	**0.3445**	0.3224	0.0586	0.2235	0.1323	0.1274	0.1457	0.1226
Birds	**0.4230**	0.4144	0.4009	0.4066	0.1225	0.0211	0.1568	0.3235	0.0038	0.0495
Image	0.2710	**0.2813**	0.2760	0.2163	0.1246	0.2234	0.2077	0.1973	0.2075	0.1980
Average	0.1805	**0.1851**	0.1809	0.1721	0.0467	0.0760	0.0806	0.1066	0.0525	0.0596

Table 4 presents the average accuracy of feature selection across 10 different groupings, evaluated with 40% labeled data. Similar to Table 3, the results reveal that our three proposed methods consistently achieve higher accuracy compared to other methods. In this scenario, the simultaneous application of both opposition techniques—opposition in pheromone update and in vector V—frequently leads to the best accuracy. The use of 40% labeled data enhances the effectiveness of semi-supervised learning, allowing the algorithms to make more informed decisions and select more relevant features. This results in a more accurate feature selection process and improved performance over traditional methods. The data underscores the efficacy of integrating both opposition techniques in enhancing accuracy with a higher proportion of labeled data.

Table 4. Average classification performance in terms of Accuracy on 10 feature subsets for 40% of labeled data (the higher the better). ↑

	Op-ph-v	Op-ph	Op-v	Ant-TD	PPT-MI	MSSL	Cluster	MDFS	LRFS	MCLS
Cooking	0.1202	**0.1217**	0.1189	0.1195	0.0722	0.1115	0.0985	0.0903	0.0195	0.0324
Chess	0.0558	0.0581	**0.0583**	0.0580	0.0069	0.0128	0.0022	0.0276	0.0193	0.0178
Arts	0.1152	0.1165	**0.1193**	0.1191	0.0267	0.0477	0.0213	0.0971	0.0092	0.0465
Language	0.1355	**0.1489**	0.1398	0.1417	0.0074	0.0067	0.0081	0.0096	0.0086	0.0077
Science	0.0292	0.0307	**0.0338**	0.0186	0.0202	0.0248	0.0097	0.0205	0.0103	0.0246
Scene	0.3633	**0.3695**	0.3635	0.3601	0.0585	0.2533	0.1708	0.1267	0.1636	0.0934
Birds	0.4087	0.4148	**0.4171**	0.4093	0.1027	0.0192	0.1398	0.2602	0.0439	0.2142
Image	0.2423	**0.2572**	0.2384	0.2450	0.0523	0.1251	0.0941	0.0792	0.1352	0.0789
Average	0.1838	**0.1897**	0.1861	0.1839	0.0434	0.0751	0.0681	0.0889	0.0512	0.0644

Table 5 provides a detailed comparison of Hamming loss for the feature selection methods evaluated with 20% labeled data. The table reveals that our proposed semi-supervised algorithms—incorporating OBL—demonstrate a notable reduction in Hamming loss compared to other methods. The lower Hamming loss indicates fewer incorrect label assignments and, consequently, a more precise feature selection process. Among the proposed methods, those utilizing opposition in the pheromone update, opposition in the vector V, and the simultaneous use of both techniques consistently show lower Hamming loss values. This outcome highlights the advantage of using semi-supervised learning in conjunction with OBL updates to enhance the algorithm's ability to correctly classify features, leading to more accurate predictions and reduced errors.

Table 5. Average classification performance in terms of Hamming loss on 10 feature subsets for 20% of labeled data (the smaller the better). ↓

	Op-ph-v	Op-ph	Op-v	Ant-TD	PPT-MI	MSSL	Cluster	MDFS	LRFS	MCLS
Cooking	0.0043	**0.0042**	0.0042	0.0042	0.0055	0.0053	0.0054	0.0054	0.0055	0.0054
Chess	0.0090	**0.0090**	0.0089	0.0090	0.0101	0.0100	0.0103	0.0100	0.0102	0.0103
Arts	0.0525	0.0522	**0.0524**	0.0627	0.0633	0.0625	0.0627	0.0622	0.0634	0.0632
Language	0.0143	0.0144	0.0143	0.0157	**0.0158**	0.0155	0.0155	0.0156	0.0155	0.0157
Science	**0.0359**	0.0363	0.0366	0.0367	0.0430	0.0429	0.0429	0.0424	0.0431	0.0430
Scene	0.1456	0.1423	**0.1439**	0.1443	0.2926	0.2527	0.2760	0.3421	0.2701	0.2825
Birds	**0.0524**	0.0529	0.0558	0.0535	0.0762	0.0813	0.0746	0.2577	0.0832	0.0794
Image	0.1999	**0.1994**	0.2006	0.2178	0.2282	0.2179	0.2222	0.2309	0.2322	0.2265
Average	0.0642	**0.0638**	0.0646	0.0680	0.0918	0.0860	0.0887	0.1208	0.0904	0.0908

Table 6 illustrates the Hamming loss for feature selection methods evaluated with 40% labeled data. As with Table 5, our proposed semi-supervised algorithms exhibit a significant reduction in Hamming loss, indicating an improvement in classification accuracy and a decrease in incorrect label assignments. The integration of 40% labeled data provides a richer context for feature selection, enhancing the performance of the algorithms. In this table, the methods employing both opposition techniques—opposition in pheromone updates and in vector V—show the lowest Hamming loss values. This suggests that using a larger proportion of labeled data further refines the algorithm's feature selection process, resulting in fewer errors and more reliable predictions. The results underscore the effectiveness of semi-supervised learning combined with OBL approaches in minimizing Hamming loss and improving overall accuracy.

