

FlameFinder: Illuminating Obscured Fire Through Smoke With Attentive Deep Metric Learning

Hossein Rajoli, *Graduate Student Member, IEEE*, Sahand Khoshdel, Fatemeh Afghah^{ID}, *Senior Member, IEEE*, and Xiaolong Ma^{ID}

Abstract—FlameFinder, a novel deep metric learning (DML) framework, accurately detects RGB-obscured flames using thermal images from firefighter drones during wildfire monitoring. In contrast to RGB, thermal cameras can capture smoke-obscured flame features but they lack absolute thermal reference points, detecting many nonflame hot spots as false positives. This issue suggests that extracting features from both modalities in unobscured cases can reduce the model's bias to relative thermal gradients. Following this idea, our proposed model utilizes paired thermal–RGB images captured onboard drones for training, learning latent flame features from smoke-free samples. In testing, it identifies flames in smoky patches based on their equivalent thermal-domain distribution, improving performance with supervised and distance-based clustering metrics. The approach includes a flame segmentation method and a DML-aided detection framework with center loss (CL), triplet CL (TCL), and triplet cosine CL (TCCL), to find the optimal cluster representatives for classification. Evaluation of FLAME2 and FLAME3 datasets shows the method's effectiveness in diverse fire and no-fire scenarios. However, the CL dominates the two other losses, resulting in the model missing features that are sensitive to them. To overcome this issue, an attention mechanism is proposed making nonuniform feature contribution possible and amplifying the critical role of cosine and triplet loss in the DML framework. Plus, the attentive DML shows improved interpretability, class discrimination, and decreased intraclass variance exploiting several other flame-related features. The proposed model surpasses the baseline with a binary classifier by 4.4% in FLAME2 and 7% in FLAME3 datasets for unobscured flame detection accuracy. It also demonstrates enhanced class separation in obscured scenarios compared to fine-tuned VGG19, ResNet18, and three other backbone models tailored for flame detection.

Index Terms—Attention, deep metric learning, flame detection.

I. INTRODUCTION

FOREST fires cause property damage, injuries, and deaths, harming people and the economy. Extreme wildfire events pose greater risks and challenges. Therefore, efficient wildfire detection and management are essential to mitigate these impacts [1]. Existing technologies are not yet capable of accurate and timely detection of wildfire noting the limited

lifetime and sensing range of thermal/smoke sensors, delays, and low spatial resolution associated with satellites' observations [2]. In the context of wildfire detection, unmanned aerial vehicles (UAVs) have emerged as a promising technology for wildfire detection and management. UAVs offer high mobility, flexibility [3], low deployment cost, and real-time data collection capabilities, making them well-suited for monitoring wildfires in remote and challenging terrains [4], [5]. Deploying UAV systems for wildfire management has not been limited to wildfire detection [6], [7]. The performance of UAVs in other tasks such as active wildfire monitoring with decision-making systems [8] and spread modeling of fire frontiers rely on how accurate their input is provided by a flame detection system [7]. While many improvements have been made in terms of general flame detection, commercial UAVs with visual cameras cannot yet accurately identify the flame locations when smoke and water vapor block visible-light spot fires, resulting in false negatives. It should be noted that while smoke detected in RGB images can be an indicator of fire, it cannot accurately pinpoint the fire location and intensity due to the complex dynamics of smoke, wind fire, and several other effective variables. As a result, utilizing thermal images was hypothesized as a solution for drone-based flame mapping [9].

While some studies focus on thermal-domain fire detection images by exploring sensor response and image-processing techniques, flame detection for smoke-obscured samples is understudied and remains a major gap in this area [10]. On the other hand, the lack of absolute thermal reference points and the existence of other thermal sources make models trained solely in thermal domains inaccurate. Shortcomings of pure RGB and pure thermal-domain methods pose the fact that the problem of obscured flame detection naturally needs complementary modalities. Following this idea, to extract features associated with flames and not just hotspots occurring in the wildland (e.g., carried smoke), an idea would be to transfer knowledge (using labels), from unobscured samples in the RGB domain (as an indirect supervision) to the unannotated target thermal domain. Deep neural networks (DNNs) have shown potential for UAV-based flame detection by integrating data from RGB and thermal cameras [11]. Such integration provides complementary information, improving the discrimination of flames from other heat sources. Transfer learning and fine-tuning on DNN architectures have demonstrated enhanced flame detection accuracy in UAV-collected images [12]. This process can be initiated by providing annotations with a segmentation algorithm on the source domain. Next, to aggregate the embedded information of the thermal domain, a model

Received 4 April 2024; revised 27 May 2024; accepted 17 July 2024. Date of publication 12 August 2024; date of current version 29 October 2024. This work was supported in part by the Air Force Office of Scientific Research under Award FA9550-20-1-0090, in part by the National Aeronautics and Space Administration (NASA) under Award 80NSSC23K1393, and in part by the National Science Foundation under Grant CNS-2232048 and Grant CNS-2204445. (Corresponding author: Fatemeh Afghah.)

The authors are with the Holcombe Department of Electrical and Computer Engineering, Clemson University, Clemson, SC 29634 USA (e-mail: hrajoli@clemson.edu; skhoshd@clemson.edu; fafghah@clemson.edu; xialom@clemson.edu).

Digital Object Identifier 10.1109/TGRS.2024.3440880

1558-0644 © 2024 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See <https://www.ieee.org/publications/rights/index.html> for more information.

is trained to shape a compressed latent space with objectives that capture various flames, rather than simply getting biased toward the thermal domain intensity. Solving this challenge is core to overcome the problem of relative thermal reference points and nonflame hotspots appearing as false positives.

Among representation learning approaches proposed for clustering, deep metric learning (DML) has shown major success for learning similarity and distance metrics by combining deep learning with distance metric learning [13]. Regarding this idea, given the rich latent space constructed with DML captures enough features from both classes on the thermal domain, flames that are obscured by smoke in the RGB domain will have corresponding thermal representations that fall in the positive (flame) class of the discriminator. Moreover, samples with no flames that have similar smoke patterns (local RGB patterns) to occluded flames will not appear similar to flames in the thermal domain.

Our proposed work uses bimodal data in RGB and thermal domains such that one modality (RGB) is used for sample annotation, while the other modality (thermal) offers latent features associated with the sample's classes, making inference possible with merely the second modality (thermal), and providing flame predictions for obscured areas. Such tasks fall under the umbrella of domain adaptation, where knowledge is transferred from one domain to the other [14]. Here, it should be mentioned that discrepancies between the source and target domain distributions can affect the performance of DML models on multimodal data [15], [16].

This study employs deep metric learning (DML) to enhance the model's performance in separating class clusters within the embedding space, specifically for detecting smoke-covered flames using paired RGB–thermal images. To facilitate the generation of data patches for fire spread modeling and monitoring, we introduce a segmentation and annotation framework (Section III-A). This framework labels a patch as “flame” if it contains at least one flame-annotated pixel and as “no-flame” otherwise. While many advanced domain adaptation methods offer innovative solutions to align features of the complementary domains, our simple indirect supervision approach benefits from the pairwise nature of patches in the source and target domains to make the convergence path smoother than unsupervised cases. After annotation is provided, a DML method is utilized for the annotated data, training a feature extractor to distinguish between flame and nonflame patches. A classifier is then employed to predict flames in smoke-obscured areas where no labels are available. The framework incorporates three loss functions in the DML framework to learn an optimal embedding function in the latent space. Recognizing the dominance of center loss (CL), particularly for magnitude-sensitive features, we propose an attention mechanism to balance feature contributions across the three DML loss gradients, harnessing the full potential of the constructed embedding space. This feature-selective loss smoothing enhances class discrimination in the latent feature space, addressing the issue of loss domination in multiobjective optimization. The main contributions of this work are as follows.

