

# FUME: Fused Unified Multi-Gas Emission Network for Livestock Rumen Acidosis Detection

Taminul Islam<sup>1</sup>, Toqi Tahamid Sarker<sup>1</sup>, Mohamed Embaby<sup>2</sup>, Khaled R Ahmed<sup>1</sup>, Amer AbuGhazaleh<sup>1</sup>

<sup>1</sup>Southern Illinois University, Carbondale, <sup>2</sup>University of California, Davis

{taminul.islam, toqitahamid.sarker, khaled.ahmed, aabugha}@siu.edu, membaby@ucdavis.edu

## Abstract

*Ruminal acidosis is a prevalent metabolic disorder in dairy cattle causing significant economic losses and animal welfare concerns. Current diagnostic methods rely on invasive pH measurement, limiting scalability for continuous monitoring. We present FUME (Fused Unified Multi-gas Emission Network), the first deep learning approach for rumen acidosis detection from dual-gas optical imaging under in vitro conditions. Our method leverages complementary carbon dioxide (CO<sub>2</sub>) and methane (CH<sub>4</sub>) emission patterns captured by infrared cameras to classify rumen health into Healthy, Transitional, and Acidotic states. FUME employs a lightweight dual-stream architecture with weight-shared encoders, modality-specific self-attention, and channel attention fusion, jointly optimizing gas plume segmentation and classification of dairy cattle health. We introduce the first dual-gas Optical Gas Imaging (OGI) dataset comprising 8,967 annotated frames across six pH levels with pixel-level segmentation masks. Experiments demonstrate that FUME achieves 80.99% mIoU and 98.82% classification accuracy while using only 1.28M parameters and 1.97G MACs—outperforming state-of-the-art methods in segmentation quality with 10× lower computational cost. Ablation studies reveal that CO<sub>2</sub> provides the primary discriminative signal and dual-task learning is essential for optimal performance. Our work establishes the feasibility of gas emission-based livestock health monitoring, paving the way for practical, in vitro acidosis detection systems. Codes are available at <https://github.com/taminulislam/fume>.*

## 1. Introduction

Ruminal acidosis represents one of the most economically significant metabolic disorders in modern dairy and beef cattle production, affecting an estimated 19–26% of dairy cows in intensive farming systems worldwide [28, 39]. This condition, characterized by abnormally low rumen pH (typ-

ically below 5.8), results from the accumulation of volatile fatty acids, lactic acid in particular, during rapid fermentation of high-concentrate diets. The consequences extend beyond immediate digestive dysfunction: chronic acidosis leads to reduced feed intake, decreased milk production (losses of 1–3 kg/day), laminitis, liver abscesses, systemic inflammation, and in severe cases, mortality [26]. Subacute ruminal acidosis costs the North American dairy industry an estimated \$500 million to \$1 billion annually, with per-cow losses of approximately \$400–475 per lactation [7], figures that likely underestimate current losses given inflation and increased production intensity [38].

Early detection of acidosis is critical for timely intervention and prevention of irreversible health damage. However, current diagnostic methods present significant practical limitations. The gold standard—direct rumen pH measurement via rumenocentesis or indwelling bolus sensors—is invasive, expensive, requires veterinary intervention, and causes animal stress [10]. Indirect indicators such as milk composition, fecal pH, and behavioral monitoring lack the sensitivity and specificity required for reliable early-stage detection [4]. Consequently, acidosis often remains undiagnosed until clinical signs become apparent, by which point substantial metabolic damage has already occurred. The agricultural industry urgently needs real-time and cost-effective tools for continuous rumen health monitoring at scale.

Rumen fermentation produces two primary gaseous products: carbon dioxide (CO<sub>2</sub>) and methane (CH<sub>4</sub>), whose production rates and relative proportions are directly influenced by rumen pH and microbial activity [15]. Under acidotic conditions, methanogenic archaea activity is suppressed, reducing CH<sub>4</sub> production, while CO<sub>2</sub> generation from acidic fermentation pathways increases [23]. These characteristic gas emission patterns, captured by infrared optical gas imaging (OGI) cameras, encode valuable information about underlying rumen health status. However, translating raw thermal imagery into actionable health diagnostics requires sophisticated pattern recognition beyond human visual interpretation—a challenge ideally suited to deep learning.




Figure 1. Efficiency vs. segmentation quality trade-off across all compared methods including OGI-specific baselines. The upper-left region represents optimal performance (high mIoU, low computational cost). FUME achieves the best trade-off, delivering state-of-the-art segmentation quality while requiring significantly fewer MACs than competing methods.

Before deploying gas emission monitoring systems in complex farm environments with variable conditions, free-moving animals, and multiple confounding factors, it is essential to first validate the fundamental feasibility of pH prediction from gas signatures under controlled conditions. In vitro fermentation systems, such as the ANKOM RF Gas Production System [6], provide precisely regulated pH environments with continuous gas monitoring, enabling the establishment of ground-truth relationships between pH levels and observable gas emission patterns. This controlled validation is critical for several reasons: in vitro systems eliminate confounding variables present in live animals—diet variation, individual metabolic differences, stress responses, and environmental fluctuations—allowing isolation of the pH-gas emission relationship. They also enable systematic collection of data across the full pH spectrum, including severe acidotic states ( $\text{pH} < 5.0$ ) that would be unethical to induce in live animals. Furthermore, they provide the labeled training data necessary for supervised deep learning, as ground-truth pH measurements can be precisely controlled and verified.

In this paper, we present the first deep learning approach for rumen acidosis detection from dual-gas optical imaging under controlled laboratory conditions. Our method leverages the complementary information encoded in simultaneous  $\text{CO}_2$  and  $\text{CH}_4$  emissions to predict rumen health status across three clinically relevant categories: Healthy, Transitional, and Acidotic. We formulate this as a multi-task learning problem, jointly optimizing pixel-level gas

plume segmentation and image-level livestock health classification. The segmentation task forces the model to learn spatially precise representations of gas distributions, which in turn improves classification by focusing on emission-specific features rather than irrelevant background or equipment artifacts. Our dual-stream architecture processes  $\text{CO}_2$  and  $\text{CH}_4$  channels independently before fusing them via channel attention, enabling the model to learn gas-specific patterns as well as their interactions. As illustrated in Fig. 1, FUME achieves the best accuracy-efficiency trade-off among all compared methods, occupying the optimal upper-left region with highest segmentation quality at minimal computational cost.

The main contributions of this work are twofold: (1) we introduce the first dual-gas OGI dataset for rumen health assessment, comprising 8,967 annotated frame pairs ( $\text{CO}_2$  and  $\text{CH}_4$ ) across six pH levels (Table 2), with pixel-level segmentation masks and image-level health labels—a novel resource for the agricultural computer vision community; and (2) we propose a dual-stream multi-task deep learning architecture that jointly performs gas plume segmentation and health classification, establishing the feasibility of rumen pH inference from gas emission patterns under in vitro conditions and paving the way for practical, scalable livestock health monitoring systems.

