

Explaining the Unseen: Multimodal Vision-Language Reasoning for Situational Awareness in Underground Mining Disasters

Mizanur Rahman Jewel¹, Mohamed Elmahallawy², Sanjay Madria¹, Samuel Frimpong¹

¹Missouri University of Science and Technology, ²Washington State University

{mj9vc, madrias, frimpong}@mst.edu, mohamed.elmahallawy@wsu.edu

Abstract

Underground mining disasters produce pervasive darkness, dust, and collapses that obscure vision and make situational awareness difficult for humans and conventional systems. To address this, we propose MDSE, Multimodal Disaster Situation Explainer, a novel vision-language framework that automatically generates detailed textual explanations of post-disaster underground scenes. MDSE has three-fold innovations: (i) Context-Aware Cross-Attention for robust alignment of visual and textual features even under severe degradation; (ii) Segmentation-aware dual pathway visual encoding that fuses global and region-specific embeddings; and (iii) Resource-Efficient Transformer-Based Language Model for expressive caption generation with minimal compute cost. To support this task, we present the Underground Mine Disaster (UMD) dataset—the first image-caption corpus of real underground disaster scenes—enabling rigorous training and evaluation. Extensive experiments on UMD and related benchmarks show that MDSE substantially outperforms state-of-the-art captioning models, producing more accurate and contextually relevant descriptions that capture crucial details in obscured environments, improving situational awareness for underground emergency response. The code is at [Github](#).

1. Introduction

Underground mining disasters present uniquely harsh and complex conditions that severely hinder situational awareness [10] and locating trapped personnel and coordinating effective rescue operations. However, conventional monitoring and communication systems often become inoperable due to structural collapse, power failure, or signal disruption, leaving responders with partial or outdated information (see Fig. 1). Existing situational awareness tools, typically based on discrete sensor inputs such as gas detec-



Figure 1. Representative scenes depicting disaster scenarios in underground mining environments.

tors or seismic monitors, offer limited insight and lack the contextual understanding necessary for a holistic scene interpretation [29]. Responders must often navigate in “near-zero” visibility caused by thick smoke, airborne dust, and total darkness, limiting their ability to detect hazards, locate victims, or understand the structural layout [16].

While image captioning has seen significant progress through models trained on large-scale general-purpose datasets such as the COCO caption dataset [5] and Flickr30k [17], as well as more recent work focused on disaster imagery [25], these approaches fall short when applied to the extreme and domain-specific conditions of underground mine disasters [8, 10]. Subterranean environments present unique vision challenges: narrow, confined spaces; structurally unstable scenes; low-light to dark conditions; high particulate density; and secondary hazards such as gas leaks and flooding. These conditions not only degrade image quality but also obscure critical visual cues necessary for accurate scene understanding and caption generation. Despite their potential, no existing vision-language framework has been designed to operate under such multimodal degradation and risk. This gap underscores the urgent need for a robust framework capable of integrating noisy or degraded visual data with natural language reasoning to generate context-aware, real-time descriptions that support situ-

This research was supported by a grant from CDC–NIOSH.

ational awareness and decision-making in underground disaster response.

To this end, we propose the **Multimodal Disaster Situation Explainer (MDSE)**—the first vision-language framework to enhance post-disaster situational awareness in underground mines by integrating computer vision and natural language processing to automatically generate descriptive, actionable insights from real-time visual data captured in subterranean scenes. Trained on mining-domain data, MDSE detects critical hazards, structural damage, and potential escape routes, even under severe visual degradation. It incorporates mining-specific terminology and aligns with safety protocols to produce accurate, contextually relevant narratives interpretable by emergency responders. MDSE further introduces a novel contrast-aware cross-attention mechanism that improves feature fusion in low-light, smoke-filled, or occluded environments, enabling the system to extract subtle cues, such as cracked beams or blocked passages, and translate them into concise textual reports. In summary, the key contributions are as follows:

1. **Novel Domain-Specific Captioning Framework:** We propose the first vision-language framework for mining disaster, introducing a *context-aware cross-attention mechanism* that enhances salient visual features—such as edges and contours—in low-light, high-noise imagery, thereby improving captioning accuracy.
2. **Vision Feature Enhancement and Segmentation-Aware Dual-Pathway Visual Encoding:** We propose a strategy that extracts both global scene-level features and object-centric regional features using a combination of a vision encoder and the Segment Anything Model (SAM) [11]. SAM produces precise segmentation masks even in dark, degraded images, isolating critical regions of interest. This allows the model to concentrate on key objects despite noise or occlusion in underground scenes. The dual-pathway design captures broad visual context while emphasizing salient regions, enhancing robustness in visually ambiguous conditions.
3. **Integration of Low-Resource LLMs:** Instead of relying on an extremely large language model for caption generation, MDSE utilizes Parameter-Efficient Fine-Tuning with Low-Rank Adaptation (LoRA) on the Q-Former layers and LLM head, significantly reducing the number of trainable parameters during adaptation to downstream tasks, retaining the expressiveness of the original model while ensuring computational and memory efficiency.
4. **Underground Mine Disaster (UMD) Dataset :** We construct the first image-captioning dataset specifically for underground mining emergencies, addressing a critical gap in vision-language research for extreme, low-visibility environments. To evaluate the effectiveness of our proposed framework, MDSE, we conduct

comprehensive experiments on the UMD dataset, as well as on widely used benchmarks (COCO caption dataset [5], Flickr30k [17]) and disaster-focused datasets (Incidents1M [25], DNICC19k [33]). Comparative analysis against state-of-the-art captioning models demonstrates that MDSE delivers robust and contextually accurate performance across both general and domain-specific scenarios, particularly excelling in the challenging conditions of underground disaster scenes.