- 1) An improved thermal-based flame detection method onboard aerial vehicles using a clustering-guided feature representation learning over RGB and thermal images and domain adaptation in training to improve performance and reliability.
- 2) Proposing learnable class representatives optimized with stochastic gradient descent (SGD), resulting in enhanced discrimination of class-related latent features.
- 3) Proposing a novel attention mechanism to balance latent feature contribution in DML loss functions, resulting in diverse feature discovery.

II. RELATED WORKS

Supervised machine learning techniques are increasingly applied in wildfire detection. In [17], an approach combines local binary pattern (LBP) and SVM for smoke detection from UAV-captured RGB images and a CNN for flame detection with preprocessing steps such as histogram equalization and low-pass filtering. XtinguishNet [18] trains a CNN on IoT data for real-time wildfire detection, offering weather insights and fire intensity estimates. Nguyen et al. [19] presented a MobileNet-based single-shot detection (SSD) model for real-time UAV-based wildfire monitoring. Some other works employ heuristic masks such as using HSV domain filtering for flame-sensitive features [20]. Some other works focus on object-detection-based approaches with a simplified YOLO model [21] or MobileNet combined with YOLO-v4 [22]. Notably, none of these works explore multimodal inputs for obscured flame detection.

A. DML-Based Object Detection/Classification

DML-based object detection has become popular in computer vision for learning distance metrics as a path to learn rich feature representations [13]. In [23], the proposed method tunes the backbone network parameters to find an embedding function for multimodal data in a single training process. They use a subnet architecture for embedding function training and a multimodal mixture distribution for class posterior computation. In facial expression recognition, Rajoli et al. [13] improved classification accuracy using DML and triplet loss, with SGD optimizing class prototypes, backbone, and classifier parameters together. DML's popularity in few-shot learning is evident in [24], employing representation generation, and distance estimation modules beside DML for object representation, classification, and location estimation. Lu et al. [25] discussed decoupled metric network (DMNet) for single-object detection, introducing decoupled representation transformation (DRT) and image-level distance metric learning (IDML) to address representation disagreement and enhance generalization. DML-based methods excel in multimodal detection for capturing correlations and complex features across modalities [26]. However, challenges arise in handling high-dimensional feature spaces and maintaining a balance between dimensionality size and the model performance. The density of the embedding space influences generalization in DML-based models [27], while significant

dimension reduction may impact performance. Although various works contribute to enhancing DML-based approaches, this work pioneers the use of DML in flame detection, introducing an aggregated loss including triplet and cosine losses with CL to emphasize class discrimination aided by attention.

B. Attention in DML

Traditional DML models effectively learn embedding functions but struggle with dynamic focus on different aspects of the data and adapting to variable-length inputs [28]. Attention mechanisms in DML, as highlighted by Kim et al. [29], enhance feature embeddings and improve performance on image retrieval tasks by allowing each learner to attend to different parts of the input. Wang et al. [30] employed attention for feature balancing and class representation weighting, introducing class-aware attention (CAA) to identify noisy images and improve convergence in DML. Addressing DML challenges, Kotoenko et al. [31] introduced a cross-attention mechanism between image embeddings, establishing a hierarchy of conditional embeddings to enhance tuple representations. Seidenschwarz et al. [32] went beyond pairs and triplets in DML, proposing message-passing networks with an attention mechanism to weigh the importance of each neighbor in a mini-batch. Li et al. [33] introduced transformed attention consistency and contrastive clustering loss for robust similarity measures in DML. In several application domains, attention shows significant improvement for compressed feature embedding. In [34], mixed attention mechanisms aid histopathological image retrieval. Dong et al. [35] addressed the challenge of varying input sizes in DML with a hybrid channel and spatial attention approach for pavement distress classification. Coskun et al. [36] used an attentive RNN in human motion analysis to measure the similarity between motion patterns. Joint embedding of visual modalities, such as optical flows and RGB, is explored in [30] for near-duplicate video retrieval, employing attention modules to capture different levels of granularity in video-level representation. Specifically, for flame detection, the authors in [37] used a dynamic attention strategy that focuses on scale-aware and spatial-aware features. Recent works such as [38] integrate Transformer encoder blocks with multihead attention to better capture global and contextual feature information.

III. PROPOSED METHOD

This section describes our approach for flame detection and classification (Fig. 1) in obscured cases. A detailed description of the symbols used in the upcoming equations is available in the Nomenclature section. The objective is to learn the relationship between flame patterns in RGB and their corresponding thermal representation. The process involves three main steps.

- 1) *Unobscured Flame Annotation*: Annotating visible flames in RGB images using Algorithm 1.
- 2) *DML-Guided Learning*: Training an autoencoder and a classifier with DML on labeled thermal images (Algorithm 2).

Algorithm 1 FlameSeg: RGB-Unobscured Flame Segmentation Using Image Processing

Inputs:

- $Img_{[R,G,B]}$ (input images),
- R_m, G_m, B_m (red, green, blue channel means),
- τ_g, τ_b (green and blue channel value threshold),
- α, β (red-to-green low and high margin coefficients)

Output: Flame segmentation mask m .

1. filtering based on the distance to the mean:

$$m_1(x, y) \leftarrow \begin{cases} 0, & Img_{[R]} - R_m < Img_{[B]} - B_m \\ 0, & Img_{[R]} - R_m < Img_{[G]} - G_m \\ 1, & \text{Otherwise} \end{cases}$$

2. filtering based on the value:

$$m_2(x, y) \leftarrow \begin{cases} 0, & Img_{[G]} > \tau_g \\ 0, & Img_{[B]} > \tau_b \\ 1, & \text{Otherwise} \end{cases}$$

3. filtering based on the inter-channel distance:

$$m_3(x, y) \leftarrow \begin{cases} 0, & |Img_{[R]} - Img_{[G]}| \leq \alpha Img_{[G]} \\ 0, & |Img_{[R]} - 2 * Img_{[G]}| \geq \beta Img_{[R]} \\ 0, & |Img_{[G]} - 2 * Img_{[B]}| \leq \delta Img_{[G]} \\ 1, & \text{Otherwise} \end{cases}$$

4. applying the masks, removing single pixels, and smoothing

- 3) *DML-Guided Inference*: Using the trained model for obscured thermal patches. An attention mechanism balances the role of loss functions across features of the latent space.

A. Unobscured Flame Annotation

In response to the scarcity of annotated RGB-thermal image pairs for wildfire segmentation, we propose a tailored filtering approach for unobscured flame detection in top-down UAV-captured RGB images. The RGB-unobscured flame segmentation algorithm (Algorithm 1) creates a segmentation mask based on heuristic channel value inequalities, which can be later value. Next, smoothing and pixel-removal operations are applied, and resulting patches are labeled as “flame” or “no-flame,” providing indirect supervision for learning to detect flames in unannotated patches during the inference phase. This strategy aims to overcome challenges associated with the lack of absolute thermal reference points, mitigating misclassification in thermal images due to relative temperature gradients. It is worth highlighting the fact that this simple, yet automated segmentation is prone to error, and as a result, label noise may be limiting the performance of the model in terms of classification accuracy. A comparative analysis of the effect of segmentation noise is done in Section III-A3 to further justify the automated segmentation.

1) *RGB Mask Generation*: In light of the approaches proposed by Buza and Akagic [39] and Dzigal et al. [40] for forest fire detection, we have developed a well-fit masking method to cater to the specific type of fire encountered in real-world situations, namely, spot flames captured with a top-down perspective by UAVs. Algorithm 1 explains our




Fig. 1. (a) Training: simultaneously learning the encoder, decoder, embedding space, and classifier. The class's prototype and the embedding function are jointly trained, while each feature's loss contribution is balanced with the attention module. Finally, the thermal patches are predicted as flame/no-flame. (b) Inference: obscured thermal images are introduced to the detector, and based on the classifier's output, the corresponding cluster (patch label) is predicted.

approach for detecting unobscured flames in RGB images. This algorithm sets a criterion based on channel means, values, and interchannel distances to generate a flame segmentation mask. The combination of three masks representing RGB criteria, reflecting scene-specific features such as the dominant green background in forest fire images, facilitates the identification of flames amidst the smoke. We calculate a flame segmentation mask by combining three masks, each representing an RGB criterion for candidate flames. The reason for choosing such criteria lies in some scene-specific features, such as the dominant green texture of the background in forest fire images. In addition, the combination of green and red channels creates orange and yellow tones that are the dominant colors in the observed flames. Fig. 2 demonstrates the segmentation results followed by smoothing and pixel-removal operations.

