

Video Understanding with Attention Encoder and Multimodal Large Language Model

Anni Zheng

A thesis submitted to the Auckland University of Technology
in partial fulfillment of the requirements for the degree of
Master of Computer and Information Sciences (MCIS)

2025

School of Engineering, Computer & Mathematical Sciences

Abstract

The challenge of achieving robust video understanding has become increasingly significant with the emergence of Multimodal Large Language Models (MLLMs). While MLLMs have demonstrated significant promise, effectively capturing and reasoning about complex temporal dynamics and object-level interactions in videos remains an active area of research. This project introduces a novel framework designed to enhance video understanding capabilities. We propose a new model architecture featuring a Temporal Context Gated Attention (TCGA) encoder layer, combined with a fine-tuned MLLM, demonstrates improved performance in video event retrieval and understanding tasks. Furthermore, we present the design and implementation of a real-time system application built upon our proposed model. This work aims to contribute a specialized video processing module and system design insights, offering a valuable step towards more sophisticated and applicable video understanding within MLLMs. We hope our findings provide a foundation for future research in temporal-aware multimodal learning.

Keywords: Multimodal LLM, Attention, Video Analytics, Video Classification, Event Retrieval

Contents

| | |
|---|-----------|
| Abstract | 2 |
| Attestation of Authorship | 7 |
| Publications | 8 |
| Acknowledgements | 9 |
| Intellectual Property Rights | 10 |
| 1 Introduction | 11 |
| 1.1 Background and Motivation | 11 |
| 1.2 Research Questions | 15 |
| 1.3 Contributions | 15 |
| 2 Literature Review | 17 |
| 2.1 Introduction | 17 |
| 2.2 Pre-trained Vision-Language Models | 18 |
| 2.3 Large Language Models | 21 |
| 2.4 Multimodal Large Language Models | 23 |
| 2.5 Attention Gate and Multimodal Applications | 29 |
| 2.5.1 Attention Gating | 31 |
| 2.5.2 Attention Gate Applications on Multimodal Tasks | 34 |
| 2.5.3 Intelligent Surveillance Video Analytics | 36 |
| 3 Methodology | 43 |
| 3.1 Introduction | 43 |
| 3.1.1 Model Framework | 45 |
| 3.2 Temporal Context Gated Attention | 52 |
| 3.2.1 Implementation Details | 60 |
| 3.2.2 Application Across Different Domain | 65 |
| 3.2.3 Computational Complexity Analysis | 66 |
| 3.3 Video Event Retrieval | 67 |
| 3.3.1 Fine-Tuning MLLM | 67 |
| 3.3.2 MLLM Video Retrieval Pipeline | 71 |

| | | |
|----------|---|------------|
| 3.3.3 | Application System Design | 73 |
| 4 | Experimental Results | 81 |
| 4.1 | Introduction | 81 |
| 4.2 | Experimental Setup and Datasets | 81 |
| 4.3 | Ablations Study | 83 |
| 4.4 | Video Frame Classification | 89 |
| 4.5 | Video Search with Description | 92 |
| 5 | Discussion | 96 |
| 5.1 | Evaluation Metrics | 96 |
| 5.2 | Benchmark Model Comparison | 99 |
| 5.3 | Ablation Study Results | 102 |
| 6 | Conclusion and Future Work | 105 |
| | References | 110 |

List of Tables

| | | |
|-----|---|-----|
| 4.1 | Architectural Ablation Study for TCGA Module Components. | 84 |
| 4.2 | Ablation Study on CSA Dataset Performance Evaluation | 85 |
| 4.3 | Ablation Study on the Number of Sampled Frames (T). | 87 |
| 4.4 | Video Frame Classification Accuracy (%) on Different Dataset. | 90 |
| 4.5 | Crime Video Search and Temporal Localization Performance on CrimeVidSearch-TSL Dataset. | 94 |
| 5.1 | Detailed Performance Comparison with Benchmark Models Across Use Case Datasets. | 100 |
| 5.2 | Architectural Ablation Study Results on CCTV Event and MRI Risk Datasets. | 102 |

List of Figures

| | | |
|-----|--|----|
| 2.1 | Proposed Model Achitecture | 36 |
| 3.1 | High-level architecture design of the video classification framework. | 46 |
| 3.2 | Fist layer frame encoder with multi-resolution feature fusion | 48 |
| 3.3 | Frame encoder backbone model | 49 |
| 3.4 | An overview and example of the proposed Multimodal Large Language Model (MLLM) video retrieval pipeline. | 74 |
| 3.5 | Real-time video analytics web application system design | 75 |
| 3.6 | Real-time video streaming interface | 78 |
| 3.7 | User query interface | 79 |
| 3.8 | Video retrieval result | 80 |

Attestation of Authorship

I hereby declare that this submission is my own work and that, to the best of my knowledge and belief, it contains no material previously published or written by another person nor material which to a substantial extent has been accepted for the qualification of any other degree or diploma of a university or other institution of higher learning.

Signature of student

Acknowledgements

I would like to extend my sincere gratitude to the individuals who have played a crucial role in the completion of this thesis. Their guidance, support, and encouragement have been invaluable. First and foremost, I wish to express my profound appreciation to my supervisor, Yan. Your insightful feedback, unwavering support, and expert guidance throughout this research process have been instrumental. I owe a special and heartfelt thank you to my mother, Gui. Your unconditional love, endless encouragement, and unwavering belief in me have been my constant source of strength and motivation. Through the late nights and challenging moments, your listening ear kept me going. This achievement would not have been possible without your enduring support and sacrifices.

Intellectual Property Rights

Theses, dissertations and research projects are protected by the Copyright Act 1994 (New Zealand). This thesis, dissertation or research projects may be consulted by you, provided you comply with the provisions of the Act and the following conditions of use:

- Any use you make of these documents or images must be for research or private study purposes only, and you may not make them available to any other person.
- Authors control the copyright of their thesis, dissertation or research project. You will recognise the author's right to be identified as the author of the thesis, dissertation or research project, and due acknowledgment will be made to the author where appropriate.
- You will obtain the author's permission before publishing any material from the thesis, dissertation or research project.
- The ownership of any intellectual property rights which may be described in this thesis is vested in the Auckland University of Technology, subject to any prior agreement to the contrary, and may not be made available for use by third parties without the written permission of the University, which will prescribe the terms and conditions of any such agreement.

Copyright ©2025. Anni Zheng

Chapter 1

Introduction

1.1 Background and Motivation

The remarkable advancements in Large Language Models (LLMs) (Brown et al., 2020; Zhang et al., 2022; Touvron et al., 2023b, 2023a) have revolutionized natural language processing capabilities, demonstrating unprecedented performance in text comprehension, generation, and reasoning tasks. These models, trained on vast corpora of text data, have established new benchmarks across diverse linguistic challenges, from complex question answering to nuanced creative writing. Building upon this foundation, the field has naturally evolved toward Multimodal Large Language Models (MLLMs) (OpenAI, 2023b; Team et al., 2023; Liu et al., 2024, 2023), which extend beyond textual modalities to incorporate visual understanding.

Throughout the development of LLM, ChatGPT (OpenAI, 2023a) catalyzed a

paradigm shift in artificial intelligence research and applications. The emergence of these powerful models has not only transformed how we approach natural language processing but has also paved the way for cross-modal integration. Furthermore, cutting-edge Multimodal Large Language Models (MLLMs)(OpenAI, 2023b; Team et al., 2023; Alayrac et al., 2022) represents a significant paradigm shift in artificial intelligence research, extending the fundamental capabilities of traditional Large Language Models (LLMs) to encompass sophisticated visual comprehension functionalities. This evolutionary trajectory has culminated in the development of advanced systems demonstrating unprecedented proficiency in the seamless integration and concurrent processing of both visual and textual modalities. The architectural and functional advancements exhibited by these models have subsequently positioned MLLMs as a critical nexus for interdisciplinary scientific inquiry, catalyzing convergent research initiatives across previously disparate domains and generating substantial discourse within the broader academic community (Wu et al., 2023; Yang et al., 2023; Wu et al., 2023; Wake et al., 2023).