## 2. Related Work

Our work intersects three research areas: gas detection and environmental monitoring, multi-modal fusion and multi-task learning for vision tasks, and agricultural computer vision. We review recent advances in each domain and position our contributions relative to the state-of-the-art.

**Gas Detection and Environmental Monitoring.** Recent work has addressed gas plume segmentation in OGI domains. Gasformer [32] proposes lightweight transformer hybrids for  $\text{CH}_4$  detection with attention to boundary quality. GasTwinFormer [35] combines segmentation with classification in multi-task OGI. CarboFormer [14] addresses  $\text{CO}_2$  plume segmentation. Other progress includes methane leak detection [22], wildfire smoke detection [40], and atmospheric pollutant estimation [34]. While these methods target single-gas industrial scenarios, FUME addresses dual-gas fusion for metabolic state inference—a fundamentally different problem requiring cross-modal reasoning for health assessment rather than binary presence detection.

**Multi-Modal Fusion and Multi-Task Learning.** Multi-modal fusion has proven effective for tasks requiring complementary information sources across diverse vision applications. Cross-modal attention mechanisms for RGB-thermal fusion in semantic segmentation demonstrate that learned attention outperforms fixed fusion rules [9]—a principle we adopt for  $\text{CO}_2$ - $\text{CH}_4$  fusion. Hierarchical fusion approaches for RGB-D salient object detection show that

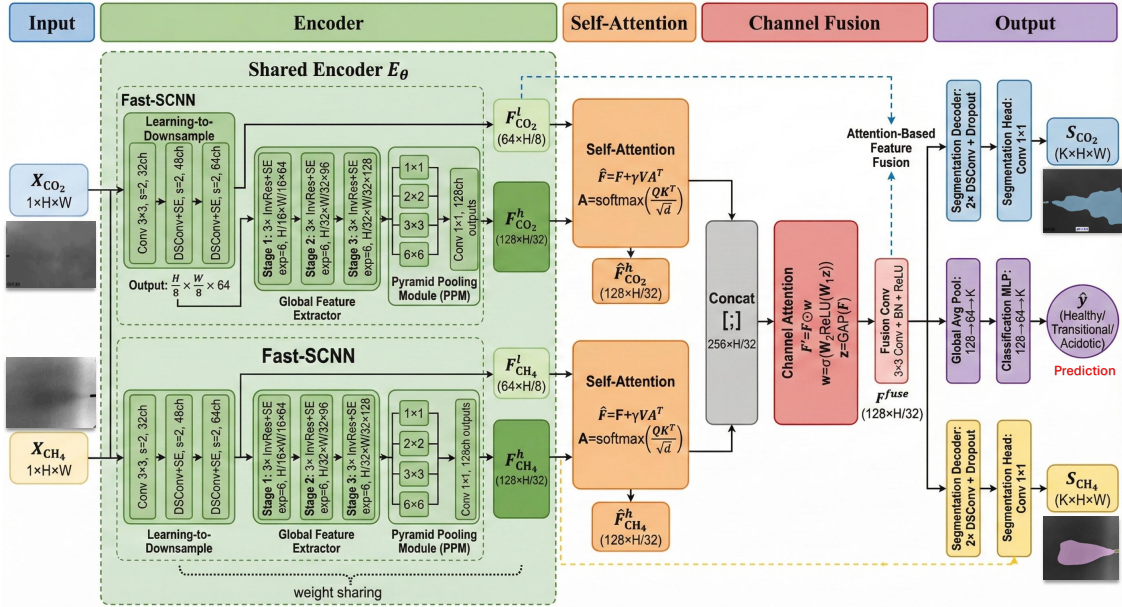


Figure 2. **FUME Architecture.** Our dual-stream network processes paired  $\text{CO}_2$  and  $\text{CH}_4$  thermal images through a weight-shared Fast-SCNN encoder with depthwise separable convolutions. The Learning-to-Downsample module reduces spatial resolution to  $\frac{H}{8} \times \frac{W}{8}$ , while the Global Feature Extractor with Pyramid Pooling Module captures multi-scale context at  $\frac{H}{32} \times \frac{W}{32}$ . Modality-specific self-attention modules refine each stream before channel attention fusion combines the representations. The fused features feed dual segmentation decoders and a classification head.

multi-scale aggregation improves robustness to modality-specific artifacts [2], while dual-stream architectures for multi-modal tumor segmentation validate the effectiveness of parallel encoders with late fusion for heterogeneous inputs [12]. Adaptive fusion networks that dynamically weight modality contributions based on input quality [3] and comparisons of feature-level versus decision-level fusion [41] provide insights that inform our fusion module design. Concurrently, multi-task learning (MTL) has emerged as a powerful paradigm for leveraging task synergies, with end-to-end MTL frameworks for autonomous driving [42], task-balanced learning strategies addressing conflicting gradients [21], joint segmentation and classification of medical lesions demonstrating complementary supervision signals [19], multi-task attention networks for scene understanding [43], and efficient MTL architectures for resource-constrained deployment [31]. Despite these advances, existing fusion and MTL approaches target RGB, depth, or thermal modalities rather than multi-gas emissions with distinct physical meanings, leaving the domain of gas-based health inference unexplored.

**Agricultural Computer Vision.** Computer vision for livestock monitoring has focused primarily on external observations, including deep learning for behavior recognition [46], automated body condition scoring [1], lameness detection from gait analysis [20], and vision-based health mon-

itoring systems [16]. In crop and weed monitoring, researchers have developed precision agriculture systems using multi-spectral imagery for disease detection [25] and weed segmentation with hierarchical vision transformers [13, 33]. While attention mechanisms have revolutionized computer vision through cross-modal learning [18] and improved transformer architectures for dense prediction [36], their data requirements make them less suitable for limited annotated data—motivating our CNN-based approach with targeted attention modules. Critically, all existing agricultural vision methods observe external animal characteristics and cannot directly assess internal metabolic states such as rumen pH.

**Research Gap.** Despite significant advances in gas detection, multi-modal fusion, and agricultural computer vision, a critical gap remains: no existing work has explored the use of gas emission imaging for livestock health assessment. Industrial gas detection systems focus on binary presence detection rather than physiological state inference. Agricultural vision systems monitor external observations but cannot assess internal metabolic conditions. Our work bridges this gap by combining dual-gas fusion with multi-task learning specifically for rumen health monitoring under in vitro conditions, demonstrating that  $\text{CO}_2$  and  $\text{CH}_4$  emission patterns captured through OGI technology encode sufficient information to distinguish healthy, transitional,

Table 1. Quantitative comparison with state-of-the-art methods. † indicates best performance. For HD95, ASD, and Latency, lower is better.