2) *Thermal Patch Labeling*: Eventually, any patch containing part of a detected flame in the output image will be labeled positive as “flame,” and all other patches will be labeled negative as “no-flame.” All equivalent thermal patches of the pair thermal image will be annotated likewise. To demonstrate the effectiveness of the algorithm, a set of such patches is pixel-annotated by a human expert, and the corresponding results are presented in Section IV-B. It should be noted that the mask is not used as an annotation module in the inference phase. We intend to classify smoky patches as flame/no-flame where no reliable class is provided. Thus, no labels are given or produced for the inference phase. In other words, the mask provides indirect supervision for a dataset of unannotated patches, in which the associated side-by-side thermal features of flames are learned to detect flames for classification on smoky patches in the inference phase. This indirect thermal patch labeling aims to avoid using the relative thermal spectrum captured in thermal images. This issue in




Fig. 2. Unobscured flame segmentation. Sample images demonstrating the performance of the proposed unobscured flame segmentation method in identifying flame instances covered by smoke.

thermal imagery is rooted in the lack of absolute thermal reference points. As a result, thermal values indicate the relative temperature of an area compared to its surroundings. It should be noted that in the case of using thermal images directly, any other local temperature gradient can lead the patch to be misclassified as a flame. In conclusion, the joint usage of RGB and thermal modalities in segmentation and classification, respectively, surpasses the discussed issues raised in unimodal RGB/thermal detection. It is noteworthy that for optimal segmentation performance, the threshold values (τ_g , τ_b) and parameters (α , β), which are only empirical hyperparameters,

need to be fine-tuned to accommodate for misty flames arising due to smoke.

3) *Comparative Annotation Analysis*: For a set of 8640 patches, originating from 135 images, the proposed annotation mechanism (Section III-A) was able to identify 81% of the flames detected by an expert. One main concern would be the effect of mislabeled patches on the final performance of the model. To be more specific, when a model is trained with noisy annotated data because the inference phase performance will also be measured with noisy labels, the annotation noise will be hindered by the good results, while the model may not perform well in reality. To address this concern and measure the annotation noise effect, a comparison between the inference phase performance should be done between the two cases of expert-annotated and automated annotated data. As training is impractical with expert-annotated data due to the large volume of required data, we have conducted a simple experiment, where we tested the fine-tuned model (on automated annotation) on the expert-labeled data. Details of the performance are shown in the confusion matrix presented in Table I. Class imbalance is due to the very few number of flame patches extracted from no-flame images, whereas many no-flame patches are present in flame-labeled images. As expected, the accuracy decreases on expert-annotated data as the model predicts several instances of flames (misannotated due to annotation noise) as “no-flame” and vice versa.

Following the performance degradation on expert-annotated data in terms of the classification accuracy (82.9% in Table I, compared to 97.8% in Table IV), the following facts should be kept in mind when utilizing the automated annotation algorithm proposed in Algorithm 1.

- 1) Some flames may be hindered by an opaque object such as a tree branch. Despite there being limited to the currently available datasets, future works may use consecutive frames of a video, multiview imaging, or other techniques to reduce the chances of object obfuscation.
- 2) Although the expert annotation is more accurate, the fact that manual annotation of the whole dataset is time-consuming due to a large number of final patches prevents us from a more comprehensive comparative study between manual and automated annotation.
- 3) In some cases, the magnitude spatial correlation between thermal and RGB patterns and also the correlation between thermal patterns in obscured versus unobscured patches may be reduced due to various reasons such as smoke mobility, the proximity of smoke to the camera, and so on. This label noise may increase the misclassified instances, as seen in Fig. 8. However, most such cases are unavoidable while training and should be detected and controlled with equipment in the data collection phase (e.g., utilizing more modalities in the field can prevent the model learn noisy associations and boost the inference accuracy.)
- 4) There may be algorithms with much higher accuracy than the proposed heuristic approach, but this approach suits the real-time implementation limitations for computational efficiency as it only serves as the preprocessing (annotating) part of the design.

TABLE I
CONFUSION MATRIX OF THE FlameSeg ALGORITHM. PREDICTIONS MADE BY THE ALGORITHM, COMPARED TO EXPERT-LABELED DATA (GROUND TRUTH)

		Predicted		
		Flame	No-Flame	Total
Actual	Flame	2093	739	2832
	No-Flame	643	4625	5268
Total		2736	5364	8100
Precision (Flame)		76.5 %		
Recall (Flame)		73.9 %		
Accuracy (Flame)		82.9 %		

B. DML Loss Derivatives

Here, the derivatives of the loss functions used in DML losses are shown in the following equations:

$$\frac{\partial \mathcal{L}_{TL}}{\partial e} = \frac{1}{m} \sum_{x^{(i)} \in \mathcal{K}} \sum_{j=1}^{c_d} a_{ij}^{TL} \left(\left(c_{y^{(i)}j}^n \right) + \left(c_{y^{(i)}j}^p \right) \right) H(\mathcal{L}_{TL}) \quad (1)$$

$$\frac{\partial \mathcal{L}_{TL}}{\partial a_{TL}} = \left(\frac{1}{2m} \sum_{i=0}^m \sum_{j=1}^{c_d} \left(\left\| \left(e_j^{(i)} \right) - \left(c_{y^{(i)}j}^p \right) \right\|_2^2 - \left\| \left(e_j^{(i)} \right) - \left(c_{y^{(i)}j}^n \right) \right\|_2^2 + \alpha \right) \right) H(\mathcal{L}_{TL}) \quad (2)$$

where $H(\mathcal{L}_{TL})$ represents the Heaviside step function. It should be noted that in (1)–(6), a_{ij}^{TL} , a_{ij}^{CL} , $e_j^{(i)}$, $c_{y^{(i)}j}^p$, and $c_{y^{(i)}j}^n$ are scalars, whereas a_i^{CosL} , $e^{(i)}$, $c_{y^{(i)}}^n$, and $c_{y^{(i)}}^p$ are vectors with length c_d . Plus, the “ $\odot$ ” and “ $\cdot$ ” symbols used in (1)–(3) denote element-wise multiplication and inner product operations, respectively,

$$\begin{aligned} \frac{\partial \mathcal{L}_{CosL}}{\partial e} = & \frac{-1}{2m} \sum_{i=1}^m a_i^{CosL} \\ & \odot \left(\left(\frac{c_{y^{(i)}}^p}{\|e^{(i)}\| \|c_{y^{(i)}}^p\|} - \frac{(e^{(i)} \cdot c_{y^{(i)}}^p) e^{(i)}}{\|e^{(i)}\|^3 \|c_{y^{(i)}}^p\|} \right) \right. \\ & \left. - \left(\frac{c_{y^{(i)}}^n}{\|e^{(i)}\| \|c_{y^{(i)}}^n\|} - \frac{(e^{(i)} \cdot c_{y^{(i)}}^n) e^{(i)}}{\|e^{(i)}\|^3 \|c_{y^{(i)}}^n\|} \right) \right) \end{aligned} \quad (3)$$

$$\frac{\partial \mathcal{L}_{CosL}}{\partial a_{CosL}} = -\frac{1}{2m} \sum_{i=0}^m \sum_{j=1}^{c_d} \left(\frac{(e_j^{(i)} c_{y^{(i)}j}^p)}{\|e^{(i)}\| \|c_{y^{(i)}}^p\|} - \frac{(e_j^{(i)} c_{y^{(i)}j}^n)}{\|e^{(i)}\| \|c_{y^{(i)}}^n\|} \right) \quad (4)$$

$$\frac{\partial \mathcal{L}_{CL}}{\partial e} = \frac{1}{m} \sum_{i=1}^m \sum_{j=1}^{c_d} a_{ij}^{CL} \left(\left(e_j^{(i)} \right) - \left(c_{y^{(i)}j} \right) \right) \quad (5)$$

$$\frac{\partial \mathcal{L}_{CL}}{\partial a_{CL}} = \frac{1}{2m} \sum_{i=1}^m \sum_{j=1}^{c_d} \left\| \left(e_j^{(i)} \right) - \left(c_{y^{(i)}j} \right) \right\|_2^2. \quad (6)$$