2. Related Work

Vision-Language Models and Image Captioning. The task of generating descriptive sentences from visual input has advanced significantly with the advent of deep learning, especially transformer-based architectures and vision-language pretraining [1, 6, 28]. CLIP [18] pioneered the alignment of textual and visual embeddings in a shared latent space, enabling strong *zero-shot* performance across multimodal tasks. BLIP [12] and its successor BLIP-2 [13] further improved caption quality by incorporating multimodal encoders and decoders. Similarly, Flamingo [3] offered a frozen vision-language transformer that set new benchmarks in few-shot and open-ended captioning tasks.

Specialized and Instruction-Tuned Models. Detective networks [19] introduced attention-shifting mechanisms to enhance disaster image recognition, showing promise for improving contextual understanding in visually complex environments. LLaVA [14] demonstrated the utility of instruction tuning for aligning vision-language models with user intent, boosting generalization in unseen domains. CrisisViT [15] leveraged Vision Transformers with adversarial training on Incidents1M [25] and CrisisMMD [2], achieving state-of-the-art performance across several disaster-related visual classification tasks.

Captioning Under Degraded Visual Conditions. Captioning in low-light environments poses unique challenges due to diminished visual cues and increased noise. While research on captioning in such conditions is limited, related work in image enhancement, denoising, and low-light reconstruction has shown effectiveness in restoring perceptual quality. Techniques like (SAM) [11] offer robust segmentation capabilities, which can be integrated into captioning pipelines to help isolate critical scene components, even in dark or obscured settings.

Limitations in Underground Disaster Captioning. Existing image captioning models have not addressed the unique challenges of domain-specific tasks such as disasters in underground mining, where extreme visual degradation and environmental hazards render standard approaches ineffective. This work fills that gap by proposing a tailored vision-language framework to enhance situational awareness and support decision-making in subterranean emergencies.




Figure 2. Overview of the proposed MDSE framework and its main components.

3. Methodology

This section presents our proposed MDSE (Multimodal Disaster Situation Explainer) framework shown in Fig. 2.

3.1. Feature Enhancement and Segmentation-Aware Dual-Pathway Encoding

Given an input image I , a feature enhancement module $\mathcal{F}_E$ first compensates for low light and noise by producing an enhanced image as $\hat{I} = \mathcal{F}_E(I)$. The feature enhancement module is designed as a sequential pipeline that improves image quality through several established operations. First, multi-exposure fusion combines differently exposed inputs using illumination-aware weighting to produce a well-balanced base image [24]. Next, spatial tone mapping and local contrast enhancement boost detail visibility by center-surround filtering and multiscale contrast adjustment [21, 22]. White balance is then corrected using a combination of Gray-World and White-Point assumptions with a power transformation to remove color cast [23]. Finally, chromatic saturation is adaptively scaled to restore vividness and perceptual naturalness, yielding the enhanced output image. This ensures the vision encoder receives improved feature representations.

A pretrained SAM $\mathcal{S}$ then generates segmentation masks M highlighting key regions from the enhanced image $\hat{I}$ as

$M = \mathcal{S}(\hat{I})$. The segmentation mask guides the model to focus on relevant objects or regions, reducing background noise during feature extraction. Even under poor visibility (darkness or smoke), SAM can delineate key objects and structures (e.g., equipment, hazards, tunnel walls) by generating multiple region-of-interest (ROI) masks in a single image. This allows MDSE to isolate numerous relevant sub-scenes within the disaster image for focused analysis. The enhanced image $\hat{I}$ and segmentation mask M are independently encoded by a CLIP ViT-based vision transformer $\mathcal{V}_E$ [18] as:

$$E_G = \mathcal{V}_E(\hat{I}), \quad E_S = \mathcal{V}_E(M). \quad (1)$$

This dual-pathway extraction produces robust visual embeddings at two scales: E_G captures the global scene embedding, while E_S encodes fine-grained region-specific features within segmented sub-regions. By integrating these complementary representations, MDSE maintains broad situational awareness and preserves critical localized evidence in challenging, low-visibility environments.

Next, we fuse the global and segmented embeddings using a learnable projection matrix W_f as:

$$E_F = W_f [E_G; E_S], \quad (2)$$

where $[E_G; E_S]$ denotes the concatenation of the global and segment-level embeddings. This fusion operation integrates

complementary information from both pathways, yielding a rich and comprehensive feature representation.

MDSE processes both global and region-specific views to ensure that subtle peripheral details are retained. SAM-guided masks combined with CLIP encodings yield a multi-scale representation: the global embedding captures overall context, while the region-level embeddings preserve localized evidence—like a cracked beam or a miner’s helmet lamp—even under extreme low-visibility conditions.