Model	Classification	Segmentation		Boundary		Efficiency			
	Acc (%)	mIoU (%)	Dice (%)	HD95 (px)	ASD (px)	Params (M)	MACs (G)	Latency (ms)	FPS
BiSeNetV2 [44]	<b>99.04</b> †	75.91	85.85	72.13	20.89	2.77	21.01	3.54	282.49
ENet [27]	97.86	46.67	54.13	101.16	34.24	<b>0.26</b> †	3.18	2.71	369.00
DDRNetSlim [11]	94.98	56.97	70.55	140.50	46.13	3.39	16.66	<b>1.66</b> †	<b>602.41</b> †
ESPNetV2 [24]	90.81	51.88	58.00	92.32	26.55	2.22	18.39	2.58	387.60
DANet [8]	48.50	47.23	54.49	103.85	34.65	1.43	5.92	2.99	334.45
RTFNet [30]	46.47	48.56	55.55	107.90	33.22	35.01	50.28	3.88	257.73
MTINet [37]	46.37	44.53	52.27	113.11	39.20	1.80	3.32	3.18	314.47
CMX [45]	33.65	26.69	29.64	187.34	61.28	4.41	13.18	1.98	505.05
Gasformer [32]	89.56	62.09	74.31	42.51	10.56	3.65	5.11	5.63	177.50
GasTwinFormer [35]	86.27	47.80	55.77	47.77	14.18	3.35	4.51	6.32	158.20
CarboFormer [14]	87.92	54.95	65.04	41.14	10.37	3.81	3.09	9.31	107.40
<b>FUME (Ours)</b>	98.82	<b>80.99</b> †	<b>89.23</b> †	<b>46.58</b> †	<b>14.00</b> †	1.28	<b>1.97</b> †	3.06	326.80

and acidotic states—establishing proof-of-concept for gas-based metabolic assessment. Specifically, we address the efficiency limitations of existing multi-modal fusion approaches by proposing channel attention fusion instead of computationally expensive cross-modal attention, achieving state-of-the-art segmentation quality with significantly reduced computational cost suitable for real-time edge deployment.

### 3. Method

We present FUME (Fused Unified Multi-gas Emission Network), a lightweight dual-stream architecture for acidosis detection from paired CO<sub>2</sub> and CH<sub>4</sub> thermal emissions. While individual components (Fast-SCNN, self-attention, channel attention) are established, our key contributions are: (1) *novel problem formulation*—first to demonstrate that dual-gas emission patterns encode metabolic state, enabling non-invasive health inference; (2) *physiologically-motivated fusion*—channel attention over cross-attention based on our finding that CO<sub>2</sub> and CH<sub>4</sub> lack pixel-wise correspondence but exhibit correlated aggregate statistics; and (3) *multi-task synergy*—joint segmentation and classification where spatial precision improves classification focus.

#### 3.1. Problem Formulation

Given paired thermal imaging frames  $\mathbf{X}_{\text{CO}_2}, \mathbf{X}_{\text{CH}_4} \in \mathbb{R}^{1 \times H \times W}$  capturing CO<sub>2</sub> and CH<sub>4</sub> gas emissions, our objective is to jointly predict: (i) pixel-wise segmentation masks  $\mathbf{M}_{\text{CO}_2}, \mathbf{M}_{\text{CH}_4} \in \{0, 1, 2\}^{H \times W}$  identifying background, tube, and gas emission regions, and (ii) a global classification label  $y \in \{\text{Healthy}, \text{Transitional}, \text{Acidotic}\}$  indicating the metabolic acidosis state. The fundamental challenge lies in the **cross-modal semantic gap**: while CO<sub>2</sub> and CH<sub>4</sub> concentrations are physiologically correlated, their spatial distributions differ substantially. We

propose a solution that captures inter-modal dependencies through efficient channel-wise attention while preserving intra-modal spatial structure via self-attention.

#### 3.2. Architecture Overview

FUME comprises four components: (1) a weight-shared Fast-SCNN [29] encoder, (2) modality-specific self-attention modules, (3) a channel attention fusion mechanism, and (4) dual-task decoders. Fig. 2 illustrates the architecture. The forward pass is:

$$\mathbf{F}_m^l, \mathbf{F}_m^h = \mathcal{E}_\theta(\mathbf{X}_m), \quad \hat{\mathbf{F}}_m^h = \text{SelfAttn}_m(\mathbf{F}_m^h) \quad (1)$$

$$\mathbf{F}^{\text{fuse}} = \text{ChannelAttn}([\hat{\mathbf{F}}_{\text{CO}_2}^h; \hat{\mathbf{F}}_{\text{CH}_4}^h]) \quad (2)$$

where  $\mathbf{F}_m^l \in \mathbb{R}^{64 \times \frac{H}{8} \times \frac{W}{8}}$  and  $\mathbf{F}_m^h \in \mathbb{R}^{128 \times \frac{H}{32} \times \frac{W}{32}}$  denote low-level and high-level features respectively.

#### 3.3. Dual-Stream Encoder and Fusion

**Shared Encoder.** A critical design decision is weight sharing between CO<sub>2</sub> and CH<sub>4</sub> encoder streams. This is motivated by two observations: (i) both thermal modalities share common low-level visual primitives (edges, textures, intensity gradients) that benefit from joint representation learning, and (ii) weight sharing reduces encoder parameters by 50%, enabling deployment on resource-constrained devices. We adopt Fast-SCNN [29] as our backbone, which employs a learning-to-downsample module achieving 8× downsampling through depthwise separable convolutions:

$$\mathbf{F}^l = \text{DSCConv}_{64}(\text{DSCConv}_{48}(\text{Conv}_{32}(\mathbf{X}))) \quad (3)$$

The computational efficiency of depthwise separable convolutions provides approximately 8× reduction compared to standard convolutions. The global feature extractor applies inverted residual blocks with expansion ratio 6 across

Table 2. Dataset distribution across pH levels, health classes, and splits.

pH	Class	Total	CO <sub>2</sub>	CH <sub>4</sub>	Train (70%)	Val (15%)	Test (15%)
6.5	Healthy	1,008	693	315	706	151	151
6.2	Healthy	1,050	721	329	735	158	157
5.9	Transitional	945	651	294	662	142	141
5.6	Acidotic	3,087	2,121	966	2,161	463	463
5.3	Acidotic	1,512	1,039	473	1,058	227	227
5.0	Acidotic	1,365	938	427	955	205	205
<b>Total</b>		8,967	6,163	2,804	6,277	1,346	1,344

three bottleneck stages, progressively extracting semantic features at  $\frac{1}{16}$  and  $\frac{1}{32}$  resolution. A Pyramid Pooling Module aggregates multi-scale context through adaptive pooling at scales  $\{1, 2, 3, 6\}$ , enabling the network to capture both local gas emission patterns and global distribution characteristics.