Algorithm 2 DML-Aided Training**Inputs:**

- Batch of labeled cropped patches as training dataset, $D_t = \{(x^{(i)}, y^{(i)}) | i = 1, 2, \dots, N_t\}$,
- Initialized encoder-decoder network parameters (θ_{en}, θ_{de} , and θ_{cls}),
- Reconstruction loss, DML, and softmax loss parameters ($\theta_{re}, \theta_{DML}, \theta_s$),
- Class (cluster) prototype, $c = \{c_k | k = 1, 2; c_k \in \mathbb{R}^{1 \times C_d}\}$,
- Hyperparameters ($\lambda_1, \lambda_2, \gamma_1, \alpha$),
- $t \leftarrow 0$.

Output:

- Trained model and updated $\theta_{en}, \theta_{de}, \theta_{re}, \theta_s, \theta_{cls}$, and C .

While not converged **do**:

- 1 Feed the input dataset minibatch $B_m = \{(x^{(i)}, y^{(i)}) | i = 0, 1, \dots, m\}$ into the network.
- 2 In feature space, compute the corresponding embeddings $e \in \mathbb{R}^{C_f \times 1 \times 1}$.
- 3 Compute the CE loss utilizing (7).
- 4 Compute the AE loss utilizing (8).
- 5 Compute the DML loss utilizing (10) to (13).
- 6 Compute loss function gradient using (1) to (6) and chain rule.
- 7 Compute Δ_C using (14) to (16).
- 8 $t \leftarrow t + 1$.
- 9 $C^{t+1} \leftarrow C^t - \alpha \Delta_C$.

End while**C. DML-Guided Learning**

During the training phase (Algorithm 2), it is crucial to sample inputs from smoke-free images to ensure a comprehensive exploration of flame representations in the thermal domain. We oversample the flame patches to address the class imbalance in the dataset and weigh the loss accordingly. Our final dataset comprises cropped patches from thermal images labeled as either “flame” or “no-flame.”

The proposed framework comprises an encoder and a decoder (based on the U-net autoencoder architecture [41]), a classifier, and a DML module that work together to learn features of the input data through a multitask loss function. The loss function is a weighted sum of three terms: cross-entropy (CE) classification loss, autoencoder reconstruction loss, and the DML losses described afterward. Let $D_m = \{(x^{(i)}, y^{(i)}) | i = 1, 2, \dots, m\}$ be a mini-batch of m samples, where $x^{(i)} \in \mathbb{X}$ represents pairs of thermal patches and their corresponding flame masks from the training set, and $y^{(i)} \in \{1, 2\}$ is the corresponding categorical label (flame or no-flame). At the output of the encoder, the sample representation is shown by $x^* \in \mathbb{R}^{C_f \times H_f \times W_f}$, where C_f , H_f , and W_f are the number of channels, rows, and columns of the encoded samples, respectively. A convolutional layer and a global average pooling layer follow the encoder. The encoder maps the encoded samples of the mini-batch to their feature space representation, denoted by $e \in \mathbb{R}^{C_d \times 1 \times 1}$, where C_d is the

number of channels (the length of the feature embedding vector).

A fully connected (FC) layer in the classifier network maps the embedding to logits of $x^z \in \mathbb{R}^2$. The Softmax function is applied to obtain a probability distribution $\Pr(y = j | x^{(i)})$. The discrepancy between the predicted label $\hat{y}$ and the true label $y^{(i)}$ is calculated using the CE loss function, denoted by $\mathcal{L}_{BCE}$, as follows:

$$\mathcal{L}_{BCE} = -\frac{1}{m} \sum_{i=1}^m \sum_{j=1}^2 y^{(i)} \log \Pr(y^{(i)} = j | x^{(i)}). \quad (7)$$

By implying the reconstruction loss, we also ensure the latent representation learned by the encoder, which is core to classifying and clustering, contains the essential features of the input domain, preserving the input information as much as possible

$$\mathcal{L}_{rec} = \frac{1}{m} \sum_{i=1}^m \|x^{(i)} - f(g(x^{(i)}))\|^2 \quad (8)$$

where m is the number of examples in the dataset, $x^{(i)}$ is the i th input example, and g and f are the encoder and the decoder, respectively.

1) *Multitask Loss Function*: The multitask loss function aims to optimize multiple objectives simultaneously. This loss function includes the direct CE ($\mathcal{L}_{BCE}$) loss for classification, the reconstruction loss for an autoencoder ($\mathcal{L}_{rec}$), and the distance loss for DML ($\mathcal{L}_{DML}$), which considers the distance of each sample’s latent representation to the cluster representatives

$$\mathcal{L} = \mathcal{L}_{BCE} + \lambda \mathcal{L}_{rec} + \mathcal{L}_{DML}. \quad (9)$$

2) *DML Metrics*: Triplet loss optimizes anchor, positive, and negative sample distances, along with emphasizing discriminative features, while cosine similarity is based on feature vector angle suiting high-dimensional spaces where vector magnitude is less significant. Combining triplet loss and cosine distance metrics can improve feature space representation resulting in the model’s improved classification as well as enhancing robustness to illumination variations and other factors affecting flame appearance, leading to enhanced generalization to unseen joint distributions in modalities, and increasing intraclass variance. Note that in (13), we define a novel form of cosine distance (denoted by $\mathcal{L}_{CosL}$) inspired by the concept of triplet loss ($\mathcal{L}_{TL}$). It is computed based on the average similarity between the sample embedding’s features ($e_j^{(i)}$) and the same feature of the sample’s corresponding class prototype ($c_{y^{(i)}}^p$) utilized by (11). In (11) and (13), the superscripts p and n denote the current samples’ corresponding class prototype and the opposite class prototype, respectively. At last, to force the representative to be close to the distribution of the clusters’ embeddings, we apply another DML, called CL [denoted by $\mathcal{L}_{CL}$ (attentive)]. However, our experiments reveal that this unmodified version of $\mathcal{L}_{centL}$ is too dominant and hinders the positive contribution of $\mathcal{L}_{TL}$ and $\mathcal{L}_{CosL}$ in the learning process. As a result, the model will focus on discrimination in feature magnitudes, pushing representatives out of the cluster space to increase their difference magnitude.