Table 6. Average classification performance in terms of Hamming loss on 10 feature subsets for 40% of labeled data (the smaller the better). ↓

	Op-ph-v	Op-ph	Op-v	Ant-TD	PPT-MI	MSSL	Cluster	MDFS	LRFS	MCLS
Cooking	0.0038	**0.0037**	0.0038	0.0038	0.0051	0.0051	0.0052	0.0051	0.0055	0.0054
Chess	0.0090	0.0090	0.0089	**0.0089**	0.0107	0.0109	0.0112	0.0107	0.0107	0.0109
Arts	**0.0417**	0.0418	0.0418	0.0628	0.0619	0.0609	0.0623	0.0598	0.0623	0.0611
Language	0.0159	**0.0156**	0.0157	0.0158	0.0357	0.0353	0.1352	0.1355	0.0357	0.0354
Science	**0.0360**	0.0361	0.0361	0.0361	0.0430	0.0429	0.0429	0.0424	0.0431	0.0430
Scene	**0.1376**	0.1378	0.1392	0.1386	0.2882	0.2452	0.2636	0.2733	0.2654	0.2846
Birds	**0.0533**	0.0536	0.0534	0.0535	0.0778	0.0817	0.0744	0.0823	0.0809	0.0688
Image	0.2136	**0.2092**	0.2134	0.2123	0.4454	0.4320	0.4358	0.4391	0.4273	0.4466
Average	0.0639	**0.0634**	0.0640	0.0660	0.1210	0.1140	0.1296	0.1316	0.1164	0.1197

In this framework, points in the graph that are deemed less important due to either the ant's movement choices or the evaporation process are strategically included in the feature selection cycle by the algorithm. If this newly explored path results in the selection of relevant features, it is reinforced by the deposition of pheromones, making it more likely to be chosen by other ants in subsequent iterations. Conversely, if the path does not lead to the selection of useful features, it is effectively excluded from the selection cycle by other ants, ensuring that the search process remains focused on promising areas of the feature space.

A detailed analysis of the tables and accuracy charts reveals that the top three methods for accuracy are those using opposition for updating the ACO algorithm. The differences among these three methods are minimal. A closer look at Tables 3 and 4 shows that, regarding accuracy for semi-supervised data, using opposition for pheromone updates yields the highest

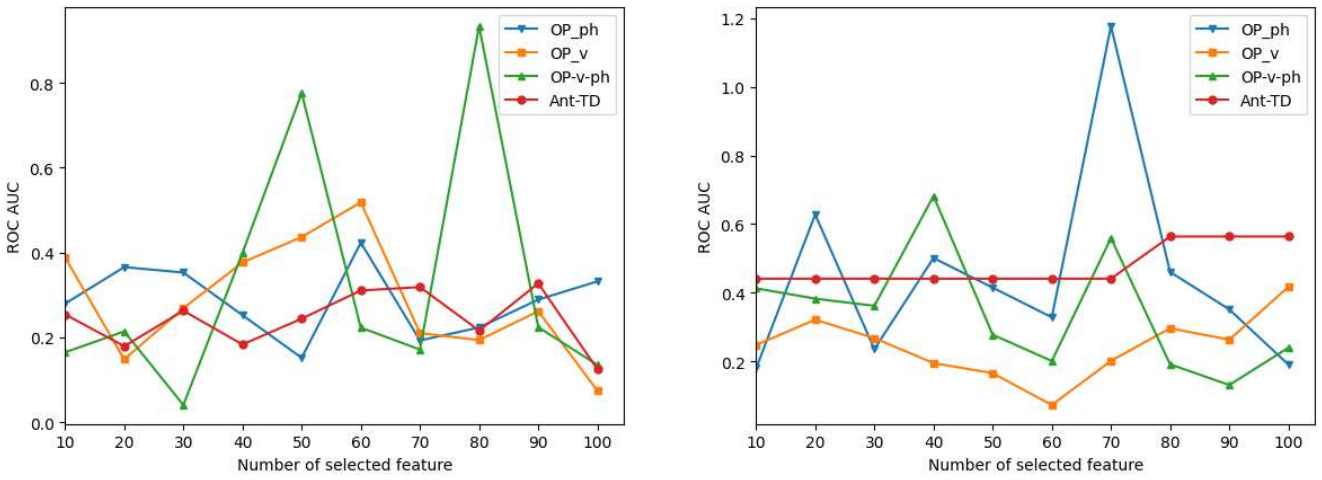
accuracy. The second-best accuracy is achieved by using opposition in updating the heuristic value V, while the third position goes to the simultaneous use of opposition for both pheromone and V updates. Furthermore, examining the impact of using 20% and 40% labeled data, as depicted in Figures 2 and 3 and and Tabless 3 and 4, shows that the accuracy of the algorithm increases as the amount of labeled data grows. These results confirm the validity of the earlier analysis, demonstrating that the more labeled data available, the better the accuracy of the algorithm.