**Self-Attention.** At the deepest feature level ( $\frac{H}{32} \times \frac{W}{32}$ ), we apply self-attention independently to each modality stream. This addresses a fundamental limitation of convolutional encoders: while convolutions excel at capturing local patterns, gas emission analysis requires modeling long-range spatial dependencies—a diffuse gas plume may span the entire field of view, and its shape provides diagnostic information that local receptive fields cannot capture:

$$\mathbf{A}_m = \text{softmax} \left( \frac{\mathbf{Q}_m \mathbf{K}_m^\top}{\sqrt{d_k}} \right), \quad \hat{\mathbf{F}}_m^h = \mathbf{F}_m^h + \gamma_m \cdot \mathbf{V}_m \mathbf{A}_m^\top \quad (4)$$

where  $\gamma_m$  is a learnable scalar initialized to zero for training stability. Critically, we apply self-attention *separately* to each modality rather than cross-modal attention. This decision is grounded in our empirical finding that CO<sub>2</sub> and CH<sub>4</sub> emissions, while correlated in aggregate statistics, do not exhibit pixel-wise spatial correspondence—their relationship is better captured through channel-wise statistics.

**Channel Attention Fusion.** After self-attention refinement, we concatenate modality-specific features and apply channel attention to model inter-gas dependencies. Unlike spatial cross-attention which asks “which CO<sub>2</sub> locations correspond to which CH<sub>4</sub> locations,” channel attention asks “which feature channels from CO<sub>2</sub> should be combined with which channels from CH<sub>4</sub>”—a more appropriate inductive bias for our task:

$$\mathbf{z} = \text{GAP}([\hat{\mathbf{F}}_{\text{CO}_2}^h; \hat{\mathbf{F}}_{\text{CH}_4}^h]), \quad \mathbf{w} = \sigma(\mathbf{W}_2 \cdot \text{ReLU}(\mathbf{W}_1 \cdot \mathbf{z})) \quad (5)$$

$$\mathbf{F}^{\text{fuse}} = \text{Conv}_{3 \times 3}([\hat{\mathbf{F}}_{\text{CO}_2}^h; \hat{\mathbf{F}}_{\text{CH}_4}^h] \odot \mathbf{w}) \quad (6)$$

where  $\mathbf{W}_1, \mathbf{W}_2$  form a bottleneck MLP with reduction ratio 16. Global average pooling aggregates statistics over the

Table 3. Per-class F1 scores (%) and balanced accuracy.

Model	Healthy	Trans.	Acidotic	Macro F1	Bal. Acc
BiSeNetV2 [44]	<b>99.04</b>	<b>94.59</b>	99.32	<b>97.65</b>	<b>99.24</b>
ENet [27]	97.91	80.46	99.14	92.50	98.32
DDRNetSlim [11]	93.96	69.31	97.75	87.01	96.01
ESPNetV2 [24]	87.80	81.97	93.18	87.65	85.71
DANet [8]	55.90	75.61	33.99	55.17	68.41
RTFNet [30]	55.70	0.87	33.99	29.90	40.16
MTINet [37]	55.65	1.12	33.76	29.80	40.10
CMX [45]	50.36	1.23	1.45	16.79	33.33
GasFormer [32]	88.00	82.50	88.00	86.16	89.56
GasTwinFormer [35]	86.00	79.00	87.00	83.98	86.27
CarboFormer [14]	87.00	81.00	87.20	85.07	87.92
<b>FUME (Ours)</b>	98.72	94.59	<b>99.15</b>	97.49	99.03

entire spatial extent, enabling the network to learn correlations like “high mean CO<sub>2</sub> intensity combined with low CH<sub>4</sub> variance indicates acidotic state.”

### 3.4. Dual-Task Decoding and Training

**Segmentation.** Separate decoding heads for CO<sub>2</sub> and CH<sub>4</sub> employ Feature Fusion Modules combining high-level features with low-level encoder features via skip connections:

$$\mathbf{S}_m = \text{Classifier}(\text{FFM}(\mathbf{F}_m^l, \hat{\mathbf{F}}_m^h)) \quad (7)$$

The classifier consists of depthwise separable convolutions followed by a  $1 \times 1$  projection to  $K = 3$  classes.

**Classification.** Global pH state prediction aggregates  $\mathbf{F}^{\text{fuse}}$  via adaptive average pooling through a two-layer MLP with dropout.

**Training Objective.** We jointly optimize through multi-task loss:

$$\mathcal{L} = \mathcal{L}_{\text{seg}}^{\text{CO}_2} + \mathcal{L}_{\text{seg}}^{\text{CH}_4} + \lambda \mathcal{L}_{\text{cls}} \quad (8)$$

where each  $\mathcal{L}_{\text{seg}}^m$  combines Focal loss and Dice loss ( $0.5 \cdot \mathcal{L}_{\text{focal}} + 0.5 \cdot \mathcal{L}_{\text{dice}}$ ) to handle class imbalance and boundary precision,  $\mathcal{L}_{\text{cls}}$  is Focal loss with class weights, and  $\lambda = 0.5$ . Self-attention uses single-head with  $d_k = C/8 = 16$  and learnable  $\gamma$  initialized to zero; channel attention uses reduction ratio 16. FUME achieves 1.28M parameters and 1.97G MACs as seen in Fig. 1 through weight sharing and efficient fusion.

## 4. Experiments

### 4.1. Dataset

Rumen fluid was collected from fistulated cattle following protocols, and fermentation substrates were prepared using standard in vitro rumen fermentation procedures with the ANKOM [5]. Samples were maintained at six controlled pH levels (6.5, 6.2, 5.9, 5.6, 5.3, and 5.0), representing the spectrum from severe acidosis to healthy rumen conditions. We captured gas emissions using two specialized FLIR OGI cameras: the FLIR GF343 for CO<sub>2</sub> detection (4.2–4.4  $\mu\text{m}$

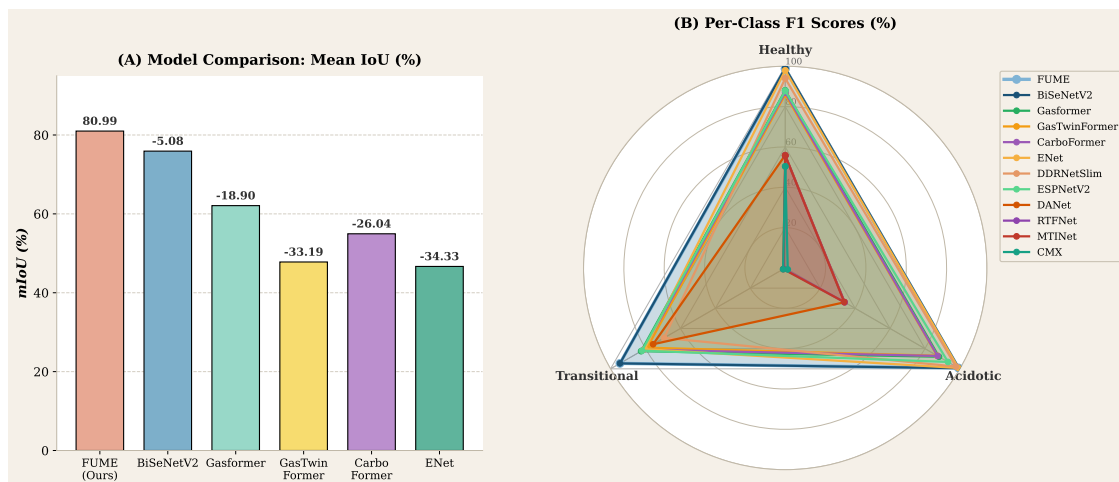


Figure 3. Comparison with state-of-the-art methods. (A) Mean IoU comparison showing FUME substantially outperforms both general segmentation baselines and OGI-specific methods (Gasformer, GasTwinFormer, CarboFormer). (B) Per-class F1 radar chart demonstrating balanced classification performance across Healthy, Transitional, and Acidotic classes, where FUME maintains consistent high performance while several baselines exhibit severe degradation on challenging classes.