To address this concern, we integrate the deep attentive CL ($\mathcal{L}_{CL}$), as proposed by Zachariadis et al. [42]. It is noteworthy that the adapted CL in (12) is employed to alleviate the prevailing constraints of the vanilla CL (a_{ij} represents feature attention weights). The attention mechanism enables each of the loss functions to scale features of the cluster representatives in distance calculation and thus modifies the trajectory of the representative in the latent space, maximizing interclass separation and minimizing intraclass separation. The calculation of a_{ij}^{TL} , a_{ij}^{CL} , a_{ij}^{CosL} in (11)–(13) is shown in the attention module in (Fig. 1). In these equations, the nonattentive (vanilla) implementation is equivalent to setting all attention coefficients to 1 ($a_{ij}^{TL} = a_{ij}^{CL} = a_{ij}^{CosL} = 1$). Moreover, all norms used are Euclidian norms. It should be noted that a_{ij}^{TL} , a_{ij}^{CosL} , a_{ij}^{CL} , $e_j^{(i)}$, and $c_{y^{(i)}j}^p$ [used in (11)–(13) and later in (14) and (16)] are all scalars, whereas a_{ij}^{CosL} , $e^{(i)}$, $c_{y^{(i)}}^n$, and $c_{y^{(i)}}^p$ [used in (15)] are vectors with length c_d , and the “ $\odot$ ” and “ $\cdot$ ” symbols used in (15) denote element-wise multiplication and inner product operations, respectively,

$$\mathcal{L}_{DML} = \gamma_1 \mathcal{L}_{TL} + \gamma_2 \mathcal{L}_{CosL} + \gamma_3 \mathcal{L}_{CL} \quad (10)$$

$$\mathcal{L}_{TL} = \max \left(0, \frac{1}{2m} \sum_{i=0}^m \sum_{j=1}^{c_d} \left[a_{ij}^{TL} \left(\left\| (e_j^{(i)}) - (c_{y^{(i)}j}^p) \right\|_2^2 - \left\| (e_j^{(i)}) - (c_{y^{(i)}j}^n) \right\|_2^2 + \alpha \right) \right] \right) \quad (11)$$

$$\mathcal{L}_{CL} = \frac{1}{2m} \sum_{i=0}^m \sum_{j=1}^{c_d} a_{ij}^{CL} \left\| (e_j^{(i)}) - (c_{y^{(i)}j}^p) \right\|_2^2 \quad (12)$$

$$\text{s.t. } 0 < a_{ij} \leq 1 \quad \forall j | j \in \{1, \dots, c_d\}$$

$$\mathcal{L}_{CosL} = -\frac{1}{2m} \sum_{i=0}^m \sum_{j=1}^{c_d} a_{ij}^{CosL} \left[\left(\frac{(e_j^{(i)}) (c_{y^{(i)}j}^p)}{\| (e^{(i)}) \| \| (c_{y^{(i)}}^p) \|} - \frac{(e_j^{(i)}) (c_{y^{(i)}j}^n)}{\| (e^{(i)}) \| \| (c_{y^{(i)}}^n) \|} \right) \right] + 2. \quad (13)$$

3) *DML Metrics Derivatives*: Our proposed method (Algorithm 2) trains the integral parameters of the loss functions, including the cluster representatives, in an end-to-end manner during supervised training using SGD. In this regard, (1)–(4) demonstrate the required closed-form derivatives for performing gradient descent. The closed-form derivatives of (11) described in (1) and (2) is required as the cluster representatives are getting updated every step and build the foundation for the inference phase. An important fact is that the computational graph cannot yield the derivatives of the loss functions concerning the representative; thus, closed-form derivatives are required for the moving-average centroid updates described in (14), (15), and (16). Parameter k in (1) is the number of samples that do not satisfy the margin defined in (11), and so they contribute to learning, in other words, $k = \text{card}(\mathcal{K})$ where, $\mathcal{K} = \{x^{(i)} | \mathcal{L}_{TL}(x^{(i)}) > 0\}$. Inspired by Zachariadis et al. [42], we follow a moving average strategy to update the class/cluster representatives, C . In this context,

we divide the set $\mathcal{K}$ into two subsets of flame samples $\mathcal{K}_1$ and nonflame samples $\mathcal{K}_2$ given that $\mathcal{K}_1 \cup \mathcal{K}_2 = \mathcal{K}$. In the same way, $\mathcal{K}' = \{x^{(i)} | \mathcal{L}_{CosL}(x^{(i)}) > 0\}$, and $\mathcal{K}'_1$ and $\mathcal{K}'_2$ are defined as subsets of $\mathcal{K}'$ that include flame and no-flame samples, respectively. Following the same notation, $C_{TL}^{K_1}$ denotes the flame cluster representative, while $C_{TL}^{K_2}$ denotes the no-flame cluster representative. It should be noted that both flame and no-flame samples affect the cluster representative gradient

$$\begin{aligned} \Delta_{C_{TL}^{K_1}} &= \frac{-1}{k} \left(\sum_{x^{(i)} \in \mathcal{K}_1} \sum_{j=1}^{c_d} a_{ij}^{TL} \left((e_j^{(i)}) - (c_{y^{(i)}j}^p) \right) - \sum_{x^{(i)} \in \mathcal{K}_2} \sum_{j=1}^{c_d} a_{ij}^{TL} \left((e_j^{(i)}) - (c_{y^{(i)}j}^n) \right) \right) \quad (14) \\ \Delta_{C_{CosL}^{K'_1}} &= \frac{-1}{2k'} \left(\sum_{x^{(i)} \in \mathcal{K}'_1} a_i^{CosL} \odot \left(\frac{e^{(i)}}{\|e^{(i)}\| \|c_{y^{(i)}}^p\|} - \frac{(e^{(i)} \cdot c_{y^{(i)}}^p) c_{y^{(i)}}^p}{\|e^{(i)}\| \|c_{y^{(i)}}^p\|^3} \right) - \sum_{x^{(i)} \in \mathcal{K}'_2} a_i^{CosL} \odot \left(\frac{e^{(i)}}{\|e^{(i)}\| \|c_{y^{(i)}}^n\|} - \frac{(e^{(i)} \cdot c_{y^{(i)}}^n) c_{y^{(i)}}^n}{\|e^{(i)}\| \|c_{y^{(i)}}^n\|^3} \right) \right) \quad (15) \end{aligned}$$

$$\Delta_{C_{CL}} = \frac{-1}{m} \sum_{i=1}^m \sum_{j=1}^{c_d} a_{ij}^{CL} \left((e_j^{(i)}) - (c_{y^{(i)}j}) \right) \quad (16)$$

where in (14) and (15); $(l, q) \in \{(1, 2), (2, 1)\}$.

D. DML-Guided Inference

Once the autoencoder is trained, the decoder is removed, leaving the classifier. The thermal-domain cluster prototypes (Fig. 1) represent classes for “flame” and “no-flame” in the latent feature space. Back in training, the classifier had received the latent feature vector of a sample, while DML losses had made the embedding function generate more separable clusters and better prototypes. The unsupervised metrics used in supervised training exist to improve generalization for obscured flames with thermal patterns (where no ground truth is available). After training the optimal cluster representative (Algorithm 2), the model predicts the class for obscured samples. The cluster representatives trained by the proposed DML approach set the prediction criteria (with fixed attention weights) and test samples are compared to each representative to make predictions. It should be noted that only thermal patches are used in the inference phase and the predicted labels for each thermal patch will eventually correspond to the equivalent RGB patch, revealing possible flames hindered by the smoke (Fig. 3).

As there are no patch labels to evaluate the model’s performance on the test set, we use intraclass variability in the




Fig. 3. Obscured fire detection during inference. Only thermal patches are used for inference. The predicted labels correspond to the equivalent RGB patches and reveal potentially obscured flames.

flame “class” as an unsupervised indicator to show how DML and attention make clusters as separable as possible by moving the prototypes and shaping the embedding function (Table III). Higher separation suggests that RGB-domain labels are good indicators of thermal-domain features, indicating successful knowledge transfer. It should be noted that successful knowledge transfer relies on the fact that obscured and unobscured thermal-domain representations of flame samples are significantly correlated. Fig. 4 shows two frames captured 7-s apart, from FLAME2, showing the same flame, but with different smoke occlusion densities and patterns around the flame. RGB representation looks very different due to smoke occlusion, yet the IR domain greatly preserves the distribution, showing class-specific feature transparency through the smoke. In addition, RGB representations of the obscured images look pretty similar (low interclass variance), and class-specific features are obscured by smoke. Thus, using RGB along IR for inference acts as noise, making class separation more difficult. This is the motivation behind using only thermal-domain patches for inference.