In the discussion of Hamming loss across various datasets, as outlined in Tables 5 and 6 and illustrated in Figures 3 and 4, the top three methods with the lowest Hamming loss are those using opposition for updating the ACO algorithm. The Hamming loss values for these three methods are also very close to each other. Regarding Hamming loss, the best-performing method for most datasets is the one that uses opposition for simultaneous updates of both pheromone and V. The second-best method involves using opposition for updating the pheromone matrix, and the third-best method uses opposition for updating V. The analysis of using 20% and 40% labeled data, as shown in Figures 3 and 4 and Tables 5 and 6, again indicates that the accuracy of the algorithm increases with the proportion of labeled data. This consistency across different metrics and datasets further confirms the reliability of the previous analysis, highlighting that more labeled data results in improved performance in terms of accuracy and reduced Hamming loss.

In both scenarios, it is observed that for the Chess and Cooking datasets—which have the highest number of features among the evaluated datasets—the algorithm slows down as the sample size increases. This highlights the higher computational complexity of these datasets due to their large label counts, requiring more processing time. In contrast, for datasets with fewer features (Arts, Scene, Birds, Science, Image), the accuracy curve rises and Hamming loss curve declines as the number of selected features increases. This demonstrates the algorithm's ability to enhance performance through effective feature selection in low-label-density environments. For the Language dataset, which has an intermediate features count, the accuracy curve still rises and Hamming loss decreases, but the rate of accuracy improvement is slower compared to datasets with fewer labels. This indicates that while the algorithm continues to improve performance via proper feature selection even with higher label counts, the speed of accuracy gains diminishes as label complexity increases. These trends underscore the method's adaptability across varying label complexities while emphasizing the trade-offs between computational efficiency and precision in multi-label settings.

Figure 6 illustrates the pheromone matrix and V values that drive the algorithm's update mechanism across different iterations for "Birds" dataset. These values are plotted for both semi-supervised scenarios (20% and 40% labeled data). As observed, the baseline algorithm maintains nearly constant values throughout all iterations in both scenarios, indicating limited adaptability and exploration. In contrast, integrating OBL introduces dynamic exploration capabilities, enabling the proposed methods to escape local optima and enhance global search efficiency. This aligns with expectations, as OBL injects diversity into the search process, preventing stagnation and fostering more robust convergence.

The final matrix values for the three proposed methods exhibit distinct "jumps" in the update trajectories, particularly visible in later iterations. These fluctuations reflect OBL's role in balancing exploration and exploitation, where opposition-driven perturbations guide the algorithm toward regions of higher accuracy and reliability. The abrupt shifts in the matrix values signify moments of critical adaptation, where the algorithm leverages opposition samples to refine its search strategy. This behavior underscores OBL's effectiveness in improving algorithmic power and precision, especially in semi-supervised settings where labeled data scarcity amplifies the risk of suboptimal convergence. The results highlight how OBL not only enhances exploration but also strengthens the algorithm's ability to generalize, making it a vital component for overcoming challenges in dynamic or noisy environments.



(a) 20% labeled data **(b)** 40% labeled data

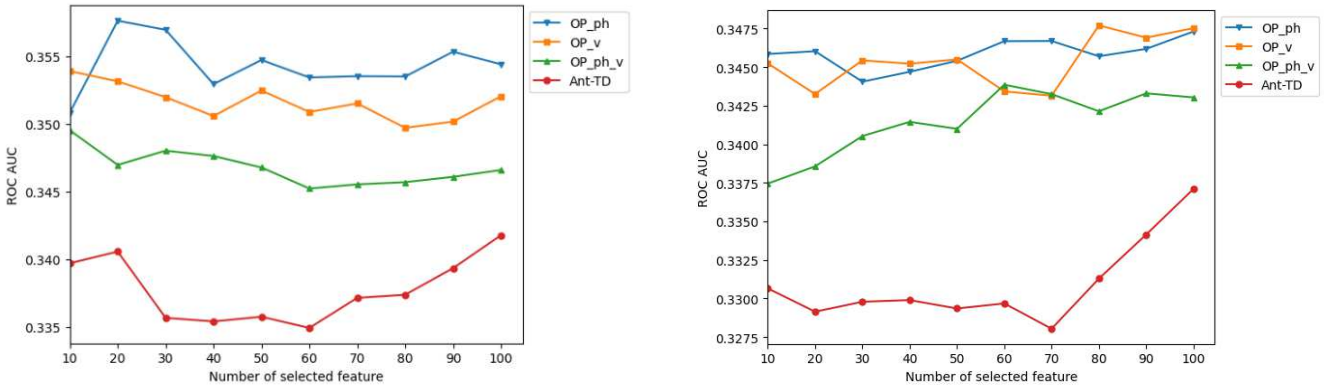
Figure 6. Convergence Dynamics of Pheromone Matrix Updates in Semi-Supervised Scenarios