mid-wave infrared) [14], and the FLIR GF320 for  $\text{CH}_4$  detection (3.2–3.4  $\mu\text{m}$ ) [35]. Both cameras were co-located at 20 cm from the gas source, recording simultaneously at  $256 \times 256$  resolution; frames are bilinearly upsampled to  $512 \times 512$  for training.

We map pH levels to three clinically relevant health classes based on established rumen physiology [28]: **Healthy** ( $\text{pH} \geq 6.0$ ), **Transitional** ( $5.8 \leq \text{pH} < 6.0$ ), and **Acidotic** ( $\text{pH} < 5.8$ ). Since gas emissions are intermittent, we extracted only frames containing visible gas plumes and applied data augmentation. Each frame was manually annotated with pixel-level masks ( $K=3$  classes: background, tube, gas), with gas plumes occupying only 5.3% of pixels. Table 2 shows the dataset distribution with stratified 70-15-15 splits; data was partitioned by pH level to prevent temporal leakage. Additionally, frames from the same fermentation session were kept within a single split to ensure independence between training and evaluation sets. The unequal  $\text{CO}_2$  vs  $\text{CH}_4$  counts reflect physiological differences:  $\text{CH}_4$  emissions are less visible at acidotic pH. Unpaired samples use zero-padding with a modality mask indicating available inputs.

## 4.2. Baseline Models and Implementation Details

We compare FUME against state-of-the-art efficient segmentation architectures including BiSeNetV2 [44], ENet [27], DDRNetSlim [11], ESPNetV2 [24], DANet [8], RTFNet [30], MTINet [37], CMX [45], and OGI-specific methods: Gasformer [32], GasTwinFormer [35], and CarboFormer [14]. Single-stream baselines receive 2-channel concatenated inputs ( $\text{CO}_2 + \text{CH}_4$ ); multimodal and OGI baselines use dual-stream configuration matching FUME.

All models are implemented in PyTorch and trained on NVIDIA A6000 GPU using AdamW optimizer ( $\text{lr}=0.001$ ), batch size 16, 20 epochs with cosine annealing. Data augmentation includes random flipping, rotation ( $\pm 15^\circ$ ), and color jittering. Latency is measured with batch size 1, FP32 precision, CUDA synchronized timing (100 warmup + 1000 iterations), `cudaNN.benchmark=True`. mIoU is computed per-modality across 3 classes and averaged over  $\text{CO}_2$  and  $\text{CH}_4$  heads.

## 4.3. Evaluation Metrics

**Classification:** Accuracy and per-class F1 scores (Macro F1 averaged across classes).

**Segmentation:** Mean IoU and Dice coefficient, where  $\text{IoU} = |P \cap G| / |P \cup G|$  and  $P, G$  denote predicted and ground-truth masks.

**Boundary Quality:** Hausdorff Distance 95<sup>th</sup> percentile (HD95) [17] and Average Surface Distance (ASD) [17]; lower values indicate better boundary delineation.

**Efficiency:** Parameters (M), MACs (G), FPS, and GPU latency (ms).

## 5. Results

### 5.1. Comparison with State-of-the-Art Methods

Table 1 presents comprehensive quantitative comparisons including OGI-specific methods. FUME achieves the best segmentation quality with 80.99% mIoU, outperforming both general segmentation baselines and OGI-specific methods: Gasformer [32] (62.09%), CarboFormer [14] (54.95%), and GasTwinFormer [35] (47.80%). FUME also achieves the best boundary metrics (HD95: 46.58, ASD:

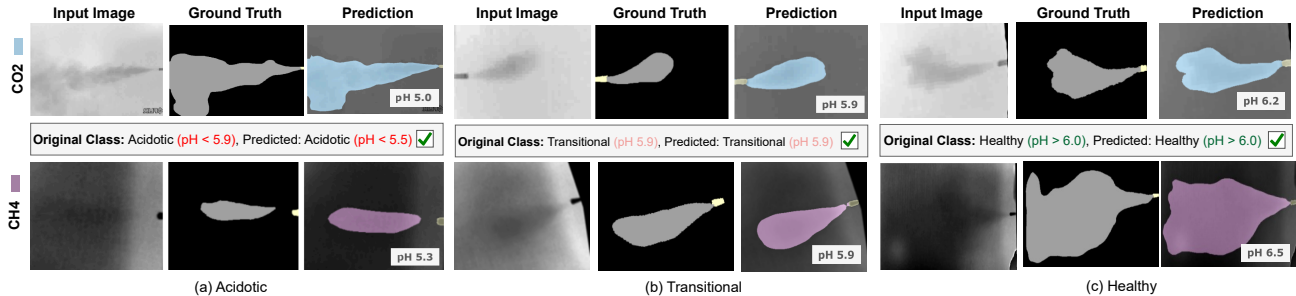


Figure 4. Qualitative segmentation results across Acidotic, Transitional, and Healthy classes for both  $\text{CO}_2$  and  $\text{CH}_4$  modalities. FUME achieves precise boundary delineation of gas plumes with predictions closely matching ground truth annotations, demonstrating robustness to varying plume morphologies and emission intensities.

14.00) and high classification accuracy (98.82%), substantially higher than OGI methods (86–90%), while using fewer parameters (1.28M vs 3.3–3.8M) and lowest MACs (1.97G). At 326.80 FPS, FUME is 1.8–3 $\times$  faster than OGI baselines, making it suitable for real-time edge deployment.

Fig. 3 provides visual comparison. The bar chart (A) illustrates the mIoU gap between methods, showing FUME achieves 80.99% while others fall short by 5.08–54.30 pp. The radar chart (B) compares per-class F1 scores across all methods, demonstrating FUME’s balanced performance across Healthy, Transitional, and Acidotic classes, while several baselines exhibit severe degradation on challenging classes. Detailed per-class metrics in Table 3 show FUME achieves 97.49% macro F1 and 99.03% balanced accuracy, with the highest Acidotic class F1 (99.15%), critical for acidosis detection.