The architecture of existing MLLMs can be delineated into three fundamental components: the pre-trained vision encoder, CLIP’s ViT-L (Radford et al., 2021) or EVA-CLIP’s ViT-G (Sun et al., 2023), which extracts meaningful representations from visual inputs; the pre-trained LLM, OPT (Zhang et al., 2022), Llama (Touvron et al., 2023a), Vicuna (Chiang et al., 2023), which processes and generates text based on contextual understanding; and the connector Q-former (Alayrac et al., 2022; Li et al., 2023a) or linear projection (Liu et al., 2024, 2023) trained from scratch to bridge the semantic gap between vision and language models. This tripartite structure has become the standard paradigm for contemporary MLLM architectures, with each component playing a crucial role in facilitating cross-modal understanding.

From the findings in computer vision literature, where classic models Densenet (Huang et al., 2017), FPN (Lin et al., 2017), ResNet (He et al., 2016a) have demonstrated the efficacy of utilizing multi-layer features to enhance visual representations for downstream tasks.

Recent efforts have sought to explicitly enhance visual information by increasing image resolution (Li et al., 2024; Bai et al., 2023b; Liu et al., 2024; Li et al., 2023; Wang et al., 2023; McKinzie et al., 2024) or introducing additional visual encoders (Jiang et al., 2023; Tong et al., 2024). However, these methods often depend solely on high-level visual features for the final embeddings in MLLMs and frequently introduce significant computational overhead or architectural complexity. In contrast, our approach proposes leveraging the inherent richness of representations across different layers of the existing visual encoder offering substantial improvements in visual understanding.

In MLLMs, vision encoders are often frozen to avoid the high costs of end-to-end training. We leverage this by using offline features from various layers of the frozen encoder, effectively enhancing visual information at no extra cost in parameters or inference computation. This method also complements techniques that directly boost visual signals, such as increasing image resolution (Li et al., 2024; Bai et al., 2023b; Liu et al., 2024; Li et al., 2023; Wang et al., 2023; McKinzie et al., 2024) or introducing additional visual encoders (Jiang et al., 2023; Tong et al., 2024; Li et al., 2024). The methodology distinguishes itself through remarkable simplicity and computational efficiency. Unlike complex architectural modifications, this approach delivers performance gains without elaborate implementation requirements. Its most compelling attribute lies in design agnosticism, the technique integrates seamlessly across diverse MLLM architectures with minimal adaptation. By interfacing with standard components present in varied MLLM frameworks, the method maintains broad

applicability across research and production environments. The technique requires no specialized structures, instead operating within established architectural boundaries while still delivering tangible improvements to cross-modal reasoning capabilities.

1.2 Research Questions

In light of these observations, we propose the Temporal Context Gated Attention, a novel video frame encoder for downstream video understanding tasks of video classification and video event retrieval.

In this research we focused on following research questions:

- Whether the accuracy of video classification can be improved by novel attention mechanics (TCGA) with better object detail encoding layer?
- Compared with the existing model, how does the novel model architecture with fine-tuned MLLM improve performance on the task of video event retrieval?

1.3 Contributions

We summarize our contributions as follows:

- We propose a novel model framework for video classification tasks that leverages the attention layer to enhance temporal context learning in video events.
- We improved model performance by fine-tuning the QWen MLLM, which was integrated with our innovative attention module, which significantly enhances the model's ability to capture temporal dynamics in video content.
- We design and implement a real-time system application based on our proposed model, demonstrating its practical utility and efficiency in deployment scenarios.

Through experiments evaluation, we demonstrate that our approach consistently

improves performance across diverse domains, including billiard game videos and surveillance recording videos, while maintaining computational efficiency suitable for real-time applications.

Chapter 2

Literature Review

2.1 Introduction

This section presents a comprehensive overview of the relevant literature that forms the foundation for our research. We explore the evolution of vision models, language models, and their integration into multimodal systems, with particular emphasis on video understanding frameworks. This review contextualizes our contributions within the broader research landscape and highlights the gaps our work aims to address.

2.2 Pre-trained Vision-Language Models

The advent of pre-trained Vision Transformers (ViT) (Dosovitskiy et al., 2020) has significantly propelled the advancement of computer vision, fundamentally transforming how visual information is processed and understood by deep learning models (Yan, 2023, 2019a; Zheng & Yan, 2025, 2024). The introduction of ViT by Dosovitskiy et al. represented a pivotal shift away from the convolutional neural network (CNN) paradigm that had dominated computer vision for nearly a decade. By adapting the transformer architecture—originally designed for natural language processing, to visual data, ViT demonstrated that self-attention mechanisms could effectively capture long-range dependencies in images without the inductive biases inherent in CNNs (Liang & Yan, 2022). The evolution of vision transformers has been marked by several key developments that have progressively enhanced their capabilities. Early implementations faced challenges related to data efficiency and computational requirements, but subsequent iterations like DeiT (Touvron et al., 2021) introduced distillation techniques that improved training efficiency. Swin Transformer (Liu et al., 2021) further refined the architecture by introducing hierarchical representation with shifted windows, effectively addressing the quadratic complexity issues while maintaining the benefits of self-attention mechanisms.

Furthermore, pre-training ViT models on web-scale image-text pairs, CLIP (Radford et al., 2021) and its subsequent iterations (Sun et al., 2023; Zhai et al., 2023; Cherti et al., 2023; Yao et al., 2021), where vision and text encoders are simultaneously trained. This contrastive learning approach has enabled models to develop robust visual representations that generalize remarkably well across unseen domains and tasks.

The CLIP architecture represents a significant advancement in multimodal learning by establishing robust connections between visual and linguistic data. Through its implementation of contrastive learning across an extensive dataset of millions image-text pairs sourced from internet collections, CLIP has achieved remarkable capabilities in zero-shot transfer learning (Radford et al.). This framework operates by projecting both images and text into a unified embedding space, where semantic relationships between modalities are preserved and can be leveraged for cross-domain inference tasks. The underlying methodology enables CLIP to perform effectively on downstream applications without requiring traditional task-specific fine-tuning processes that have historically characterized computer vision systems. By training simultaneously on visual and textual information, CLIP creates representations that capture deep semantic alignments between what objects look like and how they are described in natural language. The shared embedding space facilitates direct comparison between previously unseen images and arbitrary text descriptions, a fundamental capability that underpins CLIP's transferability across diverse vision tasks. This paradigm shift in vision-language pre-training demonstrates how contrastive objectives can yield models with broad generalization abilities that transcend the limitations of conventional supervised learning approaches.

Building upon CLIP's foundation, subsequent works have further refined the contrastive learning approach. EVA (Sun et al., 2023) extended the framework by incorporating masked image modeling alongside contrastive learning, resulting in more robust representations. OpenCLIP (Cherti et al., 2023) democratized access to CLIP-like models by providing open-source implementations trained on publicly available datasets, facilitating broader research participation. FILIP (Yao et al., 2021) introduced fine-grained token-level interactions between image and text representations, enabling more nuanced alignment between modalities.