3.2. Context-Aware Cross-Attention Fusion

Once robust global and region-specific embeddings are obtained, MDSE fuses them with linguistic context via a contrast-aware cross-attention module built on a BLIP-2 Q-Former [13]. The Q-Former comprises two transformer submodules: i) *An image transformer* that refines and projects visual embeddings; and ii) *A text transformer* that encodes prompt tokens and autoregressively decodes captions. This design allows the model to dynamically focus on complementary visual features and generate coherent, context-aware descriptions.

External context C is transformed by a learnable projection f_{ctx} as $C' = f_{\text{ctx}}(C)$ that integrates auxiliary context—such as scene type or task-specific cues—directly into the attention keys, enabling the model to prioritize visual features most relevant to the downstream task. Given the fused visual embedding E_F and context vector C' , the context-enhanced key embeddings K' is computed as:

$$K' = E_F + \sigma(W_g [E_F; C']), \quad (3)$$

where E_F denotes final feature embeddings, σ is the sigmoid activation, and W_g is a learnable gating projection. This gated fusion controls the degree of context influence on the visual features, allowing the model to modulate between relying on raw visual features and contextually enriched representations based on the task requirements.

A fixed set of N learnable query embeddings is provided as input to the image transformer (the number of queries and other Q-Former hyperparameters are specified in Sec. 4). These query embeddings interact with each other through self-attention layers and engage with image features via context-aware cross-attention layers, which are inserted at regular intervals throughout the transformer blocks. Next, given queries Q , the attention computation is computed as:

$$\text{Attention}(Q, K', V') = \text{softmax} \left(\frac{QK'^T}{\sqrt{d_k}} \right) V' \quad (4)$$

This operation enables dynamic interaction between query tokens and context-enriched visual embeddings, refining the attention focus based on the provided context. Furthermore, the queries can interact with text tokens through shared self-attention layers. Unlike the standard BLIP-2 implementation, where Q-Former uses fixed cross-attention with vision

features, MDSE extends this by integrating a learned context signal directly into the key and value representations. This allows the cross-attention layers to adapt to both the visual input and the contextual cues, making the model more flexible and better suited for complex vision-language tasks, especially in diverse or challenging environments.

3.3. Loss Function

Following the Q-Former training paradigm, we jointly optimize a weighted sum of three loss terms:

$$\mathcal{L}_{\text{total}} = \lambda_1 \cdot \mathcal{L}_{\text{ITC}} + \lambda_2 \cdot \mathcal{L}_{\text{ITM}} + \lambda_3 \cdot \mathcal{L}_{\text{ITG}} \quad (5)$$

where λ_i is a weighting coefficient that balances the objectives, preventing any single task from dominating the learning process. The values of λ_i are tuned empirically to optimize performance across retrieval, matching, and captioning tasks. The individual loss terms are defined as follows:

$$\mathcal{L}_{\text{ITC}} = -\frac{1}{N} \sum_{i=1}^N \sum_{j=1}^N y_{ij} \log \frac{\exp(\text{sim}(f_I^i, f_T^j)/\tau)}{\sum_{k=1}^N \exp(\text{sim}(f_I^i, f_T^k)/\tau)}, \quad (6)$$

$$\mathcal{L}_{\text{ITM}} = -\frac{1}{N} \sum_{i=1}^N \sum_{c=1}^C y_{i,c} \log \left(\frac{\exp(\text{ITM.Logits}_{i,c})}{\sum_{k=1}^C \exp(\text{ITM.Logits}_{i,k})} \right), \quad (7)$$

$$\mathcal{L}_{\text{ITG}} = -\sum_{t=1}^T \log P(w_t | w_{<t}, E_F). \quad (8)$$

Here $\mathcal{L}_{\text{ITC}}$ denotes the *Image–Text Contrastive* loss, which aligns corresponding image and text features f_I and f_T . The function $\text{sim}(\cdot, \cdot)$ computes the cosine similarity between visual embedding v_i and text embedding t_i , scaled by a temperature τ . This objective encourages matching image–text pairs to have high similarity while pushing apart mismatched pairs. $\mathcal{L}_{\text{ITM}}$ denotes the *Image–Text Matching* loss. The ITM.Logits are obtained from a linear classification head applied to the averaged query token embeddings produced by the image transformer after attending to the fused visual and textual features. This loss fine-tunes the model’s ability to distinguish between matched and mismatched image–text pairs, thereby directly enhancing performance on retrieval and matching tasks. $\mathcal{L}_{\text{ITG}}$ denotes the *Image-grounded Text Generation* loss, which corresponds to the autoregressive language modeling objective (i.e., captioning loss) over a sequence of length T . At each time step t , the model predicts the next word w_t conditioned on the previously generated tokens $w_{<t}$ and the fused visual features E_F . The probability $P(w_t | w_{<t}, E_F)$ is computed by the autoregressive decoder. This loss guides the model to generate fluent, semantically accurate, and contextually grounded captions, word by word. The above combined objectives ensure cross-modal alignment, robust language understanding, and accurate caption generation.