The Friedman test outcomes highlight significant disparities in the performance of the four methods across the eight datasets, with notable variations depending on the evaluation metric (accuracy or Hamming loss) and the amount of labeled data. For datasets like "Birds" and "Image", the high Friedman statistics suggest strong disagreements among the methods, implying that at least one proposed method significantly outperforms the baseline ant-td in both accuracy and error reduction. In contrast, datasets such as "Chess" and "Language" exhibit minimal variation between methods, indicating that the proposed approaches fail to meaningfully improve upon the baseline, likely due to inherent dataset complexity or insufficient feature representation. The results also reveal a trade-off between accuracy and Hamming loss in certain domains "Science" and "Arts", where method rankings diverge sharply between metrics, underscoring the challenge of balancing precision and error avoidance in semi-supervised learning. Increasing labeled data from 20% to 40% amplifies method consistency in some domains "Cooking" but exacerbates disparities in others "Language", reflecting dataset-specific sensitivities to supervision. For instance, "Scene" shows improved agreement in Hamming loss with more labeled data, suggesting reduced error variability, while its accuracy results become more polarized, hinting at unresolved class imbalance or feature noise. The stagnation in "Language" and erratic behavior in Arts across both metrics emphasize the need for adaptive feature engineering or domain-specific constraints. Overall, the Friedman rankings in Table 7 underscore that method efficacy is highly context-dependent, with no universal solution, and advocate for hybrid strategies that leverage dataset-specific strengths while addressing limitations through regularization or active learning.

Table 7. Friedman Test Results for Comparative Performance of Proposed Algorithms

Scenario	Metric	Chess	Cooking	Arts	Birds	Language	Science	Scene	Image
20%	Accuracy	14.03	14.75	6.96	30.0	16.21	13.31	13.31	19.56
	Hamming Loss	8.75	16.43	21.72	30.0	27.12	27.12	12.12	19.91
40%	Accuracy	14.03	9.96	2.93	28.07	28.91	19.56	5.03	16.68
	Hamming Loss	8.75	10.79	18.47	28.91	30.0	3.71	11.15	18.84

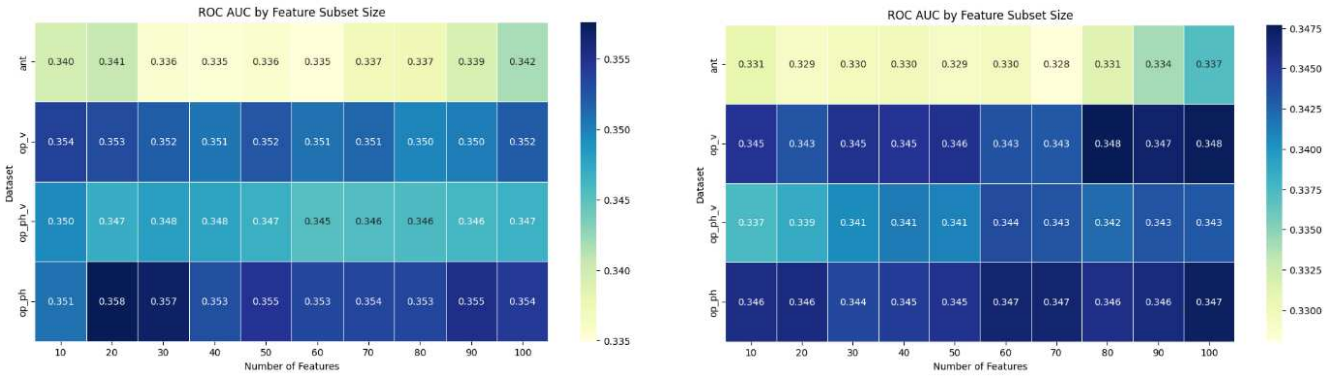
Figure 7, which compares the ROC curves for the two semi-supervised scenarios for "Birds" dataset, clearly illustrates the performance hierarchy of the methods. In this analysis, the hybrid method combining operational and photometric features (op_ph) emerged as the top-performing model, excelling in discriminiscriminative power and predictive accuracy, while the baseline method (ant) demonstrated significantly weaker performance and instability. The op_v and op_ph_v methods ranked second and third, respectively, with the former balancing stability and accuracy and the latter prioritizing result consistency, making them suitable alternatives depending on priorities (high accuracy vs. reliability). The curves further confirm a marked gap between the top-performing methods and the baseline, underscoring the need to refine or redesign weaker approaches to address their limitations.



(a) 20% labeled data **(b)** 40% labeled data

Figure 7. ROC Curve Comparison of Proposed Algorithms in Semi-Supervised Scenarios