Since their introduction, CLIP-like models have served as effective initializations and have been incorporated into various vision-language cross-modal models, for example video-text alignment (Fang et al., 2023; Wu et al., 2023, 2023; ?, ?) and large vision-language models (Li et al., 2023a; Liu et al., 2024; Zhu et al., 2023). The transferability of these pre-trained representations has accelerated progress across numerous domains, from image generation to visual reasoning and multimodal understanding.

Recently, SigLIP (Zhai et al., 2023) introduced pairwise sigmoid loss during training, representing a departure from the traditional softmax-based contrastive loss used in CLIP. This architectural innovation has enabled the visual encoder to demonstrate more advanced visual perception capabilities, particularly in fine-grained recognition tasks. The sigmoid loss function provides a more stable training dynamic and better handles hard negative examples, resulting in more discriminative visual representations.

Beyond contrastive learning approaches, alternative self-supervised learning frameworks have emerged that focus exclusively on visual data without requiring paired textual information. Models like DINO and MoCo utilize self-distillation and momentum contrast techniques, respectively, to learn powerful visual representations. These approaches have demonstrated complementary strengths to contrastive vision-language models, often excelling at capturing local structural information and semantic consistency.

Recent studies have also explored the integration of multiple pre-training paradigms, combining the strengths of different approaches. For instance, combining features from DINO and CLIP has shown promising results, as each captures distinct and complementary aspects of visual information. Such hybrid approaches highlight the importance of considering diverse representational perspectives in visual understanding.

To validate the compatibility and versatility of our proposed model, this paper conducted extensive experiments on different visual encoders, including those of CLIP (Radford et al., 2021) and SigLIP (Zhai et al., 2023). This experimental design allows us to evaluate the generalizability of our approach across different visual representation paradigms and to identify potential synergies between our method and specific encoder architectures.

2.3 Large Language Models

The exceptional text understanding and generation capabilities demonstrated by autoregressive Large Language Models (LLMs) (Brown et al., 2020; Raffel et al., 2020; Yang et al., 2019) have garnered significant attention in recent years, fundamentally reshaping research priorities in natural language processing and artificial intelligence more broadly. The emergence of models like GPT-3 (Brown et al., 2020) with 175 billion parameters marked a turning point, demonstrating that scaling up model size and training data could lead to emergent capabilities not explicitly engineered into the architecture.

Subsequently, a plethora of LLMs (Touvron et al., 2023b, 2023a; Zhang et al., 2022; Chowdhery et al., 2023) have emerged, with notable open-source efforts like LLaMA (Touvron et al., 2023b) greatly propelling community contributions to LLMs research. Meta AI's release of LLaMA (Touvron et al., 2023b) and its successor LLaMA-2 (Touvron et al., 2023a) has been particularly influential, providing researchers with access to state-of-the-art foundation models ranging from 7 billion to 70 billion parameters. These open-source initiatives have democratized access to cutting-edge

language models, enabling a diverse ecosystem of adaptations and applications.

The architectural evolution of LLMs has been characterized by refinements to the transformer architecture. Innovations like rotary positional embeddings, flash attention, and mixture-of-experts architectures have addressed key limitations related to context length, computational efficiency, and parameter utilization, respectively. These technical advancements have collectively enabled models to process longer sequences, train more efficiently, and achieve better performance with the same computational budget.

Through instruction fine-tuning techniques (Ouyang et al., 2022; Wei et al., 2021), these models showcase human-like language interaction abilities, further propelling advancements in natural language processing. The paradigm of instruction tuning, where models are explicitly trained to follow natural language instructions, has proven particularly effective in aligning model behavior with human expectations. This approach, pioneered by works like FLAN (Wei et al., 2021) and InstructGPT (Ouyang et al., 2022), has become standard practice in developing user-facing language models. Recent developments have seen LLMs scaled up or down to meet various application needs, reflecting a growing recognition that different use cases may require different model sizes. Lightweight LLMs (Javaheripi et al., 2023; Zhang et al., 2024; Bai et al., 2023a; Bellagente et al., 2024) have been developed to address computational constraints, facilitating edge deployment and real-time applications. Models like TinyLLaMA (Zhang et al., 2024) and Phi-2 (Javaheripi et al., 2023) have demonstrated impressive capabilities despite their relatively small parameter counts (1-2 billion parameters), challenging assumptions about the necessity of massive scale for useful language capabilities.

Conversely, in the pursuit of exploring the upper limits of LLMs, works such as (Jiang et al., 2024; Young et al., 2024; Touvron et al., 2023a; Bai et al., 2023a)

have expanded LLM parameters, continuously pushing the boundaries of language capabilities. Notably, Mixtral 8x7B (Jiang et al., 2024) has demonstrated that sparse mixture-of-experts architectures can achieve performance comparable to much larger dense models while reducing computational requirements during inference. Meanwhile, Yi (Young et al., 2024) and Qwen (Bai et al., 2023a) have shown that careful data curation and training methodologies can lead to models that outperform others with similar parameter counts, highlighting that scale is not the only determinant of model quality.

In this study, we made experiments to compare multiple LLMs ranging from 7B to 70B parameters.

2.4 Multimodal Large Language Models

After witnessing the success of LLMs in natural language processing, researchers have shifted their focus towards enabling LLMs to understand and reason about visual signals, giving rise to Multimodal Large Language Models (MLLMs). This emerging field represents a convergence of vision and language capabilities within unified architectural frameworks, enabling more holistic understanding of multimodal information.

To achieve visual-linguistic integration, several architectural approaches have been proposed. Early methods focused on creating specialized interfaces between pre-trained vision and language models. Prior research has proposed compressing visual embeddings using Q-former (Li et al., 2023a) into query embeddings, followed by transforming them into text embeddings through linear projection, or directly employing

MLP projection (Liu et al., 2024) to connect the visual encoder with LLM. These approaches address the fundamental challenge of bridging the representational gap between visual and linguistic domains without requiring extensive retraining of the underlying models.

The Q-former approach, pioneered by BLIP-2 (Li et al., 2023a), utilizes a transformer-based query module that acts as an intermediary between the visual encoder and language model. This module learns to extract relevant visual information based on a set of learnable query tokens, effectively distilling the high-dimensional visual features into a more compact representation that can be directly consumed by the language model. This method has demonstrated strong performance while maintaining computational efficiency through reduced token count.

In contrast, the MLP projection approach employed by models like LLaVA (Liu et al., 2024) utilizes a simpler feed-forward network to directly transform visual features into the language model's embedding space. While conceptually simpler, this approach has proven surprisingly effective, particularly when combined with high-quality instruction tuning data. The relative simplicity of this method offers advantages in terms of training stability and computational efficiency, making it a popular choice for many recent MLLMs.

Furthermore, following the instruction tuning paradigm (Ouyang et al., 2022; Wei et al., 2021) that proved highly effective for aligning language models with human intent, pioneering works (Zhu et al., 2023; Liu et al., 2024; Dai et al., 2024) significantly boost the development of MLLMs through visual instruction tuning. This approach involves fine-tuning models on datasets consisting of image-text pairs accompanied by instructions that specify desired responses or reasoning patterns. Visual instruction

tuning has proven crucial for developing models that can follow complex directives while leveraging visual information.

The quality and scale of training data have emerged as critical factors in MLLM development. Subsequently, by introducing larger-scale and higher-quality datasets, efforts such as (Li et al., 2024; Liu et al., 2024; Chen et al., 2023; Bai et al., 2023b) have notably enhanced the visual understanding and reasoning capabilities of MLLMs. The ShareGPT4V dataset (Chen et al., 2023), for instance, leverages GPT-4V’s outputs as high-quality training data, creating a virtuous cycle where stronger models help train the next generation of systems. Similarly, MiniGemini (Li et al., 2024) introduced carefully curated datasets focused on complex visual reasoning tasks, pushing the boundaries of what these models can accomplish.