3.4. Low-Rank Adaptation for LLM

For a model weight W , LoRA [9] introduces low-rank matrices A and B :

$$W' = W + A \cdot B \quad (9)$$

with $A \in \mathbb{R}^{d \times r}$, $B \in \mathbb{R}^{r \times d}$, where $r \ll d$. LoRA allows efficient fine-tuning by updating only the low-rank factors, reducing computational overhead and memory usage. The parameter d denotes the original weight dimension, r is the rank deciding the compression factor, and W' is the adapted weight matrix after low-rank modification. This method retains the expressiveness of the full model while enabling fast adaptation to new tasks with minimal parameters.

In our MDSE framework, LoRA modules are injected into the linear projections of the Q-Former's attention layers (i.e., the query, key, and value transformations). Specifically, for each projection matrix W_Q, W_K, W_V in the self-attention and cross-attention layers, LoRA augments it as:

$$W'_Q = W_Q + A_Q \cdot B_Q, \quad (10)$$

$$W'_K = W_K + A_K \cdot B_K, \quad (11)$$

$$W'_V = W_V + A_V \cdot B_V \quad (12)$$

This design allows the Q-Former to adapt its attention dynamics to domain-specific scenarios (e.g., low-light disaster imagery) without altering the pretrained parameters. Since the Q-Former is the primary trainable component in BLIP-2, applying LoRA here yields efficient task-specific adaptation with minimal computational and memory overhead. On the LLM side, a LoRA adapter is placed on the projection that maps Q-Former outputs into the language model. We use LoRA rank 8, scale 16 for all Q-Former attention and MLP projections, and rank 4, scale 8 for the LLM input projector, with LoRA-dropout 0.05. In inference, the model first enhances and segments the input image, then fuses the resulting embeddings and applies context-aware cross-attention. The Q-Former and the LoRA-adapted LLM generate text autoregressively, enabling efficient multimodal reasoning in complex visual scenarios. This streamlined process ensures robust performance while maintaining computational efficiency.

4. Performance Evaluation

4.1. Experimental Setup

We do not train SAM within MDSE; instead, SAM is used as an offline tool to generate region proposals that localize salient structures. The only visual backbone is CLIP ViT-L/14, which processes 224×224 inputs with a 14×14 patch size and outputs a sequence of 257 visual tokens of dimension 1024 (including the class token). The trainable part is the Q-Former (plus a small projector into the language model) with LoRA adapters. For query representation, we




Figure 3. Representative samples from the UMD dataset with corresponding textual annotations.

use $N = 32$ learnable queries of dimension $d = 768$, yielding $\mathbf{Z} \in \mathbb{R}^{32 \times 768}$. The visual tokens are linearly projected from 1024 to 768 before cross-attention, and the Q-Former outputs are then projected to the language-model hidden size $d_L = 2560$ (OPT-2.7B) [32]. With LoRA rank $r = 8$ and scale $\alpha = 16$ on all Q-Former attention/MLP projections and on the LLM head, the total number of trainable parameters is $\approx 2.0\text{M}$. The loss weights ($\lambda_{\text{ITC}}, \lambda_{\text{ITM}}, \lambda_{\text{ITG}}$) were tuned via a small grid on the validation split where we adopt (1.0, 0.5, 0.10) based on the best performance.

Datasets. To evaluate MDSE against state-of-the-art approaches, we use several benchmark datasets for image captioning: Flickr30k [17], COCO Captions [5], Incidents1M [25], CDNIC19k [33], and our UMD dataset. Brief descriptions of these datasets are provided below:

- **Underground Mine Disaster (UMD) Dataset:** We introduce the (UMD) dataset, a domain-specific image-caption corpus designed to advance visual recognition and situational understanding in underground mining disaster scenarios, an area largely absent from existing computer vision benchmarks. Such events present unique challenges, including confined spaces, poor lighting, smoke, debris, and structural collapse, which require specialized models for reliable detection and assessment. To address the scarcity of publicly available imagery, we curated the dataset through keyword-based internet searches and supplemented it with visually realistic scenes from disaster films that simulate underground emergencies. These cinematic samples provide diverse, contextually rich depictions of hazards, rescue efforts, and structural damage. UMD contains 300 manually annotated images with five captions for each image, partially drawn from Image-Mine [10], DNICC19k [33], and

Incidents1M [25]. Representative examples are shown in Fig. 3. To standardize heterogeneous sources, all images are de-duplicated, lightly border-trimmed, color-normalized, and resized to 224×224 pixels for both training and inference.

- **Flickr30k** [17]: It is a widely used image captioning dataset with roughly 31,783 images from Flickr, each annotated with five human-written English captions. Covering diverse everyday scenes, it serves as a standard benchmark for evaluating captioning models.
- **COCO caption dataset** [5]: Microsoft COCO (Common Objects in Context) captions is a large-scale dataset comprising over 200,000 images annotated for tasks such as object detection, segmentation, and image captioning. Each image is paired with five captions, providing strong language grounding for vision-language modeling.
- **Incidents1M** [25]: A large-scale dataset with one million images of real-world disasters, accidents, and critical incidents, including floods, earthquakes, fires, and infrastructure failures. It provides annotations for incident type, damage extent, and context, enabling tasks such as disaster classification and automated damage assessment. Its various event types, regions, and conditions make it a challenging benchmark for scene understanding and advanced research in crisis response and emergency control.
- **DNICC19k** [33]: It is known as CDNIC19k, a disaster-focused dataset with $\sim 19,000$ images across 15 categories, including natural hazards, accidents, and emergencies. It provides annotations for damage assessment, incident classification, and captioning. While less widely used than COCO or Flickr30k, DNICC19k offers event-specific imagery with contextual labels, supporting research in disaster control and humanitarian applications.