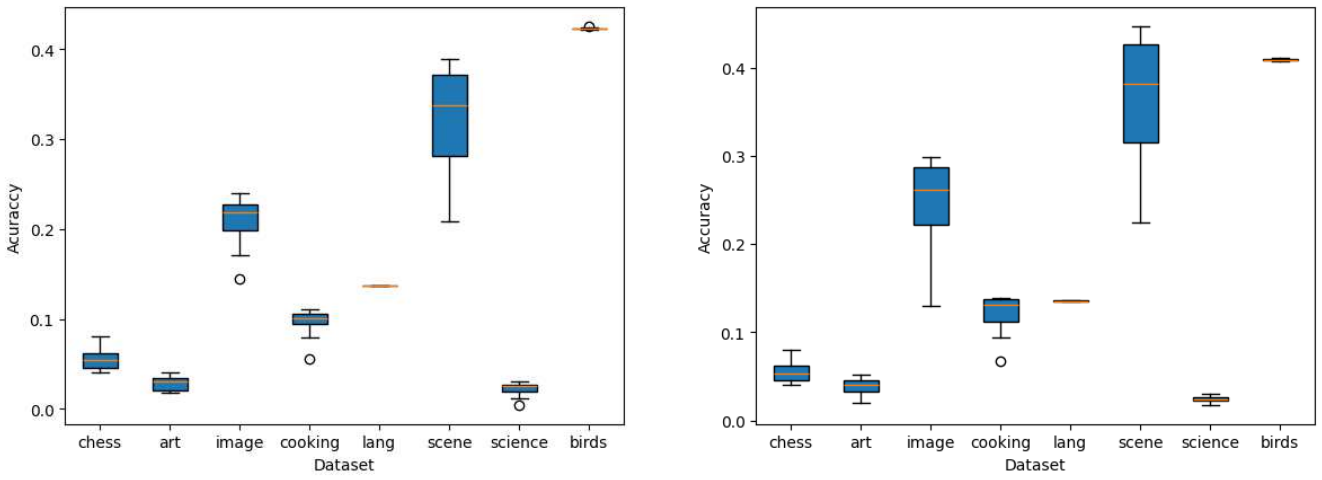
Figure 8, displaying the heatmap for feature selection in the two semi-supervised scenarios for "Birds" dataset, highlights clear patterns in model performance and feature utility. The hybrid method combining operational and photometric features (op_ph) consistently outperforms others, demonstrating robust discriminative power and stability, while the baseline approach (ant) trails significantly, marked by erratic performance and sensitivity to noise. Operational features alone (op_v) strike a balance between reliability and effectiveness, whereas the extended hybrid method (op_ph_v) shows gradual improvement, suggesting adaptive learning potential but requiring optimization to minimize variability. The heatmap underscores the critical role of feature integration in semi-supervised settings, with op_ph emerging as the most viable strategy, while ant's shortcomings emphasize the need for feature refinement or algorithmic adjustments to enhance relevance and reduce redundancy.



(a) 20% labeled data **(b)** 40% labeled data

Figure 8. Feature Selection Efficacy in Semi-Supervised Scenarios – Heatmap Analysis of ROC-Driven Algorithm Performance

Figure 9, displaying the boxplot for accuracy across the 8 datasets in the two semi-supervised scenarios, reveals distinct performance patterns and variability. The "Birds" dataset stands out as the top performer, demonstrating exceptionally high and stable values with minimal spread, indicating consistent and reliable results. "Scene" and "Image" show strong upward trends, with medians rising significantly across iterations, though "Scene" exhibits slightly wider variability in later stages. "Art" and "Cooking" display moderate performance with fluctuating results—"Art" improves gradually but remains uneven, while "Cooking" has sporadic peaks and troughs, suggesting sensitivity to specific features or conditions. In contrast, "Chess" and "Science" underperform, with consistently low medians and limited spread, reflecting poor discriminative power. "Lang" is uniquely stagnant, showing near-zero variability and flatlined medians, implying algorithmic rigidity or insufficient feature adaptation. Notably, "Image" and "Scene" exhibit the most dynamic distributions, with outliers hinting at occasional high-performing iterations. These results emphasize the trade-offs between stability and adaptability: while "Birds" excels in consistency, "Scene" and "Image" balance improvement with controlled variability. The weak performance of "Chess" and "Science" highlights potential dataset-specific limitations or the need for model redesign.



(a) 20% labeled data

(b) 40% labeled data

Figure 9. Accuracy Distribution in Semi-Supervised Scenarios – Boxplot Comparison of Proposed Algorithms Across Labeled Data Proportions

The integration of semi-supervised learning techniques allows the algorithm to better utilize available data, even when only a portion of it is labeled. This leads to a more robust and accurate feature selection process. The ability to distinguish between relevant and irrelevant features more precisely helps in reducing Hamming loss, which measures the frequency of incorrect label assignments. By continually refining the feature selection process through the use of semi-supervised data and OBL, our algorithms achieve higher accuracy and reliability compared to traditional methods. However, while the proposed algorithms (ACO-OBL-RL) outperform baseline and competing methods in terms of exploration capability and Hamming loss, they incur trade-offs in computational efficiency and resource utilization. Specifically, the calculation of opposition states for OBL introduces additional overhead, resulting in slower execution times compared to the baseline ACO. Furthermore, the need to store both original and opposition solutions during optimization increases memory consumption. The computational complexity is also heightened due to the iterative interaction between RL-guided pheromone updates, OBL-based exploration, and ACO's inherent stochastic search mechanisms. These trade-offs highlight the balance between performance gains and practical resource constraints, which may require optimization for deployment in resource-limited environments.

Conclusion

This paper introduces a novel multi-label feature selection method that integrates semi-supervised learning with ACO, enhanced by OBL and RL. MLKNN was chosen as the classifier to ensure fair benchmarking due to its inherent ability to model label correlations via k-nearest neighbors, aligning with the study's focus on multi-label tasks. The approach addresses high-dimensional data challenges by leveraging both labeled and unlabeled instances to improve accuracy and reduce redundancy. Experiments on eight diverse datasets—spanning text, images, and categorical data—evaluated two scenarios with 20% and 40% labeled data. Results demonstrated significant improvements in classification accuracy and reductions in Hamming loss compared to traditional methods. Specifically, the 20% labeled data scenario showcased the method's ability to extract meaningful information from 80% unlabeled instances, while the 40% scenario highlighted its scalability and robustness, further enhancing accuracy with increased labeled data.