Beyond data-centric improvements, architectural innovations have continued to enhance MLLM capabilities. Some approaches focus on enriching the visual signal provided to the language model. Additionally, there are works that introduce additional visual encoders (Jiang et al., 2023; Tong et al., 2024) or utilize higher-resolution images (Liu et al., 2024; Li et al., 2024; Bai et al., 2023b) to provide richer visual signal sources. For example, From CLIP2 DINO (Jiang et al., 2023) combines features from different pre-trained visual models to capture complementary aspects of visual information, while LaVIN introduced efficient parameter-tuning techniques that enable better integration of visual information without extensive retraining.

The transition from static image understanding to dynamic video comprehension represents a natural evolution for MLLMs. Meanwhile (Ma & Yan, 2024), a plethora of studies (Maaz et al., 2023; Zhang et al., 2023; Lin et al., 2023) directly extend these above image-based methods to video conversational models by leveraging video

instruction tuning datasets. Video-LLaMA (Zhang et al., 2023) adapts the Q-former architecture to handle temporal information, while Video-ChatGPT (Maaz et al., 2023) employs strategic temporal pooling to manage the increased token count associated with video inputs. These approaches demonstrate that the principles developed for image understanding can be effectively adapted to temporal domains with appropriate modifications.

A persistent challenge in video understanding is efficiently managing the increased token count and computational demands associated with processing multiple frames. Various strategies have been proposed to address this challenge, including uniform frame sampling (Alayrac et al., 2022), hierarchical temporal encoding, and specialized attention mechanisms that operate across both spatial and temporal dimensions. These approaches aim to capture meaningful temporal patterns while maintaining computational tractability.

The integration of video understanding capabilities into multimodal language models presents unique challenges compared to static image understanding. The substantial increase in input token count due to multiple frames necessitates efficient strategies for temporal information aggregation. Early approaches like Flamingo (Alayrac et al., 2022) adopted simple uniform frame sampling (typically at 1 FPS) followed by temporal pooling to compress video information before feeding it to the language model. While computationally efficient, this approach risks losing fine-grained temporal details that may be crucial for certain tasks.

More sophisticated approaches have since emerged to better preserve temporal information while managing computational constraints. Video-LLaMA (Zhang et al., 2023) introduced a specialized video Q-former that processes sampled frames

while maintaining awareness of their temporal relationships. Video-ChatGPT (Maaz et al., 2023) employed a combination of spatial and temporal attention mechanisms to selectively focus on relevant spatiotemporal regions before projection into the language model’s embedding space. Valley (Luo et al., 2023) incorporated a temporal perception module that explicitly models motion patterns and temporal transitions, enhancing the model’s ability to reason about dynamic events.

The Temporal Channel Guided Attention (TCGA) module we propose in this work builds upon these foundations while introducing several key innovations. Unlike approaches that treat temporal modeling as a separate stage, TCGA integrates temporal reasoning directly into the multimodal fusion process. By explicitly guiding attention based on temporal channel information, our module helps the model focus on relevant temporal patterns without requiring extensive parameter additions or computational overhead. The real-time system application we’ve developed demonstrates the practical utility of our approach in deployment scenarios. Unlike many academic implementations that prioritize accuracy over efficiency, our system maintains a careful balance between performance and computational requirements, enabling responsive interaction in practical settings. This is achieved through a combination of efficient model design, optimized inference strategies, and carefully tuned preprocessing pipelines that minimize latency while preserving critical information.

In summary, our work contributes to the evolving landscape of video understanding by introducing a novel model framework that effectively leverages the Dense Connector architecture in conjunction with the specialized TCGA module. By fine-tuning the QWen language model with this integrated approach, we achieve significant performance improvements on video classification tasks while maintaining computational efficiency suitable for real-time applications. Our research addresses key challenges in

temporal modeling while providing a practical implementation pathway for deployment in real-world scenarios.

Despite the significant progress in MLLMs and video understanding discussed above, several key research gaps remain that our work aims to address. These gaps represent opportunities for meaningful contribution to the field and motivate the specific research directions we pursue in this thesis. First, while substantial attention has been devoted to enhancing language model capabilities and expanding training datasets, the visual encoding side of MLLMs has received comparatively less attention. Most approaches continue to rely on high-level features from the final layers of frozen visual encoders, effectively discarding potentially valuable information encoded in intermediate representations. This oversight is particularly surprising given that classic computer vision literature has consistently demonstrated the value of multi-level feature integration in tasks ranging from object detection to segmentation.

Second, the transition from image to video understanding in MLLMs has primarily focused on managing the increased computational burden rather than fundamentally rethinking how temporal information is encoded and utilized. Many approaches simply extend image-based methods with minimal adaptations, potentially missing opportunities to leverage the unique properties of temporal data. The development of specialized architectures that efficiently capture meaningful temporal patterns while integrating seamlessly with language models remains an underexplored area.

Third, while theoretical advances in model architecture and training methodologies continue at a rapid pace, there remains a significant gap between research prototypes and deployable systems. Many state-of-the-art models are prohibitively expensive to run in real-time scenarios, limiting their practical utility. The development of efficient,

deployment-ready solutions that maintain strong performance while operating under realistic computational constraints represents a crucial research direction. Fourth, the majority of existing work relies on extensive task-specific fine-tuning or specialized architectures for video understanding, limiting the flexibility and adaptability of resulting models. Approaches that can leverage the capabilities of existing models while extending them to new domains with minimal additional training offer significant practical advantages and warrant further investigation.

2.5 Attention Gate and Multimodal Applications

The incorporation of **gating mechanisms** represents a significant and now well-established paradigm within the design and architecture of neural networks. Foundational advancements in recurrent neural architectures, most notably Long Short-Term Memory (LSTM) units (Hochreiter & Schmidhuber, 1997) and Gated Recurrent Units (GRUs) (Dey & Salem, 2017), alongside innovative feedforward structures like Highway Networks (Srivastava et al., 2015), were instrumental in pioneering the application of gating. These early systems effectively demonstrated the utility of gates for meticulously regulating the transmission of information across successive time steps in recurrent models or through hierarchical layers in deeper networks. A primary motivation was the enhancement of gradient propagation, thereby mitigating issues such as vanishing or exploding gradients which historically plagued simpler recurrent structures.

This fundamental principle of controlled information flow via gating has not only persisted but has become increasingly integral in a wide array of **contemporary neural network designs**. Recent breakthroughs in the domain of sequence modeling, for

instance, frequently rely on gating. This includes sophisticated state-space models (SSMs) like Mamba (Gu & Dao, 2023; Dao & Gu, 2024) and various instantiations of attention mechanisms that form the backbone of Transformer models and their derivatives (Hua et al., 2022; Sun et al., 2023; Qin et al., 2024a; Yang et al., 2024; Lin et al., 2025). In these modern contexts, gating is commonly applied to dynamically modulate or refine the outputs generated by token-mixer components, allowing the network to selectively emphasize or attenuate different pieces of information. Despite this pervasive adoption and the consistent empirical successes reported across diverse applications, a truly comprehensive theoretical and empirical understanding of the nuanced roles and multifaceted impacts of these gating mechanisms, extending beyond their *prima facie* conceptualization, remains notably incomplete in the existing literature.

This current deficit in profound understanding creates considerable challenges when attempting to accurately ascertain the genuine, isolated contribution of gating elements to overall model performance. The difficulty is exacerbated because the effects of gating are often intricately confounded with a multitude of other architectural variables and design choices. To illustrate this point, consider the Switch Heads architecture (Csordas et al., 2024, 2024), which introduces a sigmoid gating function specifically for the purpose of selecting a subset of top- K attention head "experts."