Baselines. For benchmarking, we consider a set of strong baseline vision-language models that allow a comprehensive assessment of MDSE’s effectiveness:

- **BLIP-2** [13]: A vision-language model using a frozen image encoder and a learnable Q-Former to bridge vision and language. Excels in multimodal tasks like captioning and visual question answering.
- **Flamingo** [3]: DeepMind’s few-shot multimodal model using gated cross-attention with a frozen image encoder and a language model for flexible multimodal tasks.
- **CLIP** [18]—A contrastive vision-language model aligning images and texts for robust zero-shot recognition and retrieval, widely used for visual understanding.
- **LLaVA** [14]: A large vision-language assistant augmenting a Vicuna LLM with visual grounding through lightweight projection and instruction-tuning on multimodal dialogues.
- **Florence-2** [26]: Microsoft’s unified vision-language transformer model trained for tasks like captioning, detection, and grounding on large-scale datasets.

- **Florence** [31]: Florence is a large-scale vision foundation model designed to learn generalizable visual representations from diverse datasets, serving as a strong baseline for tasks such as image classification, object detection, and image captioning.
- **STCNet** [33]: A disaster-domain captioner that leverages spatial context and topic guidance; it surpasses Show-Attend-and-Tell [27] and AoANet [7] on DNICC19k, making it a strong baseline for MDSE.
- **Qwen-2.5** [30]: Excels at integrating visual and textual features, producing accurate, detailed image descriptions through cross-modal reasoning.

Evaluation metrics. For performance comparison on the UMD and DNICC19k datasets, we use Consensus-based Image Description Evaluation (CIDER) [20] and Semantic Propositional Image Caption Evaluation (SPICE) [4] scores. Below is a brief description of each one:

- **CIDER** [20]: It measures similarity between a generated caption and human references using TF-IDF weighted n -gram matching, rewarding both distinctive content (via higher weights for rare n -grams) and commonly agreed phrases, yielding a robust measure of image relevance.
- **SPICE** [4]: It evaluates by comparing their semantic scene graphs—objects, attributes, and relationships—between generated and reference captions. It computes an F-score over matching tuples, focusing on meaning rather than surface n -gram overlap and aligning more closely with human judgment.

For the COCO caption and Flickr30k datasets, we evaluate models using recall-based metrics. In our experimental settings, recall at rank k ($R@k$) measures the proportion of queries for which the correct caption appears within the top k predictions. Formally, let N be the total number of queries, and let r_i be the rank at which the correct label appears for the i -th query. Then $R@k$ is defined as:

$$R@k = \frac{1}{N} \sum_{i=1}^N \mathbb{I}(r_i \leq k)$$

where $\mathbb{I}$ is the indicator function, which returns 1 when the condition $r_i \leq k$ is true and 0 otherwise. In practice, $R@1$ measures the fraction of instances where the correct answer is the top prediction, while $R@5$ and $R@10$ evaluate how often the correct answer is found within the top 5 or top 10 predictions, respectively.

Finally, for the Incidents1M dataset, we use the prompt “Describe the incidents depicted in this image among these labels”, as the dataset provides class-level labels rather than sentence-level annotations. We use mean Average Precision (mAP) as the primary evaluation metric, which is widely employed in multi-label classification and retrieval tasks. For each class, the Average Precision (AP) is computed by plotting the precision-recall curve and calculating the area under the curve (AUC). The mAP is then obtained by av-




Figure 4. MDSE inference results on UMD and DNICC19k datasets, showing its captioning performance under challenging conditions.

Table 1. Performance comparison of MDSE with baseline image captioning methods.

Model	UMD		DNICC19k	
	CIDEr	SPICE	CIDEr	SPICE
BLIP-2	0.62	0.50	0.58	0.45
Flamingo	0.53	0.42	0.53	0.41
Florence	0.58	0.47	0.57	0.43
LLaVA	0.59	0.47	0.57	0.43
STCNet	–	–	0.53	0.43
Qwen-2.5	0.66	0.51	0.62	0.47
MDSE (Ours)	0.70	0.53	0.66	0.49

eraging the AP values across all classes. This metric effectively captures both precision and recall, making it well-suited for handling class imbalance and multi-label scenarios, as present in the Incidents1M dataset.

4.2. Experimental Results

We evaluate the proposed MDSE framework against baselines using CIDEr and SPICE on the UMD and DNICC19k datasets. To test generalization, we compare performance on standard benchmarks (MSCOCO and Flickr30k) using image-text retrieval metrics. We also assess classification on the disaster-oriented Incidents1M dataset via zero-shot evaluation with a fixed prompt. All models are fine-tuned on UMD, DNICC19k, COCO, and Flickr30k.