E. Local Features Extraction (LFE) Mechanism

To capture distinctive features from specific subareas of an image, both the thermal image, denoted as $x^{(i)}$, and the masked thermal image, represented as $x_m^{(i)} = x^{(i)} \odot m^{(i)}$ (where $\odot$ signifies the Hadamard product (element-wise multiplication) and $m^{(i)}$ is the corresponding mask generated by Algorithm 1), are fed into two U-Nets with shared parameters. Subsequently, the local and global features are aggregated within the latent space as illustrated in Fig. 1. This approach enables the model to focus on masked subareas containing intricate patterns and textures associated with flame spots. Given that wild-fires’ early-stage flames typically manifest as dispersed spots, this methodology facilitates a detailed understanding of such regions through targeted analysis of the masked portions.

IV. PERFORMANCE EVALUATION

In this section, first, we will cover dataset partitioning, preprocessing, and training. Next, we will discuss the results obtained from training, validation, hyperparameter tuning, and testing in terms of supervised and unsupervised metrics.




Fig. 4. Thermal-domain spatial distribution remains nearly unchanged after smoke obscures a flame patch and its surroundings. RGB representations are, however, very different due to different smoke patterns, even in frames with small time separations.

TABLE II

DATASET DETAILS AND EXPERIMENT SETTINGS (* FLAME 3 STATS ARE NOT FINAL. ** A PORTION OF FLAME2 SAMPLES ARE UTILIZED DURING TRAINING AND TESTING. *** THE SHOETANK’S FIRE SAMPLES OF FLAME3 ARE USED)

Dataset Details	FLAME 2	FLAME 3*
# RGB/Thermal Img. Pairs	53k	20k
Experiment Settings	FLAME 2	FLAME 3
# Total Samples (Train)	4700**	880***
# Total Samples (Test)	1170	220
# Batches/Sample	64	64
# Original Sample Resolution	256 × 256	512 × 640
# Input Resolution	256 × 256	256 × 256
# Batch Resolution	32 × 32	32 × 32

A. Dataset

The FLAME2 [43] and FLAME3¹ datasets have been used for evaluation. FLAME2 includes seven side-by-side infrared and visible spectrum video pairs captured by drones from open-canopy fires in Northern Arizona in 2021. Derived from the video pairs, a set labeled original resolution RGB/thermal frame pairs and a set of 53 000 labeled 254p × 254p RGB/thermal frame pairs are provided. FLAME 2 provides two labels of fire/no-fire and smoke/no-smoke for every image. Experts have assigned fire labels based on visual fire indicators in the RGB domain. Moreover, they have assigned “smoke” and “no-smoke” labels to every RGB image, resulting in 14 760 smoke-free frames. The RGB frames are the source of visual labeling, while thermal images are provided along them to extract the thermal features of the supervised set. The corresponding details for FLAME2 and FLAME3 and the experimental settings are shown in Table II.

B. Implementation

In our work, we used 4700 frames, or one-third of the smoke-free samples of the labeled RGB/thermal image pairs in FLAME2 and 1200 samples of FLAME3 to diversify the dataset and avoid redundant frames. We start preprocessing by upsampling frames to 256p × 256p resolution. Next, we divide each image into 64 32p × 32p patches. This applies to RGB and thermal images. After that, we segment fire-containing

¹FLAME3 dataset is collected by a team of universities, U.S. Forest Services, and CALFIRE and is not publicly available.

TABLE III
PERFORMANCE COMPARISON AGAINST OTHER STUDIES. (ICV:
INTRACLASS VARIANCE, FNR: FALSE NEGATIVE RATE).
METRICS FOR FLAMES 2 AND 3 ARE REPORTED ON
THE LEFT AND RIGHT, RESPECTIVELY

Dataset	Model	Accuracy	False Negative	Intra-Class Variance
Flame2	ResNet18	93.70	17.00	48.70
	VGG19	92.90	17.40	49.70
	ShRe Xception [45]	93.50	18.60	50.30
	Ensemble model [46]	94.20	14.10	46.60
	CT-Fire [47]	94.80	13.50	43.00
	Ours	97.80	11.60	39.70
Flame3	ResNet18	81.40	19.10	50.90
	VGG19	80.50	18.80	49.20
	ShRe Xception [45]	81.90	18.20	49.00
	Ensemble model [46]	83.80	16.80	47.40
	CT-Fire [47]	85.10	16.90	45.10
	Ours	88.90	14.30	41.20

RGB images using Algorithm 1. The segmentation parameters were manually set to $\{\alpha = 0.1, \beta = 0.47, \delta = 0.14\}$. The annotation accuracy on smoke-free patches is compared to pixel-annotated patches labeled by an expert. The segmentation is compared to a human-expert perfect baseline performed on 8640 patches originating from 135 images, through which the segmentation algorithm was able to identify 81% of the patches annotated as “flame” by the expert. The gap is assumed to be mainly due to the presence of obscure and partially ignited patches seen by the expert and rejected by the RGB mask. While not achieving human expert-level performance, Algorithm 1 offers an automated segmentation for large-scale datasets. The comparative analysis between manual and automated labeling is present in Section III-A3 featuring Table I. Next, annotated thermal patches are oversampled to create a balanced dataset of “flame” and “no-flame” patches. Data augmentation is performed by random flipping vertically and horizontally and normalizing. Our model’s encoder, decoder, and classifier were trained using the procedure shown in Fig. 1. The FLAME2 dataset was used for 20 epochs of training with an initial learning rate of 0.01 for both $\mathcal{L}_{BCE}$ and $\mathcal{L}_{rec}$. The learning rate decreased by 0.1 at epochs 10 and 18 for $\mathcal{L}_{BCE}$ and $\mathcal{L}_{rec}$, respectively. Each batch of thermal patches had 512 images to train the model. Empirical tuning yielded the following training hyperparameter values: $\lambda = 1$, $\gamma_1 = 0.01$, $\gamma_2 = 0.0001$, and $\gamma_3 = 0.01$. In addition, we set DML learning rates as follows: $\alpha_{TCL} = 0.5$, $\alpha_{CL} = 0.5$, and $\alpha_{CosL} = 0.1$. Our models are trained using the PyTorch framework on an NVIDIA 3080 GPU.

C. Latent Scattering Demonstration

This study reduces latent space representation dimensionality using t-distributed stochastic neighbor embedding (t-SNE) in Fig. 7. The model is guided to the optimal subspace by the DML losses. The 2-D plots in Fig. 5 illustrate the effects of the proposed modules. The DML-guided model subfigure shows significant improvement over the non-DML baseline model. Fig. 1 highlights misclassifications in some patches due




Fig. 5. Visual performance comparison of the proposed method with baseline models. The gradient magnitude of the last convolutional layer, based on the Grad-CAM idea in [44] is upsampled to the image size and stacked on the sample image. P and L denote prediction and label, respectively.

to challenges posed by the nonoptimal annotation mechanism in Algorithm 1, making the discrimination between flame and no-flame classes in RGB space challenging. Misclassifications in flame patches arise from the lack of temperature calibration in the FLAME2 dataset, resulting in subclusters, the similarity between smoke and flame thermal distributions, and temperature gradient localized wind being generated and interacting with smoke patterns, moving them to areas far from the flame.

D. Numerical Results

In this section, we present a thorough numerical analysis to: 1) compare our model with previously proposed models for

TABLE IV

ABLATION STUDY, BINARY FLAME DETECTION ACC. FOR LFE AND ATTENTION, AND TRIPLET COSINE AND COSINE LOSSES. 1, BOTH ADDING DML AND LFE SHOW ACCURACY IMPROVEMENT, 2. TRIPLET AND COSINE LOSSES IMPROVE DML AFTER ADDING ATTENTIVE CL

Dataset	Detection Method	No DML	DML Ablation		
			CL	CL & TCL	CL & TCL & CosL
Flame2	Baseline	93.40	95.90	95.70	95.50
	Baseline + LFE	95.00	96.10	96.10	95.90
	Baseline + LFE + Attention	Not Applicable	97.20	97.50	97.80
Flame3	Baseline	81.9	84.40	85.10	84.70
	baseline + LFE	83.80	86.00	85.90	85.80
	baseline + LFE + attention	Not Applicable	88.10	88.80	88.90

flame detection as well as VGG19 and ResNet18 backbones adapted for flame detection and 2) evaluate how DML, LFE, and attention aid in boosting the classification accuracy as well as the impact of adding triplet and cosine losses to CL before and after attention.