The framework's adaptability extends to critical applications such as biomarker selection in genomics, sentiment classification in noisy text, and multimedia tagging systems. Its modular design enables seamless integration into existing machine learning pipelines (e.g., TensorFlow, PyTorch) and distributed computing environments, making it viable for real-world deployment in healthcare diagnostics, recommendation engines, or IoT systems. Economically, the method reduces computational and storage costs, offering tangible benefits for industries like finance and telecommunications that rely on high-dimensional datasets. By minimizing redundancy and optimizing feature selection, the approach balances efficiency with performance, addressing both technical and practical demands.

To address scalability, future work will explore hybrid architectures combining ACO with Particle Swarm Optimization (PSO) or Differential Evolution (DE) to enhance global search efficiency while retaining interpretability. Additionally, hardware-aware optimizations (e.g., GPU parallelization, sparse matrix representations) and adaptive iteration limits are critical to reduce runtime and memory strain. While the current framework is unsuitable for real-time applications due to its iterative,

batch-oriented design, optimizations like lightweight feature subset caching or incremental learning could enable near-real-time usage in constrained settings. Economically, the method's ability to reduce redundancy and storage costs remains advantageous for industries like finance and telecommunications, but balancing these benefits with computational trade-offs will be essential for industrial-scale deployment.

References

1. Hatami, M., Moradi, P., Sulaimany, S. & Jalili, M. Maximum relevant minimum redundant multi-label feature selection using ant colony optimization. *Eng. Appl. Artif. Intell.* **161**, 112007 (2025).
2. Ernst, E., Damien, D. & Louette, A. Introduction to reinforcement learning. *Bus. & Inf. Syst. Eng.* **66**, 111–126 (2024).
3. Wang, H. *et al.* Feature selection strategies: A comparative analysis of shap-value and importance-based methods. *J. Big Data* **11**, 44, DOI: 10.1186/s40537-024-00900-1 (2024).
4. Hashemi, A., Dowlatshahi, M. B. & Nezamabadi-Pour, H. A bipartite matching-based feature selection for multi-label learning. *Int. J. Mach. Learn. Cybern.* **12**, 459–475, DOI: 10.1007/s13042-020-01180-w (2021).
5. Yu, Y. *et al.* Feature selection for multi-label learning based on variable-degree multi-granulation decision-theoretic rough sets. *Int. J. Approx. Reason.* **169**, 109181, DOI: 10.1016/j.ijar.2024.109181 (2024).
6. Faraji, M. *et al.* Multi-label feature selection with global and local label correlation. *Expert. Syst. with Appl.* **246**, 123198, DOI: 10.1016/j.eswa.2024.123198 (2024).
7. Turukmane, A. V. & Devendiran, R. M-multisvm: An efficient feature selection assisted network intrusion detection system using machine learning. *Comput. & Secur.* **137**, 103587, DOI: 10.1016/j.cose.2024.103587 (2024).
8. Iranzad, R. & Liu, X. A review of random forest-based feature selection methods for data science education and applications. *Int. J. Data Sci. Anal.* DOI: 10.1007/s41060-024-00526-9 (2024).
9. Sari, D. W., Dwijayanti, S. & Suprapto, B. Y. Integration of regression-based guidance ant for enhanced exploration and convergence in ant colony optimization (aco). *IEEE Access* (2025).
10. Dorigo, M. & Stützle, T. *Ant Colony Optimization: Overview and Recent Advances*, 311–351 (Springer, 2019).
11. Manakkadu, S. & Dutta, S. Ant colony optimization based support vector machine for improved classification of unbalanced datasets. *Procedia Comput. Sci.* **237**, 586–593, DOI: 10.1016/j.procs.2024.03.123 (2024).
12. Tizhoosh, H. R. Opposition-based learning: A new scheme for machine intelligence. In *CIMCA-IAWTIC'06*, 695–701, DOI: 10.1109/CIMCA.2005.1631345 (2005).
13. Gopi, S. & Mohapatra, P. Fast random opposition-based learning aquila optimization algorithm. *Heliyon* **10**, e25643, DOI: 10.1016/j.heliyon.2024.e25643 (2024).
14. Qtaish, A. *et al.* Enhanced coati optimization algorithm using elite opposition-based learning and adaptive search mechanism for feature selection. *Int. J. Mach. Learn. Cybern.* **16**, 361–394, DOI: 10.1007/s13042-024-02167-7 (2025).
15. François-Lavet, V., Henderson, P., Islam, R., Bellemare, M. G. & Pineau, J. An introduction to deep reinforcement learning. *Foundations Trends Mach. Learn.* **11**, 219–354, DOI: 10.1561/2200000071 (2018).
16. Chen, C., Han, J. & Debattista, K. Virtual category learning: A semi-supervised learning method for dense prediction with extremely limited labels. *IEEE Transactions on Pattern Analysis Mach. Intell.* DOI: 10.1109/TPAMI.2024.3375188 (2024).
17. Li, R. *et al.* Semi-supervised multi-label feature selection combining nonlinear manifold structure and minimizing group sparse redundant correlation. *Expert. Syst. with Appl.* **273**, 126844, DOI: 10.1016/j.eswa.2023.126844 (2025).
18. Paniri, M., Dowlatshahi, M. B. & Nezamabadi-Pour, H. Mlaco: A multi-label feature selection algorithm based on ant colony optimization. *Knowledge-Based Syst.* **192**, 105285 (2020).
19. Hashemi, A., Dowlatshahi, M. B. & Nezamabadi-Pour, H. Mfs-mcdm: Multilabel feature selection using multi-criteria decision making. *Knowledge-Based Syst.* **206**, 106365 (2020).
20. Hashemi, A., Dowlatshahi, M. B. & Nezamabadi-Pour, H. Ensemble of feature selection algorithms: a multi-criteria decision-making approach. *Int. J. Mach. Learn. Cybern.* **13**, 49–69, DOI: 10.1007/s13042-021-01347-z (2022).
21. Karimi, F., Dowlatshahi, M. B. & Hashemi, A. Semiaco: A semi-supervised feature selection based on ant colony optimization. *Expert. Syst. with Appl.* **214**, 119130, DOI: 10.1016/j.eswa.2022.119130 (2023).