In this simplified configuration, the gate's function ostensibly reduces to merely modulating the value output of that single head. This outcome strongly suggests that the gating mechanism itself confers significant intrinsic benefits that are separate and distinct from its more apparent role as a routing or selection mechanism.

Analogous complexities arise in other recently proposed architectures. For example, while Native Sparse Attention (NSA) (Yuan et al., 2025) demonstrates commendable

overall performance improvements, the presented analyses do not systematically disentangle the specific contributions of its integrated gating mechanism from the broader effects inherent to the sparse attention design itself. It remains unclear to what extent the observed gains are attributable to the gating versus the specific pattern of sparsity or other computational aspects of the NSA framework. These illustrative examples, among others, powerfully underscore the critical and pressing need for more rigorous experimental designs and analytical methodologies. Such approaches are essential to meticulously disentangle and quantify the precise effects attributable to gating, distinct from the influence of other concurrently operating architectural components, thereby fostering a more precise and actionable understanding of their value in neural systems.

2.5.1 Attention Gating

Gating mechanisms have become a foundational component in the architecture of numerous neural networks. Pioneering efforts, exemplified by Long Short-Term Memory networks (LSTMs) (Hochreiter & Schmidhuber, 1997) and Gated Recurrent Units (GRUs) (Dey & Salem, 2017), initially employed gates to orchestrate the flow of information across temporal sequences, thereby tackling the notorious issues of vanishing or exploding gradients through the selective retention or dismissal of data. This principle was subsequently extrapolated to feedforward architectures by Highway Networks (Srivastava et al., 2015), paving the way for the successful training of substantially deeper models. More recently, SwiGLU (Shazeer, 2020) incorporated gating into the feedforward network (FFN) layers of transformers, a development credited with bolstering their expressive capabilities and subsequently establishing these mechanisms as a standard feature in many prominent open-source Large Language Models

(LLMs) (Grattafiori et al., 2024; Yang et al., 2024).

The utility of gating is further highlighted by its integration into diverse state-space models (Gu & Dao, 2023; Dao & Gu, 2024) and innovations in Linear Attention, including FLASH (Hua et al., 2022), RetNet (Sun et al., 2023), Lightning Attention (Qin et al., 2024a, 2024b; Li et al., 2025), and Gated Delta Networks (Yang et al., 2024). These frameworks employ gating modules to meticulously manage the information processed by their token-mixer sub-layers. Notably, the Forgetting Transformer (Lin et al., 2025) positions gating mechanisms directly after the softmax attention output, reporting considerable performance gains as a result. While these collective works affirm the empirical benefits of gating, a deeper, more granular understanding of the precise operational dynamics of these gates and the underlying factors contributing to their success warrants further exploration. Such investigations could not only broaden the appreciation of gating's pivotal role outside the traditional realm of RNNs but also inspire architectural innovations that more astutely capitalize on gating's inherent advantages. To illustrate, while models such as Switch Heads (Csordas et al., 2024, 2024), NSA (Yuan et al., 2025), and MoSA (Piękos et al., 2025) leverage sigmoid-based gating (Csordas et al., 2023) primarily for selection, a dedicated examination aimed at isolating the distinct impact of the gating function could prove highly informative. Moreover, comparative analyses against baselines that incorporate comparable gating within conventional transformer architectures could offer a more discerning assessment of their proposed selection techniques' true added value.

The research most analogous to our own investigation is presented in Quantizable Transformers (Bondarenko et al., 2023). This study also ascertains that deploying gating mechanisms within softmax attention serves to ameliorate issues of pronounced attention concentration and the occurrence of outlier values in the hidden states of encoder

architectures like BERT and ViT. Whereas the aforementioned study predominantly employs gating to counteract outliers to facilitate model quantization, our work distinguishes itself by offering an in-depth analysis of multiple gating configurations. We elucidate their contributions to augmenting non-linearity and sparsity, and to promoting more stable training dynamics. Drawing from these observations, we then proceed to scale gated attention models, thereby substantiating the extensive applicability and profound impact of these mechanisms.

The phenomenon denoted as the 'attention sink,' wherein certain tokens become recipients of disproportionately large attention scores, was formally described by (Xiao et al., 2023). Analogously, within the context of vision transformers, (Darcet et al., 2023) observed that some ostensibly redundant tokens adopt a role akin to 'registers,' serving as accumulators for attention scores. Further extending this line of inquiry, (Sun et al., 2024) provided evidence that exceedingly high attention scores are frequently directed towards tokens that also exhibit massive activation values. Our own findings, however, introduce a critical distinction: even when gating applied at the value projection output successfully curtails massive activations, attention sinks nevertheless endure. This observation implies that such activations are not an indispensable precursor to the emergence of attention sinks. In a similar vein, (Gu et al., 2024) characterize attention sinks as essentially non-informative 'key biases' that accrue redundant attention scores. They contend that the softmax function's inherent normalization dependency is a primary driver of this tendency.

Various experimental interventions aimed at modifying softmax attention—such as the replacement of softmax with unnormalized sigmoid attention (Ramapuram et al., 2024; Gu et al., 2024), the incorporation of a softmax attention gate or clipping mechanism (Bondarenko et al., 2023), and adjustments to the softmax computation (Zuhri

et al., 2025) or its denominator (Miller, 2023)—have shown encouraging results in diminishing the prevalence of attention sinks. The study (Qiu et al., 2025) shows that the strategic application of sparse gating, positioned after Scaled Dot-Product Attention (SDPA), effectively eradicates attention sinks. This holds true for both substantial dense models (1B parameters) and large-scale Mixture-of-Experts (MoE) architectures (15B parameters), even when subjected to extensive training on 3.5T tokens. It unveils a consequential implication: the successful mitigation of attention sinks may unlock new potentials for extending the effective context length manageable by these models.

2.5.2 Attention Gate Applications on Multimodal Tasks

The pursuit of robust and accurate models for breast cancer risk prediction using multimodal data is substantially informed by a confluence of advancements across several key domains within machine learning. Central to this endeavor is the capacity to derive highly informative representations from complex medical imaging data, such as Magnetic Resonance Imaging (MRI) scans. In this context, self-supervised contrastive learning methodologies have gained prominence. A notable example is SimCLR (Chen et al., 2020), which has demonstrated considerable efficacy in learning potent image features without the prerequisite of extensive, pixel-level human annotations. The core principle of contrastive learning involves training a model to maximize concordance between different augmented perspectives of the same input image while simultaneously minimizing concordance with other, distinct images. This paradigm is especially valuable in the medical field, where the acquisition of large, meticulously labeled datasets can be both labor-intensive and costly. Such self-supervised frameworks frequently

leverage well-established deep convolutional neural network architectures, with ResNet (He et al., 2016b) being a prominent choice. ResNet's introduction of residual connections fundamentally enabled the training of significantly deeper networks, which serve as powerful backbones for feature extraction within these contrastive learning schemes.

Beyond the representation of individual images, clinical assessments often rely on multiple image instances for a single patient—for example, different MRI sequences, distinct views, or images taken over a period. Addressing this multi-instance nature is critical for a holistic understanding. Multi-Instance Contrastive Learning (MICLe) (Azizi et al., 2021) offers a sophisticated approach specifically designed for such scenarios. It operates by maximizing the mutual information between the collective set of images originating from the same patient. This strategy encourages the model to learn embeddings that are not only discriminative but also cohesive at the patient level, thereby facilitating the generation of comprehensive representations that capture the overall patient status rather than isolated features from individual images.