1) *Comparative Analysis*: As seen in Table III, our model achieves higher accuracy, and less false negative rate (miss) compared to all other backbones, including VGG-19, ResNet-18, and three other works specified for flame detection. Our model also achieves a lower intraclass variance showing the compactness of the latent clusters as an unsupervised metric.

2) *Ablation Study*: Table IV shows an ablation study for DML, LFE, and attention as well as triplet and cosine losses. Samples acquiring less than 90% confidence in one of the two classes are omitted from the prediction. For the baseline, the cluster representatives are calculated by averaging all feature vectors in the latent space. According to Table IV, DML improves the binary classification accuracy. After adding the attention mechanism to overcome the CL domination problem, each latent feature contribution toward the proposed loss functions is balanced, demonstrating 1.9% and 3.1% accuracy improvement compared to the nonattentive version, for FLAME2 and FLAME3, respectively. Lower accuracy on the latter may be due to more diverse samples with variety in flame type, illumination, vegetation, and so on.

E. Interpretability

With the aid of two main tools presented in this section, the attentive classifier's results are explainable, and the contribution of all DML losses and global feature extraction is showcased.

- 1) *Grad-CAM*: As seen in Fig. 5, in the baseline models, according to the Grad-CAM [44], most of the learning is focused on high-intensity regions in the thermal domain, whereas most flame features in our model, especially with attention, are nonintensity-based flame features, covering around the flame. This aligns with global feature discovery and acts as a solution to the lack of absolute thermal reference points decreasing false positives.
- 2) *Attention Map*: Fig. 6(a)–(c) depicts how each of the 64 features contributes to the aggregated loss gradient.




Fig. 6. Comparison of attention values of each loss across the 64-dimensional latent feature vector. (a) Cosine-based loss attention, (b) triplet loss attention, and (c) CL attention. Attention values are indicators for including and excluding features' contribution in the DML losses value, (11)–(12), and in the backpropagation process, (1)–(4): The weight sparsity of the CL (e.g., features 9, 36, 45, and 64) shows the attention mechanism adaptively regularizes the contribution of every extracted latent feature in the loss functions, guiding the model toward the optimum in the parameter space (instead of scaling loss function derivatives uniformly with the hypertuning parameters and causing center-loss domination).




Fig. 7. 2-D demonstration of the clustered latent feature space. (a) Baseline latent scattering and (b) DML-guided latent scattering. Dimension reduction via t-SNE approach.

The intensity of each band within the map corresponds to the attention weight of the feature associated with that band. The sparsity of features in the CL attention map [lower weights—Fig. 6(c)] shows the attention mechanism successfully overcomes the domination problem and refines the model toward capturing nonmagnitude




Fig. 8. Classified and misclassified patches during validation.

features in the latent space, provoking the cosine and triplet loss, which appear quite correlated [Fig. 6(a) and (b)], in the learning process.

V. CONCLUSION AND FUTURE WORK

FlameFinder offers a comprehensive solution encompassing segmentation, DML-guided classification, and aided attention, leveraging the power of DML as a significant driving force for detecting obscured flames. The results achieved from the predictions of the DML-based model demonstrate favorable performance in terms of separating flames in both unobscured and obscured cases. Here, the attention mechanism used in DML balances the utilization of features in separating latent clusters. As the proposed DML-based model receives segmentation masks as labels, it can be generalized to other application domains. The proposed model can also be utilized for other domain-adaptation tasks on multimodal data by modifying the segmentation masks or proposing other application-specific annotation mechanisms. Despite DML's significant role in constructing better cluster representatives and compressing latent features around them, some samples may still be misclassified. As seen in Fig. 8, false positives are mostly cases with high thermal activity originating from hot smoke, neighboring flames, and other thermal sources. False negatives are mostly small flames (compared to the patch size), corresponding to early thermal patterns with different shapes than larger flames. In general, the leading causes of misclassification are assumed to include the lack of thermal reference points, similar thermal distributions in smoke and flame, thermal noise originating from temperature gradients of superimposed/interfered smoke layers in adjacent ignited areas, and, lastly, the premature thermal distribution of early flames resulting in outliers relative to larger flames. These limitations open up new research directions in many fields, such as multimodal detection, thermal imagery, and representation learning, as well as the necessity of a more challenging dataset suited for generalization to real-world scenarios.