Table 1 presents a comparison of our proposed MDSE method with baseline image captioning models on two disaster-related datasets: UMD and DNICC19k [33]. On the UMD dataset, MDSE significantly outperforms all other models. It achieves a CIDEr score of 0.70, which is notably higher than the next best model, BLIP-2, which

scored 0.62. MDSE also surpasses other baselines such as LLaVA, Qwen, and Florence. For the SPICE metric, MDSE scores 0.53, exceeding BLIP-2’s score of 0.50 and showing clear improvements over the other methods. These results demonstrate MDSE’s strong capability to produce captions that are both syntactically coherent and semantically meaningful. On the DNICC19k dataset, MDSE again shows superior performance, achieving a CIDEr score of 0.66 and a SPICE score of 0.49, both higher than those of the competing models. Additionally, Figure 4 highlights MDSE’s consistent improvements across datasets and metrics, demonstrating its effectiveness in producing human-aligned image descriptions on real-world datasets.

In Table 2, we show image-text retrieval results across COCO caption and Flickr30k datasets using Recall at k ($R@1$, $R@5$, $R@10$). These metrics reflect the percentage of queries for which a relevant item is retrieved within the top- k results, where higher values indicate better retrieval accuracy. On the COCO caption dataset, MDSE outperforms all baselines. It achieves the highest $R@1$ score of 0.88, surpassing BLIP-2, which achieves 0.84. For $R@5$, MDSE scores 0.97, slightly above BLIP-2’s 0.96. In $R@10$, both MDSE and BLIP-2 attain a perfect 0.98. Competing models like Flamingo and Florence lag behind, with $R@1$ scores of 0.66 and 0.68, respectively. Similarly, on Flickr30k, MDSE also demonstrates strong performance, matching BLIP-2’s $R@1$ score of 0.98. For both $R@5$ and $R@10$, MDSE achieves perfect scores of 1.00, equaling BLIP-2 and Flamingo. These results highlight MDSE’s ability to retrieve semantically relevant captions and its generalization across both domain-specific and general-purpose benchmarks.

Table 2. Comparison with baseline image-text retrieval methods.

Model	COCO Caption Dataset			Flickr30k Dataset		
	R@1	R@5	R@10	R@1	R@5	R@10
BLIP-2	0.84	0.96	0.98	0.98	1.00	1.00
CLIP	0.60	0.83	0.90	0.88	0.99	.99
Flamingo	0.66	0.87	0.93	0.89	0.99	1.00
Florence	0.68	0.86	–	0.91	.99	–
Qwen-2.5	0.87	0.97	0.98	0.99	1.00	1.00
MDSE (Ours)	0.88	0.97	0.98	0.98	1.00	1.00

Table 3. Zero-shot classification results on Incidents1M dataset.

Model	BLIP-2	CLIP	Flamingo	Florence-2	MDSE (Ours)
mAP	0.44	0.50	0.45	0.52	0.47

Finally, in Table 3, we report the mAP scores of our MDSE approach compared to the baselines on the Incidents1M dataset for zero-shot classification. Among the baselines, Florence-2 achieves the highest mAP score of 0.52, followed by CLIP with 0.50. Our MDSE model achieves a competitive mAP of 0.47, outperforming BLIP-2 (0.44) and Flamingo (0.45), though it does not surpass the top-performing models. Additionally, Figure 5 presents qualitative examples of MDSE’s classification outputs. The model correctly classifies three images corresponding to earthquake, landslide, and sinkhole, while the remaining examples are misclassified.

It is worth emphasizing that the primary objectives of MDSE are image captioning and image-text retrieval, where it has shown state-of-the-art performance, as detailed in earlier sections. In contrast, zero-shot classification is a core strength of models like CLIP and Florence-2, which are explicitly designed for broad visual recognition and multi-task vision-language understanding. Despite this distinction, MDSE demonstrates robust generalization to unseen incident categories and provides strong zero-shot classification performance on Incidents1M (see Figure 5). These results highlight MDSE’s versatility and balanced performance across diverse vision-language tasks.

4.3. Ablation Study

We ablate MDSE by removing one component at a time and report CIDEr/SPICE on UMD and DNICC19k (Table 4). The largest degradations arise from removing dual-path fusion (global+region) and removing SAM-based region proposals (offline), confirming the value of combining global/region evidence and localized proposals. If we disable the Context-aware cross-attention and use the cross-attention module from the original Q-former, we find lower SPICE scores for both datasets, indicating weaker relational detail. Removing the feature enhancement module yields only a small decline, suggesting FE improves robustness but is not the dominant driver of gains.



Figure 5. Demo results of MDSE on the Incidents1M dataset.

Table 4. MDSE’s evaluation on UMD and DNICC19k datasets.

Model	UMD		DNICC19k	
	CIDEr	SPICE	CIDEr	SPICE
w/o Dual-Pathway Encoding	0.65	0.5	0.62	0.47
w/o SAM	0.66	0.51	0.63	0.47
w/o Context aware cross attention	0.67	0.5	0.64	0.46
w/o FE	0.68	0.53	0.64	0.48
MDSE (whole)	0.70	0.53	0.66	0.49

Table 5. LoRA hyperparameter sweep on UMD and DNICC19k (r = rank, α = scale).