22. Halder, A., Ghosh, S. & Ghosh, A. Ant based semi-supervised classification. In *International Conference on Swarm Intelligence*, 490–497 (2010).

23. Albinati, J., Lopes, H. S., Freitas, A. & Gosling, A. An ant colony-based semi-supervised approach for learning classification rules. *Swarm Intell.* **9**, 315–341 (2015).

24. Paniri, M., Dowlatshahi, M. B. & Nezamabadi-Pour, H. Ant-td: Ant colony optimization plus temporal difference reinforcement learning for multi-label feature selection. *Swarm Evol. Comput.* **64**, 100892, DOI: 10.1016/j.swevo.2021.100892 (2021).

25. Munif, M. M., Nasir, H. B. J. A. & Ahmad, M. I. Enhanced ant colony system for water quality rules classification in smart aquaculture. In *2024 IEEE Symposium on Wireless Technology & Applications (ISWTA)*, 216–220, DOI: 10.1109/ISWTA62130.2024.10652089 (2024).

26. Hashemi, A., Joodaki, M., Joodaki, N. Z. & Dowlatshahi, M. B. Ant colony optimization equipped with an ensemble of heuristics through multi-criteria decision making: A case study in ensemble feature selection. *Appl. Soft Comput.* **124**, 109046, DOI: 10.1016/j.asoc.2022.109046 (2022).

27. Li, J., Gao, L. & Li, X. Multi-operator opposition-based learning with the neighborhood structure for numerical optimization problems and its applications. *Swarm Evol. Comput.* **84**, 101457, DOI: 10.1016/j.swevo.2023.101457 (2024).

28. Huang, P. *et al.* Orthogonal opposition-based learning honey badger algorithm with differential evolution for global optimization and engineering design problems. *Alex. Eng. J.* **91**, 348–367, DOI: 10.1016/j.aej.2024.02.024 (2024).

29. Li, W. & Wang, G. G. Improved elephant herding optimization using opposition-based learning and k-means clustering to solve numerical optimization problems. *J. Ambient Intell. Humaniz. Comput.* **14**, 1753–1784, DOI: 10.1007/s12652-021-03391-7 (2023).

30. Farahmand-Tabar, S. & Shirgir, S. *Opposed Pheromone Ant Colony Optimization for Property Identification of Nonlinear Structures* (Springer, 2024).

31. Malisia, A. R. & Tizhoosh, H. R. Applying opposition-based ideas to the ant colony system. In *2007 IEEE Swarm Intelligence Symposium*, 182–189, DOI: 10.1109/SIS.2007.368044 (2007).

32. Dorigo, M. & Blum, C. Ant colony optimization theory: A survey. *Theor. Comput. Sci.* **344**, 243–278, DOI: 10.1016/j.tcs.2005.05.020 (2005).

33. Gu, X., Li, X., Ke, W., Wang, Z. & Lv, Y. Feature selection algorithm based on ant colony optimization and mutual information. *Clust. Comput.* **29**, 180 (2026).

34. Yang, Q. *et al.* Adaptive multimodal continuous ant colony optimization. *IEEE Transactions on Evol. Comput.* **21**, 191–205, DOI: 10.1109/TEVC.2016.2591064 (2017).

35. Sun, B. *et al.* Tensor-based ant colony optimization for set meal design in online-to-offline restaurants. *IEEE Transactions on Cybern.* (2026).

36. Liu, C. *et al.* An improved heuristic mechanism ant colony optimization algorithm for solving path planning. *Knowledge-Based Syst.* **271**, 110540, DOI: 10.1016/j.knosys.2023.110540 (2023).

37. Kelidari, M. & Hamidzadeh, J. Feature selection by using chaotic cuckoo optimization algorithm with levy flight, opposition-based learning and disruption operator. *Soft Comput.* **25**, 2911–2933, DOI: 10.1007/s00500-020-05349-x (2021).

38. Tubishat, M., Idris, N., Shuib, L., Abushariah, M. A. M. & Mirjalili, S. Improved salp swarm algorithm based on opposition based learning and novel local search algorithm for feature selection. *Expert. Syst. with Appl.* **145**, 113122, DOI: 10.1016/j.eswa.2019.113122 (2020).

39. Mohapatra, S. & Mohapatra, P. Fast random opposition-based learning golden jackal optimization algorithm. *Knowledge-Based Syst.* **275**, 110679, DOI: 10.1016/j.knosys.2023.110679 (2023).