Furthermore, as breast cancer risk is influenced by a variety of factors that can be captured through disparate data modalities (e.g., imaging, clinical records, genetic information), the effective integration of these diverse information streams is a paramount challenge. The Multimodal Adaptation Gate (MAG) mechanism, as proposed by (Rahman et al., 2020), presents an adaptive and elegant solution to this fusion problem. Recognizing that the relative importance of each modality can vary significantly across different patient cases, MAG introduces learnable gating components. These gates dynamically modulate the influence of each modality's representation during the fusion process, allowing the model to learn attention weights that selectively emphasize the most salient and pertinent information from the combined multimodal input for the

specific prediction task at hand. Collectively, these established and innovative methodologies in representation learning, multi-instance learning, and adaptive multimodal fusion provide a robust theoretical and practical toolkit for advancing the development of next-generation predictive models in complex medical applications like breast cancer risk assessment.

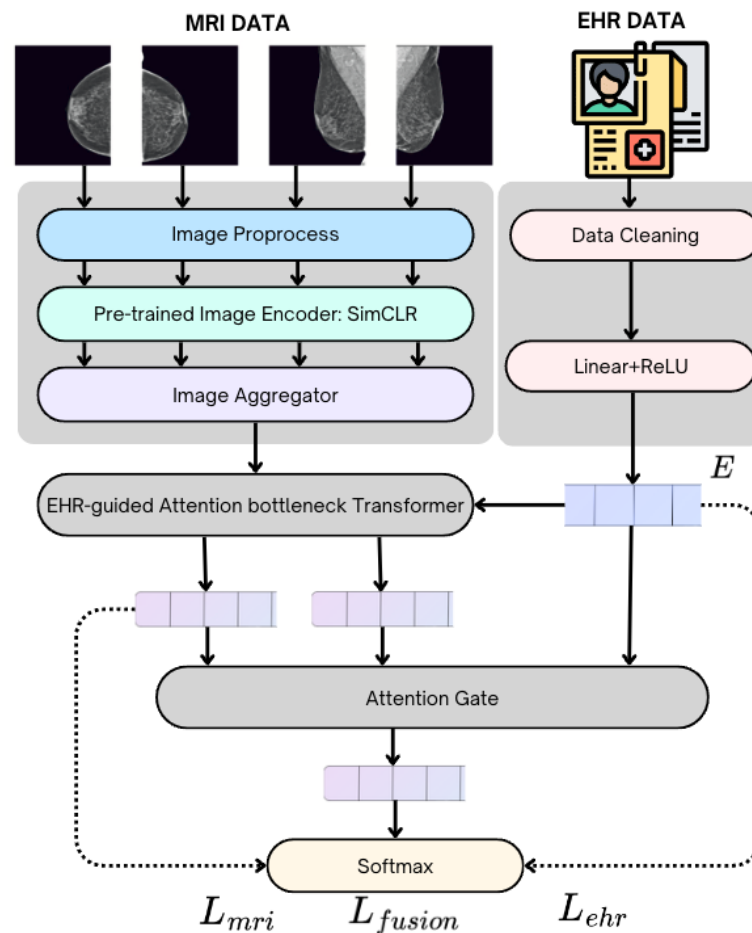


Figure 2.1: Proposed Model Architecture

2.5.3 Intelligent Surveillance Video Analytics

Intelligent Surveillance Video Analytics (ISVA) has become an indispensable field, addressing the escalating demand for automated security, real-time monitoring, and

comprehensive event understanding across diverse public and private sectors (Cao & Yan, 2023). The evolution of ISVA has been marked by significant advancements in sensor technology, computational power, and algorithmic sophistication. This section delves into the key research contributions that have charted the course of ISVA, examining its foundational principles, architectural paradigms, the pivotal role of event processing, the transformative impact of artificial intelligence, and the critical, ongoing challenges related to privacy and ethics.

Intelligent Surveillance Video Analytics (ISVA) has emerged as a critical field of research and application, driven by the increasing need for automated monitoring, security, and event understanding in various environments. This section reviews key literature that has shaped the development of ISVA, focusing on foundational concepts, system architectures, event processing, and the growing importance of privacy considerations.

The foundations of intelligent surveillance often involve adaptive mechanisms to handle dynamic environments and large volumes of data. Early research highlighted the importance of adaptive monitoring; for instance, (Wang et al., 2004) proposed schemes for adjusting video camera parameters based on feedback from video analysis, utilizing experiential sampling techniques to improve the quality of surveillance output. The concept of experiential sampling itself was further detailed by (Wang et al., 2003) as a methodology for real-time video surveillance, enabling dynamic modeling of attention to perform efficient monitoring and manage operator fatigue in multi-camera setups.

A significant thrust in ISVA has been the development of event-centric systems. The theoretical underpinnings of "Visual Event Computing" were explored by (Yan, 2019b), laying groundwork for understanding and processing events captured in video streams. Practical implementations of event-driven surveillance systems were presented

by (Kieran & Yan, 2010), who developed a framework focused on enabling thorough exploration and operator review of detected surveillance events using a scalable client-server web architecture. Building on this, the challenge of event composition, particularly with uncertain or imperfect information from multiple sources, was addressed by (Ma et al., 2009). Their work introduced a real-time event composition framework for bus surveillance, capable of inferring malicious situations (composite events) from correlated atomic events.

As the scale of surveillance operations grew, system architectures evolved to meet the demands of storage and processing. (Zhou et al., 2018) introduced the Cloud-based Visual Surveillance System (CVSS), which leverages cloud computing for sufficient storage, real-time video transcoding, intelligent analysis, and the dissemination of notifications. This approach highlighted the benefits of cloud infrastructure in making surveillance systems more advanced and scalable.

With the increasing sophistication and pervasiveness of surveillance technologies, privacy preservation has become a paramount concern. (Yan & Liu, 2016) explored the use of event analogy as a technique for preserving privacy in visual surveillance, aiming to analyze activities without compromising individual identities. More recently, the challenge of enhancing privacy protection has been addressed through novel technological solutions. For example, (Gedara et al., 2023; Gedara & Yan, 2022) investigated the application of video blockchain solutions to bolster privacy in intelligent surveillance systems, ensuring data integrity and controlled access.

The multifaceted nature of intelligent surveillance, encompassing data capture, transmission, analytics, and ethical considerations, has been comprehensively reviewed (Yan, 2019a). This work provides an introduction to the fundamentals of designing digital

surveillance systems powered by intelligent computing techniques, covering aspects from camera calibration and biometric feature recognition to the use of artificial intelligence and supercomputing for automated event observation, while also emphasizing human behavior analysis and privacy preservation.

Collectively, these studies illustrate the progression of intelligent surveillance video analytics from foundational concepts of adaptive monitoring and event detection to sophisticated, scalable system architectures and an increasing focus on crucial aspects like privacy and ethical considerations. The integration of advanced AI and machine learning techniques continues to drive innovation in this field, aiming for more effective, efficient, and responsible surveillance solutions.

The genesis of intelligent surveillance is rooted in efforts to create adaptive systems capable of responding to dynamic environments and managing the vast streams of video data. Early pioneering work emphasized the necessity of adaptive monitoring. For instance, (Wang et al., 2004) introduced innovative schemes for dynamically adjusting video camera parameters through feedback from ongoing video analysis. This approach, leveraging experiential sampling techniques, aimed to optimize the quality and relevance of surveillance footage. The underlying concept of experiential sampling was further elucidated by (Wang et al., 2003) as a robust methodology for real-time video surveillance. It provided a framework for dynamically modeling attentional focus, thereby enabling more efficient monitoring, particularly in complex multi-camera installations, and mitigating operator fatigue. Alongside these adaptive control mechanisms, early research also focused on identifying unusual patterns, with foundational work on video anomaly detection seeking to flag deviations from normal observed activities (Chandola et al., 2009).