REFERENCES

- [1] F. Tedim et al., "Defining extreme wildfire events: Difficulties, challenges, and impacts," *Fire*, vol. 1, no. 1, p. 9, Feb. 2018.
- [2] F. Afghah, A. Razi, J. Chakareski, and J. Ashdown, "Wildfire monitoring in remote areas using autonomous unmanned aerial vehicles," in *Proc. IEEE Conf. Comput. Commun. Workshops (INFOCOM WKSHPS)*, Apr. 2019, pp. 835–840.
- [3] S. Khoshdel, F. Afghah, and Q. Luo, "SkyGrid: Energy-flow optimization at harmonized aerial intersections," 2024, *arXiv:2406.13817*.
- [4] A. Shamsoshoara, F. Afghah, A. Razi, L. Zheng, P. Z. Fulé, and E. Blasch, "Aerial imagery pile burn detection using deep learning: The FLAME dataset," *Comput. Netw.*, vol. 193, Jul. 2021, Art. no. 108001.
- [5] A. Shamsoshoara, F. Lotfi, S. Mousavi, F. Afghah, and İ. Güvenc, "Joint path planning and power allocation of a cellular-connected UAV using apprenticeship learning via deep inverse reinforcement learning," *Comput. Netw.*, vol. 254, Dec. 2024, Art. no. 110789.
- [6] S. P. H. Boroujeni et al., "A comprehensive survey of research towards AI-enabled unmanned aerial systems in pre-, active-, and post-wildfire management," *Inf. Fusion*, vol. 108, Aug. 2024, Art. no. 102369.
- [7] H. Rajoli, P. Afshin, and F. Afghah, "Thermal image calibration and correction using unpaired cycle-consistent adversarial networks," in *Proc. 57th Asilomar Conf. Signals, Syst., Comput.*, Oct. 2023, pp. 1425–1429.
- [8] S. Khoshdel, Q. Luo, and F. Afghah, "PyroTrack: Belief-based deep reinforcement learning path planning for aerial wildfire monitoring in partially observable environments," 2024, *arXiv:2403.11095*.
- [9] A. Gaur, A. Singh, A. Kumar, A. Kumar, and K. Kapoor, "Video flame and smoke based fire detection algorithms: A literature review," *Fire Technol.*, vol. 56, no. 5, pp. 1943–1980, Apr. 2020.
- [10] M. J. Sousa, A. Moutinho, and M. Almeida, "Thermal infrared sensing for near real-time data-driven fire detection and monitoring systems," *Sensors*, vol. 20, no. 23, p. 6803, Nov. 2020.
- [11] X. Chen et al., "Wildland fire detection and monitoring using a drone-collected RGB/IR image dataset," *IEEE Access*, vol. 10, pp. 121301–121317, 2022.
- [12] S. Treneska and B. R. Stojkoska, "Wildfire detection from UAV collected images using transfer learning," in *Proc. 18th Int. Conf. Informat. Inf. Technol.*, Skopje, North Macedonia, 2021, pp. 6–7.
- [13] H. Rajoli, F. Lotfi, A. Atyabi, and F. Afghah, "Triplet loss-less center loss sampling strategies in facial expression recognition scenarios," in *Proc. 57th Annu. Conf. Inf. Sci. Syst. (CISS)*, Mar. 2023, pp. 1–6.
- [14] A. Farahani, S. Voghoei, K. Rasheed, and H. R. Arabnia, "A brief review of domain adaptation," in *Proc. ICDATA*, 2020, pp. 877–894.
- [15] J. Li, Y. Wu, J. Zhao, and K. Lu, "Low-rank discriminant embedding for multiview learning," *IEEE Trans. Cybern.*, vol. 47, no. 11, pp. 3516–3529, Nov. 2017.
- [16] J. Yu, X. Yang, F. Gao, and D. Tao, "Deep multimodal distance metric learning using click constraints for image ranking," *IEEE Trans. Cybern.*, vol. 47, no. 12, pp. 4014–4024, Dec. 2017.
- [17] Y. Chen et al., "UAV image-based forest fire detection approach using convolutional neural network," in *Proc. 14th IEEE Conf. Ind. Electron. Appl. (ICIEA)*, Jun. 2019, pp. 2118–2123.
- [18] S. C. Sethuraman, G. R. Tadkapally, S. P. Mohanty, and A. Subramanian, "IDrone: IoT-enabled unmanned aerial vehicles for detecting wildfires using convolutional neural networks," *Social Netw. Comput. Sci.*, vol. 3, no. 3, p. 242, May 2022.
- [19] A. Nguyen, H. Nguyen, V. Tran, H. X. Pham, and J. Pestana, "A visual real-time fire detection using single shot multibox detector for UAV-based fire surveillance," in *Proc. IEEE 8th Int. Conf. Commun. Electron. (ICCE)*, Jan. 2021, pp. 338–343.
- [20] J. Ryu and D. Kwak, "Flame detection using appearance-based pre-processing and convolutional neural network," *Appl. Sci.*, vol. 11, no. 11, p. 5138, May 2021.
- [21] D. Shen, X. Chen, M. Nguyen, and W. Q. Yan, "Flame detection using deep learning," in *Proc. 4th Int. Conf. Control Autom. Robot. (ICCAR)*, 2018, pp. 416–420.
- [22] W. Li and Z. Yu, "A lightweight convolutional neural network flame detection algorithm," in *Proc. IEEE 11th Int. Conf. Electron. Inf. Emergency Commun. (ICEIEC)*, Jun. 2021, pp. 83–86.
- [23] L. Karlinsky et al., "RepMet: Representative-based metric learning for classification and few-shot object detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 5197–5206.
- [24] B. Hong, Y. Zhou, H. Qin, Z. Wei, H. Liu, and Y. Yang, "Few-shot object detection using multimodal sensor systems of unmanned surface vehicles," *Sensors*, vol. 22, no. 4, p. 1511, Feb. 2022.
- [25] Y. Lu, X. Chen, Z. Wu, and J. Yu, "Decoupled metric network for single-stage few-shot object detection," *IEEE Trans. Cybern.*, vol. 53, no. 1, pp. 514–525, Jan. 2023.
- [26] P. Wu, S. C. H. Hoi, P. Zhao, C. Miao, and Z.-Y. Liu, "Online multi-modal distance metric learning with application to image retrieval," *IEEE Trans. Knowl. Data Eng.*, vol. 28, no. 2, pp. 454–467, Feb. 2016.

- [27] K. Roth, T. Milbich, S. Sinha, P. Gupta, B. Ommer, and J. P. Cohen, "Revisiting training strategies and generalization performance in deep metric learning," in *Proc. Int. Conf. Mach. Learn. (ICML)*, Jul. 2020, pp. 8242–8252.
- [28] F. Lotfi, F. Afghah, and J. Ashdown, "Attention-based open RAN slice management using deep reinforcement learning," in *Proc. IEEE Global Commun. Conf.*, Dec. 2023, pp. 6328–6333.
- [29] W. Kim, B. Goyal, K. Chawla, J. Lee, and K. Kwon, "Attention-based ensemble for deep metric learning," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 2018, pp. 736–751.
- [30] X. Wang, Y. Hua, E. Kodirov, G. Hu, and N. M. Robertson, "Deep metric learning by online soft mining and class-aware attention," in *Proc. AAAI*, vol. 33, 2019, pp. 5361–5368.
- [31] D. Kotovenko, P. Ma, T. Milbich, and B. Ommer, "Cross-image-attention for conditional embeddings in deep metric learning," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2023, pp. 11070–11081.
- [32] J. D. Seidenschwarz, I. Elezi, and L. Leal-Taixé, "Learning intra-batch connections for deep metric learning," in *Proc. Int. Conf. Mach. Learn.*, 2021, pp. 9410–9421.
- [33] Y. Li, S. Kan, and Z. He, "Unsupervised deep metric learning with transformed attention consistency and contrastive clustering loss," in *Proc. Eur. Conf. Comput. Vis.* Springer, 2020, pp. 141–157.
- [34] P. Yang et al., "A deep metric learning approach for histopathological image retrieval," *Methods*, vol. 179, pp. 14–25, Jul. 2020.
- [35] H. Dong, K. Song, Q. Wang, Y. Yan, and P. Jiang, "Deep metric learning-based for multi-target few-shot pavement distress classification," *IEEE Trans. Ind. Informat.*, vol. 18, no. 3, pp. 1801–1810, Mar. 2022.
- [36] H. Coskun, D. J. Tan, S. Conjeti, N. Navab, and F. Tombari, "Human motion analysis with deep metric learning," in *Proc. ECCV*, 2018, pp. 667–683.
- [37] L. Fang, Y. Zhang, H. Li, B. Chen, and Y. Sui, "Dynamic attention network for flame detection," *Electron. Lett.*, vol. 59, no. 2, Jan. 2023, Art. no. e12690.
- [38] X. Wang, Z. Pan, H. Gao, N. He, and T. Gao, "An efficient model for real-time wildfire detection in complex scenarios based on multi-head attention mechanism," *J. Real-Time Image Process.*, vol. 20, no. 4, p. 66, Aug. 2023.
- [39] E. Buza and A. Akagic, "Unsupervised method for wildfire flame segmentation and detection," *IEEE Access*, vol. 10, pp. 55213–55225, 2022.
- [40] D. Dzigal, A. Akagic, E. Buza, A. Brdjanin, and N. Dardagan, "Forest fire detection based on color spaces combination," in *Proc. 11th Int. Conf. Electr. Electron. Eng. (ELECO)*, Nov. 2019, pp. 595–599.
- [41] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III 18*, vol. 9351. Springer, 2015, pp. 234–241.
- [42] A. H. Farzaneh and X. Qi, "Facial expression recognition in the wild via deep attentive center loss," in *Proc. IEEE Winter Conf. Appl. Comput. Vis. (WACV)*, Jan. 2021, pp. 2402–2411.
- [43] P. Fule et al., "Flame 2: Fire detection and modeling: Aerial multi-spectral image dataset," IEEE DataPort, 2022. [Online]. Available: <https://ieee-dataport.org/open-access/flame-2-fire-detection-and-modeling-aerial-multi-spectral-image-dataset>
- [44] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual explanations from deep networks via gradient-based localization," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 618–626.
- [45] E. Zachariadis, M. Antonakakis, and M. Zervakis, "Wildforest fire detection with ShRe-xception network on aerial optical and infrared images," in *Proc. IEEE Int. Conf. Imag. Syst. Techn. (IST)*, Oct. 2023, pp. 1–6.
- [46] R. Ghali, M. A. Akhloufi, and W. S. Mseddi, "Deep learning and transformer approaches for UAV-based wildfire detection and segmentation," *Sensors*, vol. 22, no. 5, p. 1977, Mar. 2022.
- [47] R. Ghali and M. A. Akhloufi, "CT-fire: A CNN-transformer for wildfire classification on ground and aerial images," *Int. J. Remote Sens.*, vol. 44, no. 23, pp. 7390–7415, Dec. 2023.