## 5.2. Efficiency Analysis

Fig. 1 visualizes the efficiency-quality trade-off. FUME occupies the optimal upper-left region (highest mIoU, minimal MACs). OGI-specific methods (Gasformer [32], CarboFormer [14], GasTwinFormer [35]) cluster in the mid-range with 3–5G MACs and 47–62% mIoU, demonstrating that single-gas architectures struggle with dual-modality fusion. BiSeNetV2 [44] achieves good mIoU but requires 10 $\times$  more computation. FUME uniquely achieves state-of-the-art segmentation at 326.80 FPS with lowest computational cost.

## 5.3. Ablation Studies

To validate our architectural design choices, we conduct systematic ablation studies. Table 4 and Fig. 5 present the results of removing or modifying key components.

**Effect of Attention Mechanisms.** A counter-intuitive finding is that adding attention mechanisms—whether full cross-modal attention or self-attention only—degrades segmentation quality (−1.68 pp and −0.63 pp mIoU respectively) while significantly increasing inference latency

Table 4. Ablation study results.  $\Delta$  denotes change relative to FUME.

Variant	Acc (%)	mIoU (%)	Dice (%)	Latency	$\Delta$ mIoU
+ Full Cross-Modal Attn	99.25	79.31	88.15	4.55	−1.68
+ Self-Attention Only	98.82	80.36	88.81	3.55	−0.63
w/o $\text{CH}_4$ ( $\text{CO}_2$ Only)	98.61	80.52	88.92	1.78	−0.47
w/o $\text{CO}_2$ ( $\text{CH}_4$ Only)	48.29	—	—	1.81	—
Classification Only	97.86	—	—	3.23	—
Segmentation Only	33.65	71.21	82.67	2.95	−9.78
<b>FUME (Proposed)</b>	<b>98.82</b>	<b>80.99</b>	<b>89.23</b>	<b>3.06</b>	—

(+48% and +16%). This challenges the common assumption that attention always benefits multi-modal fusion. We hypothesize that for tasks where modalities lack pixel-wise spatial correspondence, direct feature concatenation may provide sufficient alignment; however, this finding may be sensitive to attention design choices and hyperparameter tuning, warranting further investigation across different attention variants.

**Modality Contributions.** Each gas modality contributes distinct information: removing  $\text{CO}_2$  ( $\text{CH}_4$  Only) causes catastrophic classification failure (48.29% accuracy), confirming  $\text{CO}_2$  as the primary discriminative signal for acidosis detection. Conversely, removing  $\text{CH}_4$  ( $\text{CO}_2$  Only) degrades segmentation quality (−0.47 pp mIoU, −0.31 pp Dice), indicating that  $\text{CH}_4$ ’s spatial diffusion patterns enhance boundary delineation. This asymmetric yet complementary relationship— $\text{CO}_2$  for discriminative power,  $\text{CH}_4$  for spatial refinement—validates our dual-modality design.

**Task Interaction.** Both classification and segmentation branches are necessary for optimal performance. Using only the classification branch achieves high accuracy (97.86%) but has no segmentation capability (no segmentation head), while the segmentation-only variant fails at classification (33.65% accuracy) despite maintaining reasonable segmentation (71.21% mIoU). This validates our dual-task learning framework where both tasks provide complementary supervisory signals.

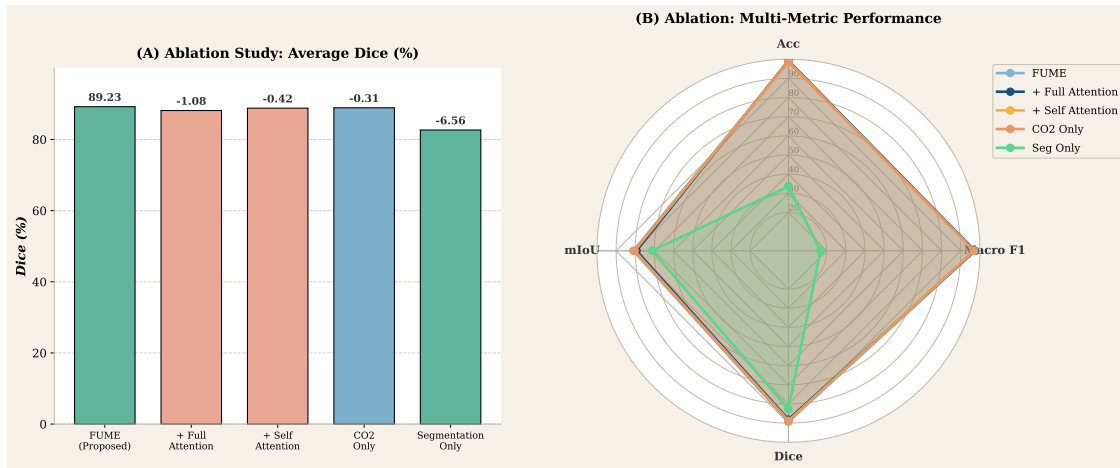


Figure 5. Ablation study results. (A) Dice coefficient comparison across architectural variants, demonstrating that the proposed FUME configuration achieves optimal segmentation performance. (B) Multi-metric radar chart visualizing the critical importance of the CO<sub>2</sub> modality and dual-task learning framework for achieving balanced performance across accuracy, segmentation quality, and classification metrics.

Fig. 4 shows qualitative results across all three classes. FUME achieves excellent boundary delineation across varying morphologies—irregular acidotic patterns, intermediate transitional states, and larger healthy regions—maintaining sharp plume boundaries even with diffuse emissions and low contrast.

## 5.4. Discussion

**Design Trade-offs.** Our results reveal a trade-off between classification accuracy and segmentation quality in attention-based fusion. Attention mechanisms boost classification (+0.43% accuracy) but degrade segmentation (−1.68 pp mIoU). For applications requiring precise boundary delineation, FUME without attention provides superior performance.

**Modality Importance.** CO<sub>2</sub> and CH<sub>4</sub> encode physiologically distinct metabolic signals: CO<sub>2</sub> reflects fermentation rate while CH<sub>4</sub> indicates methanogenic activity, which is suppressed under acidotic conditions. Our ablation confirms this complementarity—CO<sub>2</sub>-only achieves strong classification but CH<sub>4</sub> fusion improves segmentation quality (+0.47 pp mIoU, +0.31 pp Dice), particularly for boundary delineation where CH<sub>4</sub>'s diffusion patterns provide additional spatial cues. Crucially, CH<sub>4</sub>-only fails entirely (48.29% accuracy), validating that both modalities contribute non-redundant information: CO<sub>2</sub> provides discriminative power while CH<sub>4</sub> refines spatial precision. This dual-gas framework also establishes a foundation for future multi-modal livestock monitoring systems.

**Limitations.** This study validates FUME under controlled in vitro conditions; translating to in vivo settings will require addressing animal motion, environmental variability,

and intermittent plume visibility. Future work will explore temporal modeling, domain adaptation, and validation with live animals under farm conditions.