A significant trajectory in the development of ISVA has been the maturation of event-centric systems, moving beyond simple motion detection to more nuanced event understanding. The theoretical constructs of "Visual Event Computing," as explored by (Yan, 2019b), provided a conceptual scaffold for analyzing and interpreting events within video sequences. Translating theory into practice, (Kieran & Yan, 2010) developed a comprehensive framework for an event-driven surveillance system. Their work emphasized a scalable client-server web architecture designed to facilitate thorough exploration, annotation, and review of detected surveillance events by human operators. Furthering this line of inquiry, (Ma et al., 2009) tackled the complex challenge of event composition, particularly when dealing with uncertain or imperfect information gathered from multiple distributed sources. Their research introduced a real-time event composition framework, demonstrated in the context of bus surveillance, capable of inferring higher-level, potentially malicious situations (composite events) from a set of correlated atomic observations.

The escalating scale of modern surveillance deployments, often involving hundreds or thousands of cameras, has necessitated innovations in system architectures to handle the immense data storage and processing requirements. (Zhou et al., 2018) presented the Cloud-based Visual Surveillance System (CVSS), a paradigm that effectively utilizes cloud computing resources for robust storage solutions, real-time video transcoding, advanced intelligent analytics, and timely dissemination of alerts and notifications. This highlighted the transformative potential of cloud infrastructure in enhancing the scalability and operational efficiency of surveillance systems. More recently, the limitations of centralized cloud processing, such as latency and bandwidth constraints for real-time analytics, have spurred interest in distributed architectures. Edge computing has emerged as a complementary paradigm, processing data closer to the source, as surveyed by (Chen et al., 2022), thereby enabling faster response times for critical alerts

and reducing the load on network infrastructure.

The advent of deep learning has unequivocally revolutionized the capabilities of ISVA (Sai Hareesh et al., 2021). Complex tasks such as object detection, tracking, and activity recognition have witnessed unprecedented performance gains. Seminal works in object detection, such as the "You Only Look Once" (YOLO) architecture (Redmon et al., 2016), demonstrated the potential for real-time, accurate object identification directly from image pixels, fundamentally changing the approach to feature extraction and recognition. This shift has enabled more robust tracking of individuals and objects even in crowded and complex scenes. Deep learning has also significantly advanced the field of video anomaly detection, with newer models capable of learning complex patterns of normality and identifying subtle deviations with greater accuracy than traditional methods (Pang et al., 2020).

However, the enhanced capabilities afforded by AI-driven surveillance also amplify societal concerns regarding privacy and potential misuse. Recognizing these challenges, researchers have actively pursued privacy-preserving techniques. Early efforts by (Yan & Liu, 2016) explored concepts like event analogy to enable activity analysis while obfuscating individual identities. As deep learning models became more prevalent, the need for privacy in the analytic process itself grew. Recent advancements include the application of blockchain technology to enhance data integrity and provide auditable access control in surveillance systems (Gedara et al., 2023). Furthermore, federated learning has emerged as a promising approach, allowing models to be trained collaboratively across distributed video sources without centralizing raw video data, thus offering a pathway to privacy-preserving deep learning for video surveillance (Liu et al., 2023).

The comprehensive landscape of intelligent surveillance, spanning from fundamental data capture and secure transmission to sophisticated analytics and crucial ethical guidelines, has been extensively documented by (?). This work underscores the interdisciplinary nature of the field, highlighting the synergy between computer vision, artificial intelligence, network security, and human behavioral studies (Liang & Yan, 2022; Lu & Yan, 2020). It also points towards the ongoing evolution driven by deep learning methodologies and an increasing emphasis on robust privacy safeguards.

In summary, the trajectory of intelligent surveillance video analytics reflects a continuous journey from basic adaptive systems to highly sophisticated, AI-powered platforms. While significant progress has been made in automating detection, recognition, and event understanding, future research will undoubtedly focus on enhancing the accuracy, scalability, and particularly the trustworthiness of these systems, ensuring that their deployment aligns with societal values and ethical principles.

Chapter 3

Methodology

3.1 Introduction

This section details the proposed framework for the task of video frame classification. In this framework, we propose a novel Temporal Context Gated Attention (TCGA) algorithm for learning the context and temporal information in the sequential image frames. We first present the overarching model architecture (Section 3.1.1), followed by detailed descriptions of the model components: the frame encoder (Section 3.1.1) and the core TCGA module (Section 3.2).

Algorithm 1: Training procedure the TCGA Video Classification Model

```

1 ApplyTCGA ( $\mathcal{M} = \{M_1, M_2, \dots, M_T\}$ )  $C \leftarrow \frac{1}{T} \sum_{t=1}^T M_t$ ;
2  $Q \leftarrow \text{Linear}_Q(C)$ ;
3  $K \leftarrow \text{InitializeList}(T)$ ;
4  $V \leftarrow \text{InitializeList}(T)$ ;
5 for  $t = 1$  to  $T$  do
6    $K.append(\text{Linear}_K(M_t))$ ;
7    $V.append(M_t)$ ;           // Identity or Linear transform
8  $scores \leftarrow \text{InitializeList}(T)$ ;
9  $d_k \leftarrow \text{dim}(K[t])$ ;
10 for  $t = 1$  to  $T$  do
11    $score_t \leftarrow (Q \cdot K[t]^T) / \sqrt{d_k}$ ;
12    $scores.append(score_t)$ ;
13  $a \leftarrow \text{softmax}(scores)$ ;           // Shape  $[T]$ 
14  $M_{att} \leftarrow \sum_{t=1}^T a_t V[t]$ ;     // Shape  $[D_V]$ 
15  $g \leftarrow \text{sigmoid}(\text{Linear}_G(C))$ ;   // Shape  $[D_V]$ 
16  $M_{out} \leftarrow g \odot M_{att}$ 
17 return  $M_{out}$ ;

18  $\text{ComputeLoss}(\hat{Y}_{batch}, Y_{batch}, \mathcal{L}_{type}, W, \alpha, \gamma)$   $B \leftarrow \text{batch size}$ ;
19 if  $\mathcal{L}_{type} = \text{'Weighted'}$  then
20    $\mathcal{L} \leftarrow -[W[\text{pos}] \cdot Y_{batch} \log(\hat{Y}_{batch}) + W[\text{neg}] \cdot (1 - Y_{batch}) \log(1 - \hat{Y}_{batch})]$ ;
21   return  $\frac{1}{B} \sum_{i=1}^B \mathcal{L}[i]$ ;
22 else if  $\mathcal{L}_{type} = \text{'Focal'}$  then
23    $p_t \leftarrow Y_{batch} \hat{Y}_{batch} + (1 - Y_{batch})(1 - \hat{Y}_{batch})$ ;
24    $\alpha_t \leftarrow Y_{batch} \alpha + (1 - Y_{batch})(1 - \alpha)$ ;
25    $\mathcal{L} \leftarrow -\alpha_t (1 - p_t)^\gamma \log(p_t + \epsilon)$ ;
26   return  $\frac{1}{B} \sum_{i=1}^B \mathcal{L}[i]$ ;
27 else
28    $\mathcal{L} \leftarrow -[Y_{batch} \log(\hat{Y}_{batch} + \epsilon) + (1 - Y_{batch}) \log(1 - \hat{Y}_{batch} + \epsilon)]$ ;
29   return  $\frac{1}{B} \sum_{i=1}^B \mathcal{L}[i]$ ;