(r, α)	UMD		DNICC19k	
	CIDEr	SPICE	CIDEr	SPICE
(4, 4)	0.68	0.52	0.64	0.48
(4, 8)	0.69	0.52	0.65	0.48
(8, 8)	0.69	0.53	0.65	0.49
(8, 16)	0.70	0.53	0.66	0.49
(16, 16)	0.69	0.52	0.65	0.48
(16, 32)	0.68	0.51	0.64	0.47

We also sweep LoRA hyperparameters (r, α) on the Q-Former and LLM input projector Table 5. Performance is stable across settings, with a mild peak in the middle on both datasets. These results indicate that modest-rank adapters provide efficient task adaptation without large compute/memory overhead, while more aggressive ranks do not yield additional gains under our data scale.

5. Conclusion

This paper presents MDSE, a new captioning framework tailored for low-light and visually challenging underground mine disaster scenes. MDSE combines segmentation-aware dual-pathway visual encoding with context-aware cross-attention, enabling robust alignment of visual and textual data under severe degradation. Experiments on our newly introduced UMD dataset demonstrate that MDSE outperforms state-of-the-art baselines, generating more accurate and domain-appropriate descriptions, with a 12% improvement over leading vision-language models.

References

- [1] Praveesh Agrawal, Szymon Antoniak, Emma Bou Hanna, Baptiste Bout, Devendra Chaplot, Jessica Chudnovsky, Diogo Costa, Baudouin De Monicault, Saurabh Garg, Theophile Gervet, et al. Pixtral 12b. *arXiv preprint arXiv:2410.07073*, 2024. 2
- [2] Firoj Alam, Ferda Ofli, and Muhammad Imran. Crisismmd: Multimodal twitter datasets from natural disasters. In *Proceedings of the 12th International AAAI Conference on Web and Social Media (ICWSM)*, 2018. 2
- [3] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *Advances in neural information processing systems*, 35:23716–23736, 2022. 2, 6
- [4] Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. Spice: Semantic propositional image caption evaluation. In *European conference on computer vision*, pages 382–398. Springer, 2016. 6
- [5] Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*, 2015. 1, 2, 5, 6
- [6] Wenliang Dai, Junnan Li, Dongxu Li, Anthony Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale N Fung, and Steven Hoi. Instructblip: Towards general-purpose vision-language models with instruction tuning. *Advances in neural information processing systems*, 36:49250–49267, 2023. 2
- [7] Guangyu Gao, Qianxiang Wang, Jing Ge, and Yan Zhang. Aonet: attentional occlusion-aware network for occluded person re-identification. In *Proceedings of the Asian conference on computer vision*, pages 1606–1621, 2022. 6
- [8] Abhay Goyal, Sanjay Madria, and Samuel Frimpong. Oglemine: Obstacle-infused goal-conditioned learning for post-disaster navigation in underground mine. In *Proceedings of the 37th International Conference on Scalable Scientific Data Management*, pages 1–12, 2025. 1
- [9] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3, 2022. 5
- [10] Mizanur Rahman Jewel, Mohamed Elmahallawy, Sanjay Madria, and Samuel Frimpong. Dis-mine: instance segmentation for disaster-awareness in poor-light condition in underground mines. In *2024 IEEE International Conference on Big Data (BigData)*, pages 6279–6288. IEEE, 2024. 1, 5
- [11] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4015–4026, 2023. 2
- [12] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International conference on machine learning*, pages 12888–12900. PMLR, 2022. 2
- [13] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pages 19730–19742. PMLR, 2023. 2, 4, 6
- [14] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in neural information processing systems*, 36:34892–34916, 2023. 2, 6
- [15] Zijun Long, Richard McCreadie, and Muhammad Imran. Crisisvit: A robust vision transformer for crisis image classification. *arXiv preprint arXiv:2401.02838*, 2024. 2
- [16] Moshood Onifade. Towards an emergency preparedness for self-rescue from underground coal mines. *Process Safety and Environmental Protection*, 149:946–957, 2021. 1
- [17] Bryan A Plummer, Liwei Wang, Chris M Cervantes, Juan C Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *Proceedings of the IEEE international conference on computer vision*, pages 2641–2649, 2015. 1, 2, 5, 6
- [18] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 2, 3, 6
- [19] Narongthart Thanyawet, Photchara Ratsamee, Yuki Uranishi, and Haruo Takemura. Detective networks: Enhancing disaster recognition in images through attention shifting using optimal masking. In *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 8734–8743. IEEE, 2025. 2
- [20] Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4566–4575, 2015. 6
- [21] Vassilios Vonikakis and Stefan Winkler. A center-surround framework for spatial image processing. In *Proc. IS&T Human Vision and Electronic Imaging*, 2016. 3
- [22] Vassilios Vonikakis, Ioannis Andreadis, and Antonios Gasteratos. Fast centre-surround contrast modification. In *IET Image Processing*, pages 19–34, 2008. 3
- [23] Vassilios Vonikakis, Ioannis Arapakis, and Ioannis Andreadis. Combining gray-world assumption, white-point correction and power transformation for automatic white balance. In *International Workshop on Advanced Image Technology (IWAIT)*, 2011. 3
- [24] Vassilios Vonikakis, Odysseas Bouzos, and Ioannis Andreadis. Multi-exposure image fusion based on illumination estimation. In *Signal and Image Processing Applications (SIPA)*, 2011. 3
- [25] Ethan Weber, Dim P Papadopoulos, Agata Lapedriza, Ferda Ofli, Muhammad Imran, and Antonio Torralba. Incidents1m: a large-scale dataset of images with natural disasters, damage, and incidents. *IEEE transactions on pattern analysis and machine intelligence*, 45(4):4768–4781, 2022. 1, 2, 5, 6
- [26] Bin Xiao, Haiping Wu, Weijian Xu, Xiyang Dai, Houdong Hu, Yumao Lu, Michael Zeng, Ce Liu, and Lu Yuan.