## 6. Conclusion

We presented FUME, a lightweight dual-stream architecture for rumen acidosis detection from dual-gas optical imaging that jointly performs gas plume segmentation and health classification through weight-shared encoders and channel attention fusion. Our experiments demonstrate state-of-the-art segmentation quality and classification accuracy with exceptional computational efficiency, outperforming both general segmentation baselines and OGI-specific methods. Key findings reveal that CO<sub>2</sub> serves as the primary biomarker for acidosis detection, dual-task learning provides essential complementary supervision, and channel attention fusion outperforms expensive cross-modal attention while maintaining real-time inference capability. Our work establishes proof-of-concept for gas emission-based livestock health monitoring, with future directions including in vivo validation and temporal modeling. The dataset and code will be released upon acceptance to facilitate future research.

## 7. Acknowledgment

This work is supported by the United States Department of Agriculture, National Institute of Food and Agriculture (USDA-NIFA), through the Capacity Building Grants for Non-Land-Grant Colleges of Agriculture (Grant No. 2023-70001-40997).

## References

- [1] Tom Robert Angel and Sophie Anne Mahendran. Comparison of manual and automated body condition scoring of commercial dairy cattle. *Veterinary Record*, 195(7):no–no, 2024. 3
- [2] Hao Chen and Youfu Li. Progressively complementarity-aware fusion network for rgb-d salient object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3051–3060, 2018. 3
- [3] Junda Cheng, Wei Yin, Kaixuan Wang, Xiaozhi Chen, Shijie Wang, and Xin Yang. Adaptive fusion of single-view and multi-view depth for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10138–10147, 2024. 3
- [4] Georgios Christodouloupoulos. Subacute ruminal acidosis in cattle: A critical review of clinical management. *Veterinary Research Communications*, 49(5):273, 2025. 1
- [5] MG Embaby, M Günal, and A AbuGhazaleh. Effect of rumen culture adaptation to *origanum vulgare* l. essential oil on rumen methane and fermentation. *American J. Anim. Vet. Sci*, 14:190–196, 2019. 5
- [6] Mohamed G Embaby, Toqi Tahamid Sarker, Amer AbuGhazaleh, and Khaled R Ahmed. Optical gas imaging and deep learning for quantifying enteric methane emissions from rumen fermentation in vitro. *IET Image Processing*, 19(1):e13327, 2025. 2
- [7] Jörg MD Enemark. The monitoring, prevention and treatment of sub-acute ruminal acidosis (sara): A review. *The veterinary journal*, 176(1):32–43, 2008. 1
- [8] Jun Fu, Jing Liu, Haijie Tian, Yong Li, Yongjun Bao, Zhiwei Fang, and Hanqing Lu. Dual attention network for scene segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3146–3154, 2019. 4, 5, 6
- [9] Meili Fu, Huanliang Sun, Zhihan Chen, and Lulin Wei. Cafnet: Cross-modal adaptive fusion network with attention and gated weighting for rgb-t semantic segmentation. *IEEE Access*, 2025. 2
- [10] Vesna Gantner, Franjo Poljak, Boris Ljubojević, Zvonimir Steiner, and Ranko Gantner. Prevalence, persistence, and costs of acidosis in holstein cows across breeding regions. *AgroReS*, 14:158–165, 2025. 1
- [11] Yuanduo Hong, Huihui Pan, Weichao Sun, and Yisong Jia. Deep dual-resolution networks for real-time and accurate semantic segmentation of road scenes. *arXiv preprint arXiv:2101.06085*, 2021. 4, 5, 6
- [12] Tianqing Hu, Xiaofei Nan, Xiabing Zhou, Yu Shen, and Qinglei Zhou. A dual-stream feature decomposition network with weight transformation for multi-modality image fusion. *Scientific Reports*, 15(1):7467, 2025. 3
- [13] Taminul Islam, Toqi Tahamid Sarker, Khaled R Ahmed, Cristiana Bernardi Rankrape, and Karla Gage. Weedswin hierarchical vision transformer with sam-2 for multi-stage weed detection and classification. *Scientific Reports*, 15(1):23274, 2025. 3
- [14] Taminul Islam, Toqi Tahamid Sarker, Mohamed G Embaby, Khaled R Ahmed, and Amer AbuGhazaleh. Carbonext and carboformer: Dual semantic segmentation architectures for detecting and quantifying carbon dioxide emissions using optical gas imaging. *arXiv preprint arXiv:2506.05360*, 2025. 2, 4, 5, 6, 7
- [15] Peter H Janssen. Influence of hydrogen on rumen methane formation and fermentation balances through microbial growth kinetics and fermentation thermodynamics. *Animal Feed Science and Technology*, 160(1-2):1–22, 2010. 1
- [16] Pengcheng Jiao, Xinghong Ye, Chenjie Zhang, Wentao Li, and Haipeng Wang. Vision-based real-time marine and offshore structural health monitoring system using underwater robots. *Computer-Aided Civil and Infrastructure Engineering*, 39(2):281–299, 2024. 3
- [17] Harinder Kaur, Navjot Kaur, and Nirvair Neeru. Evolution of multiorgan segmentation techniques from traditional to deep learning in abdominal ct images—a systematic review. *Displays*, 73:102223, 2022. 6
- [18] Sangwon Kim, Dasom Ahn, and Byoung Chul Ko. Cross-modal learning with 3d deformable attention for action recognition. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10265–10275, 2023. 3
- [19] Jiajia Li, Pingping Zhang, Xia Yang, Lei Zhu, Teng Wang, Ping Zhang, Ruhan Liu, Bin Sheng, and Kaixuan Wang. Ssm-net: Semi-supervised multi-task network for joint lesion segmentation and classification from pancreatic eus images. *Artificial Intelligence in Medicine*, 154:102919, 2024. 3
- [20] Qian Li, Yongsheng Si, Mengyuan Chu, Ning Liu, Xi Kang, and Gang Liu. A novel lameness detection method for dairy cows based on temporal gait and spatial post features. *Computers and Electronics in Agriculture*, 224:109172, 2024. 3
- [21] Wenzhuo Liu, Wenshuo Wang, Yicheng Qiao, Qiannan Guo, Jiayin Zhu, Pengfei Li, Zilong Chen, Huiming Yang, Zhiwei Li, Lening Wang, et al. Mmtl-uniad: A unified framework for multimodal and multi-task learning in assistive driving perception. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 6864–6874, 2025. 3
- [22] Hongfang Lu, Dongmin Xi, Yaqin Xiang, Zhenhao Su, and Y Frank Cheng. Vehicle–canine collaboration for urban pipeline methane leak detection. *Nature Cities*, 2(4):336–343, 2025. 2
- [23] Jinju Mao and Lizhi Wang. Rumen acidosis in ruminants: a review of the effects of high-concentrate diets and the potential modulatory role of rumen foam. *Frontiers in Veterinary Science*, 12:1595615, 2025. 1
- [24] Sachin Mehta, Mohammad Rastegari, Linda Shapiro, and Hannaneh Hajishirzi. Espnetv2: A light-weight, power efficient, and general purpose convolutional neural network. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9190–9200, 2019. 4, 5, 6
- [25] Habiba N Ngugi, Absalom E Ezugwu, Andronicus A Akinyelu, and Laith Abualigah. Revolutionizing crop disease detection with computational deep learning: a comprehensive review. *Environmental Monitoring and Assessment*, 196(3):302, 2024. 3
- [26] FN Owens, DS Secrist, WJ Hill, and DR Gill. Acidosis in