40. Wiering, M. A. & Van Hasselt, H. Two novel on-policy reinforcement learning algorithms based on td(λ)-methods. In *IEEE International Symposium on Approximate Dynamic Programming and Reinforcement Learning*, DOI: 10.1109/ADPRL.2007.368199 (2007).

41. Awan, H. H. & Shahzad, W. Semi-supervised associative classification using ant colony optimization algorithm. *PeerJ Comput. Sci.* **7**, e676 (2021).

42. Chen, J. *et al.* Open-set semi-supervised text classification via adversarial disagreement maximization. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics*, 574–586 (2024).

43. Pluim, J. P. W., Maintz, J. B. A. & Viergever, M. A. Mutual-information-based registration of medical images: a survey. *IEEE Transactions on Med. Imaging* **22**, 986–1004, DOI: 10.1109/TMI.2003.815867 (2003).

44. Xia, P., Zhang, L. & Li, F. Learning similarity with cosine similarity ensemble. *Inf. Sci.* **307**, 39–52, DOI: 10.1016/j.ins.2015.02.016 (2015).

45. Liu, Y. & Cao, B. A novel ant colony optimization algorithm with levy flight. *IEEE Access* **8**, 67205–67213, DOI: 10.1109/ACCESS.2020.2986932 (2020).

46. Tsoumakas, G., Spyromitros-Xioufis, E., Vilcek, J. & Vlahavas, I. Mulan: A java library for multi-label learning. *J. Mach. Learn. Res.* **12**, 2411–2414 (2011).

47. Tsoumakas, G., Katakis, I. & Vlahavas, I. Mining multi-label data. In *Data Mining and Knowledge Discovery Handbook*, DOI: 10.1007/978-0-387-09823-4 (Springer, 2009).

48. Zhang, M.-L. & Zhou, Z.-H. Ml-knn: A lazy learning approach to multi-label learning. *Pattern Recognit.* **40**, 2038–2048, DOI: 10.1016/j.patcog.2006.12.019 (2007).

49. Swets, J. A. Measuring the accuracy of diagnostic systems. *Science* **240**, 1285–1293, DOI: 10.1126/science.3287615 (1988).

50. Norouzi, M., Fleet, D. J. & Salakhutdinov, R. R. Hamming distance metric learning. In *Advances in Neural Information Processing Systems* (2012).

51. Kashef, S. & Nezamabadi-Pour, H. A label-specific multi-label feature selection algorithm based on the pareto dominance concept. *Pattern Recognit.* **88**, 654–667, DOI: 10.1016/j.patcog.2018.12.020 (2019).

52. Zhang, P., Liu, G. & Gao, W. Distinguishing two types of labels for multi-label feature selection. *Pattern Recognit.* **95**, 72–82, DOI: 10.1016/j.patcog.2019.05.023 (2019).

53. Doquire, G. & Verleysen, M. Mutual information-based feature selection for multilabel classification. *Neurocomputing* **122**, 148–155, DOI: 10.1016/j.neucom.2013.04.041 (2013).

54. Huang, R., Jiang, W. & Sun, G. Manifold-based constraint laplacian score for multi-label feature selection. *Pattern Recognit. Lett.* **112**, 346–352, DOI: 10.1016/j.patrec.2018.08.009 (2018).

55. Zhang, J. *et al.* Manifold regularized discriminative feature selection for multi-label learning. *Pattern Recognit.* **95**, 136–150, DOI: 10.1016/j.patcog.2019.06.003 (2019).

56. Bayati, H., Dowlatshahi, M. B. & Hashemi, A. Mssl: A memetic-based sparse subspace learning algorithm for multi-label classification. *Int. J. Mach. Learn. Cybern.* **13**, 3607–3624, DOI: 10.1007/s13042-022-01616-5 (2022).

Author contributions statement

S. S. B: Writing, Original draft preparation, Methodology, Software, Data curation. V. M: Supervision, Conceptualization, Writing-Reviewing, and Editing, Methodology. M. B. D: Supervision, Conceptualization, Writing-Reviewing and Editing, Methodology.

Additional Information

Data Availability Statement

The datasets used in this study are publicly available. All datasets, including Birds, Scene, Arts, Science, Chess, Cooking, Image, and Language, can be accessed through the following repositories:

- Boomer Datasets: `https://github.com/mrapp-ke/Boomer-Datasets`
- MULAN Datasets: `https://mulan-dataset.github.io/`

Both repositories provide open access to the data used in this research, and no special permissions are required to obtain them.

Ethical Approval

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper. We undersigned declare that this manuscript is original, has not been published before, and is not currently being considered for publication elsewhere.

Funding
No funds were used to prepare this article.

Author Agreement Statement
We the undersigned declare that this manuscript is original, has not been published before, and is not currently being considered for publication elsewhere. We confirm that the manuscript has been read and approved by all named authors and that there are no other persons who satisfied the criteria for authorship but are not listed. We further confirm that the order of authors listed in the manuscript has been approved by all of us. We understand that the Corresponding Author is the sole contact for the Editorial process. He is responsible for communicating with the other authors about progress and submissions of revisions. We confirm that the manuscript has been read and approved by all named authors and that there are no other persons who satisfied the criteria for authorship but are not listed. We further confirm that the order of authors listed in the manuscript has been approved by all of us.