- Florence-2: Advancing a unified representation for a variety of vision tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4818–4829, 2024. [6](#)
- [27] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *International conference on machine learning*, pages 2048–2057. PMLR, 2015. [6](#)
- [28] Le Xue, Manli Shu, Anas Awadalla, Jun Wang, An Yan, Senthil Purushwalkam, Honglu Zhou, Viraj Prabhu, Yutong Dai, Michael S Ryoo, et al. xgen-mm (blip-3): A family of open large multimodal models. *arXiv preprint arXiv:2408.08872*, 2024. [2](#)
- [29] Manish Yadav, Mohamed Elmahallawy, Sanjay Madria, and Samuel Frimpong. Predicting battery levels of sensor nodes using reinforcement learning in harsh underground mining environments. In *Proceedings of the 40th ACM/SIGAPP Symposium on Applied Computing*, pages 2048–2057, 2025. [1](#)
- [30] An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025. [6](#)
- [31] Lu Yuan, Dongdong Chen, Yi-Ling Chen, Noel Codella, Xiyang Dai, Jianfeng Gao, Houdong Hu, Xuedong Huang, Boxin Li, Chunyuan Li, et al. Florence: A new foundation model for computer vision. *arXiv preprint arXiv:2111.11432*, 2021. [6](#)
- [32] Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*, 2022. [5](#)
- [33] Jinfei Zhou, Yaping Zhu, Yana Zhang, Cheng Yang, and Hong Pan. Spatial-aware topic-driven-based image chinese caption for disaster news. *Neural Computing and Applications*, 35(13):9481–9500, 2023. [2](#), [5](#), [6](#), [7](#)

Proceedings

**2026 IEEE/CVF Winter Conference on
Applications of Computer Vision
WACV 2026**

Proceedings

**2026 IEEE/CVF Winter Conference on
Applications of Computer Vision
WACV 2026**

**Tucson, Arizona, USA
6-10 March 2026**



IEEE COMPUTER SOCIETY
**CONFERENCE
PUBLISHING
SERVICES**

Los Alamitos, California
Washington • Tokyo



IEEE
**COMPUTER
SOCIETY**

Copyright © 2026 by The Institute of Electrical and Electronics Engineers, Inc.
All rights reserved.

Copyright and Reprint Permissions: Abstracting is permitted with credit to the source. Libraries may photocopy beyond the limits of US copyright law, for private use of patrons, those articles in this volume that carry a code at the bottom of the first page, provided that the per-copy fee indicated in the code is paid through the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923.

Other copying, reprint, or republication requests should be addressed to: IEEE Copyrights Manager, IEEE Service Center, 445 Hoes Lane, P.O. Box 133, Piscataway, NJ 08855-1331.

The papers in this book comprise the proceedings of the meeting mentioned on the cover and title page. They reflect the authors' opinions and, in the interests of timely dissemination, are published as presented and without change. Their inclusion in this publication does not necessarily constitute endorsement by the editors, the IEEE Computer Society, or the Institute of Electrical and Electronics Engineers, Inc.

ISBN-13: 979-8-3315-5511-5

ISSN: 2642-9381

BMS Part # CFP26082-ART

Additional copies may be ordered from:

IEEE Computer Society
Customer Service Center
10662 Los Vaqueros Circle
P.O. Box 3014
Los Alamitos, CA 90720-1314
Tel: +1 800 272 6657
Fax: +1 714 821 4641
<http://computer.org/cspress>
csbooks@computer.org

IEEE Service Center
445 Hoes Lane
P.O. Box 1331
Piscataway, NJ 08855-1331
Tel: +1 732 981 0060
Fax: +1 732 981 9667
[http://shop.ieee.org/store/
customer-service@ieee.org](http://shop.ieee.org/store/customer-service@ieee.org)

IEEE Computer Society
Asia/Pacific Office
Watanabe Bldg., 1-4-2
Minami-Aoyama
Minato-ku, Tokyo 107-0062
JAPAN
Tel: +81 3 3408 3118
Fax: +81 3 3408 3553
tokyo.ofc@computer.org

Individual paper REPRINTS may be ordered at: <reprints@computer.org>

Editorial production by Cristina Ceballos
Cover art production by Hector Torres



**IEEE
COMPUTER
SOCIETY**



**IEEE COMPUTER SOCIETY
CONFERENCE
PUBLISHING
SERVICES**

*IEEE Computer Society
Conference Publishing Services (CPS)*
<http://www.computer.org/cps>