- cattle: a review. *Journal of animal science*, 76(1):275–286, 1998. 1
- [27] Adam Paszke, Abhishek Chaurasia, Sangpil Kim, and Eugenio Culurciello. Enet: A deep neural network architecture for real-time semantic segmentation. *arXiv preprint arXiv:1606.02147*, 2016. 4, 5, 6
- [28] JC Plaizier, DO Krause, GN Gozho, and BW McBride. Subacute ruminal acidosis in dairy cows: The physiological causes, incidence and consequences. *The Veterinary Journal*, 176(1):21–31, 2008. 1, 6
- [29] Rudra PK Poudel, Stephan Liwicki, and Roberto Cipolla. Fast-scnn: Fast semantic segmentation network. *arXiv preprint arXiv:1902.04502*, 2019. 4
- [30] Bintang Pradana Erlangga Putra, Heri Prasetyo, and Esti Suryani. Residual transformer fusion network for salt and pepper image denoising. *arXiv preprint arXiv:2502.09000*, 2025. 4, 5, 6
- [31] Shambhavi Balamuthu Sampath, Sami Sawani, Moritz Thoma, Lukas Frickenstein, Pierpaolo Mori, Nael Fafous, Manoj Rohit Vemparala, Alexander Frickenstein, Ulf Schlichtmann, Claudio Passerone, et al. La-mtl: Latency-aware automated multi-task learning. In *2025 62nd ACM/IEEE Design Automation Conference (DAC)*, pages 1–7. IEEE, 2025. 3
- [32] Toqi Tahamid Sarker, Mohamed G Embaby, Khaled R Ahmed, and Amer AbuGhazaleh. Gasformer: A transformer-based architecture for segmenting methane emissions from livestock in optical gas imaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5489–5497, 2024. 2, 4, 5, 6, 7
- [33] Toqi Tahamid Sarker, Khaled R Ahmed, Taminul Islam, Cristiana Bernardi Rankrape, and Karla Gage. Weedsense: Multi-task learning for weed segmentation, height estimation, and growth stage classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7180–7190, 2025. 3
- [34] Zohreh Sohrabi and Jamshid Maleki. Fusing satellite imagery and ground-based observations for pm2. 5 air pollution modeling in iran using a deep learning approach. *Scientific Reports*, 15(1):21449, 2025. 2
- [35] Toqi Tahamid Sarker, Mohamed Embaby, Taminul Islam, Amer AbuGhazaleh, and Khaled R Ahmed. Gastwinformer: A hybrid vision transformer for livestock methane emission segmentation and dietary classification in optical gas imaging. *arXiv e-prints*, pages arXiv–2508, 2025. 2, 4, 5, 6, 7
- [36] Song Tang, Chuang Li, Pu Zhang, and RongNian Tang. Swinlstm: Improving spatiotemporal prediction accuracy using swin transformer and lstm. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 13470–13479, 2023. 3
- [37] Simon Vandenhende, Stamatios Georgoulis, and Luc Van Gool. Mti-net: Multi-scale task interaction networks for multi-task learning. In *European conference on computer vision*, pages 527–543. Springer, 2020. 4, 5, 6
- [38] Nikolaos Voulgarakis, Dimitrios Gougoulis, Dimitra Psalla, Georgios Papakonstantinou, Panagiotis-Dimitrios Katsoulos, Konstantinos Katsoulis, Mariana Angelidou-Tsifida, Labrini Athanasiou, Vasilios Papatsiros, and Georgios Christodoulopoulos. Can computerized rumen mucosal colorimetry serve as an effective field test for managing subacute ruminal acidosis in feedlot cattle? *Veterinary Research Communications*, 48(1):475–484, 2024. 1
- [39] Nikolaos Voulgarakis, Dimitrios A Gougoulis, Dimitra Psalla, Georgios I Papakonstantinou, Konstantinos Katsoulis, Mariana Angelidou-Tsifida, Labrini V Athanasiou, Vasileios G Papatsiros, and Georgios Christodoulopoulos. Subacute rumen acidosis in greek dairy sheep: Prevalence, impact and colorimetry management. *Animals*, 14(14):2061, 2024. 1
- [40] Guanbo Wang, Haiyan Li, Shuhua Ye, Hongzhi Zhao, Hongwei Ding, and Shidong Xie. Rfwnet: A multiscale remote sensing forest wildfire detection network with digital twinning, adaptive spatial aggregation, and dynamic sparse features. *IEEE Transactions on Geoscience and Remote Sensing*, 62:1–23, 2024. 2
- [41] Chuheng Wei, Ziyang Zhang, Haishan Liu, Guoyuan Wu, and Matthew Barth. Hierarchical multi-modal fusion for roadside vru detection: Method complementarity under sparse label constraints. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5343–5350, 2025. 3
- [42] Shuo Xing, Chengyuan Qian, Yuping Wang, Hongyuan Hua, Kexin Tian, Yang Zhou, and Zhengzhong Tu. Openemma: Open-source multimodal model for end-to-end autonomous driving. In *Proceedings of the Winter Conference on Applications of Computer Vision*, pages 1001–1009, 2025. 3
- [43] Wei Xu, Chunsheng Shi, Sifan Tu, Xin Zhou, Dingkan Liang, and Xiang Bai. A unified framework for 3d scene understanding. *Advances in Neural Information Processing Systems*, 37:59468–59490, 2024. 3
- [44] Changqian Yu, Changxin Gao, Jingbo Wang, Gang Yu, Chunhua Shen, and Nong Sang. Bisenet v2: Bilateral network with guided aggregation for real-time semantic segmentation. *International journal of computer vision*, 129(11):3051–3068, 2021. 4, 5, 6, 7
- [45] Jiaming Zhang, Huayao Liu, Kailun Yang, Xinxin Hu, Ruiping Liu, and Rainer Stiefelham. Cmx: Cross-modal fusion for rgb-x semantic segmentation with transformers. *IEEE Transactions on intelligent transportation systems*, 24(12):14679–14694, 2023. 4, 5, 6
- [46] Yan Zhang, Xinze Yang, Yufei Liu, Junyu Zhou, Yihong Huang, Jiapeng Li, Longxiang Zhang, and Qin Ma. A time-series neural network for pig feeding behavior recognition and dangerous detection from videos. *Computers and Electronics in Agriculture*, 218:108710, 2024. 3