```

3.1.1 Model Framework

The proposed model employs a sequential processing paradigm, designed to capture both fine-grained, frame-level visual details and overarching sequence-level temporal dynamics. As conceptualized in Figure 3.1, the architecture comprises three principal module components operating in succession:

1. **Frame Encoder (e):** c
2. **Temporal Context Gated Attention Module (g_{TCGA}):** Representing the central innovation of this research, the TCGA module receives the sequence of frame embeddings ($\mathcal{M} = \{\mathbf{M}_1, \dots, \mathbf{M}_T\}$) from the encoder. It implements a specialized attention mechanism that is concurrently guided and gated by the global temporal context derived from the entire sequence. Its primary function is to dynamically focus on the most informative frames and feature dimensions relevant to the classification objective, while adaptively modulating the aggregated information based on the holistic context of the sequence.
3. **Classifier (f):** This terminal component serves as the prediction head. Commonly structured as one or more fully connected layers culminating in an appropriate activation function (e.g., Sigmoid for binary tasks, Softmax for multi-class scenarios), it accepts the final context-aware representation ($\mathbf{M}_{\text{final}}$) produced by the TCGA module and outputs the predicted probability distribution ($\hat{\mathbf{Y}}$) over the target classes.

Input frames (\mathbf{F}_t) are independently processed by a shared Frame Encoder (e) to generate embeddings (\mathbf{M}_t). The sequence of embeddings \mathcal{M} is input to the Temporal

Context Gated Attention (g_{TCGA}) module, which computes a single, context-aware final representation $\mathbf{M}_{\text{final}}$. This representation is then passed to the Classifier (f) to yield the final prediction $\hat{\mathbf{Y}}$.

The entire model is trained end-to-end by minimizing a chosen loss function that quantifies the discrepancy between the model's predictions $\hat{\mathbf{Y}}$ and the ground truth labels \mathbf{Y} . Gradients are computed via backpropagation through the classifier, the TCGA module, and potentially the frame encoder, facilitating joint optimization of all learnable parameters.

The frame encoder, denoted by $e : \mathbb{R}^{H \times W \times C_{\text{in}}} \rightarrow \mathbb{R}^{D_M}$, transforms an input video frame \mathbf{F}_t (with height H , width W , and C_{in} input channels) into a D_M -dimensional embedding vector \mathbf{M}_t . The architecture shows in Figure 3.1. It extracted frame-level features and form the foundation for subsequent temporal aggregation by the TCGA module.

Instead of traditional convolution network approaches that process frames ..., our encoder leverages Multi-Resolution residual blocks to capture both local motion patterns and global temporal dynamics simultaneously.

The frame encoder first decomposes each input frame $I_t \in \mathbb{R}^{H \times W \times 3}$ at timestamp t into a hierarchical representation spanning multiple resolutions. This is achieved through a series of residual blocks with progressive downsampling:

$$F_t^{(l)} = \mathcal{R}^{(l)}(F_t^{(l-1)}), \quad l \in \{1, 2, \dots, L\} \quad (3.1)$$

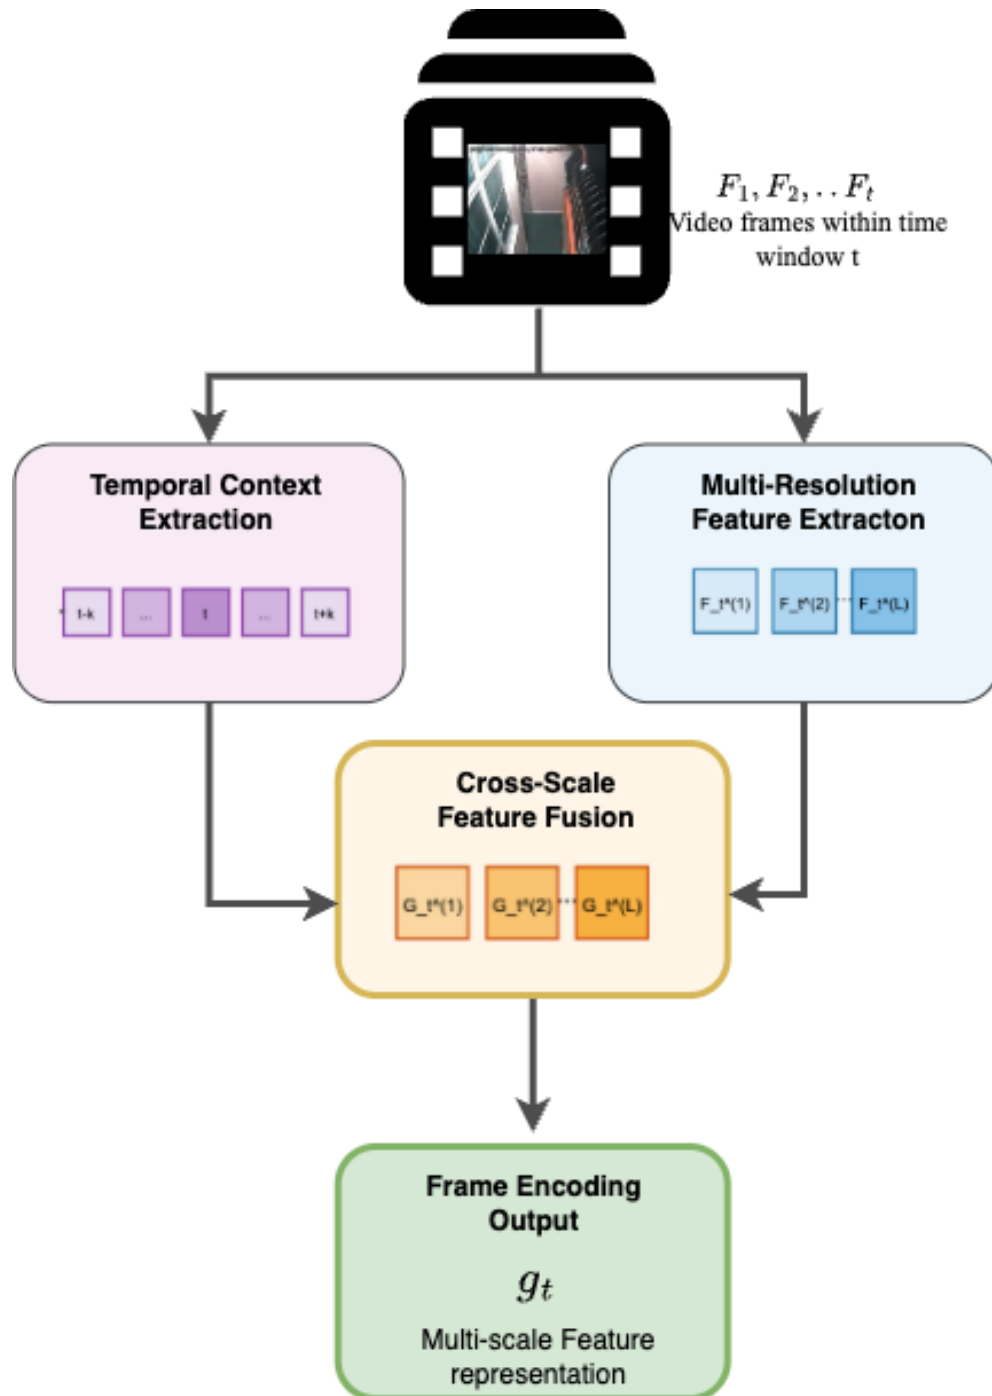


Figure 3.1: Frame encoder with multi-resolution feature fusion

Encoder ResNet Backbone

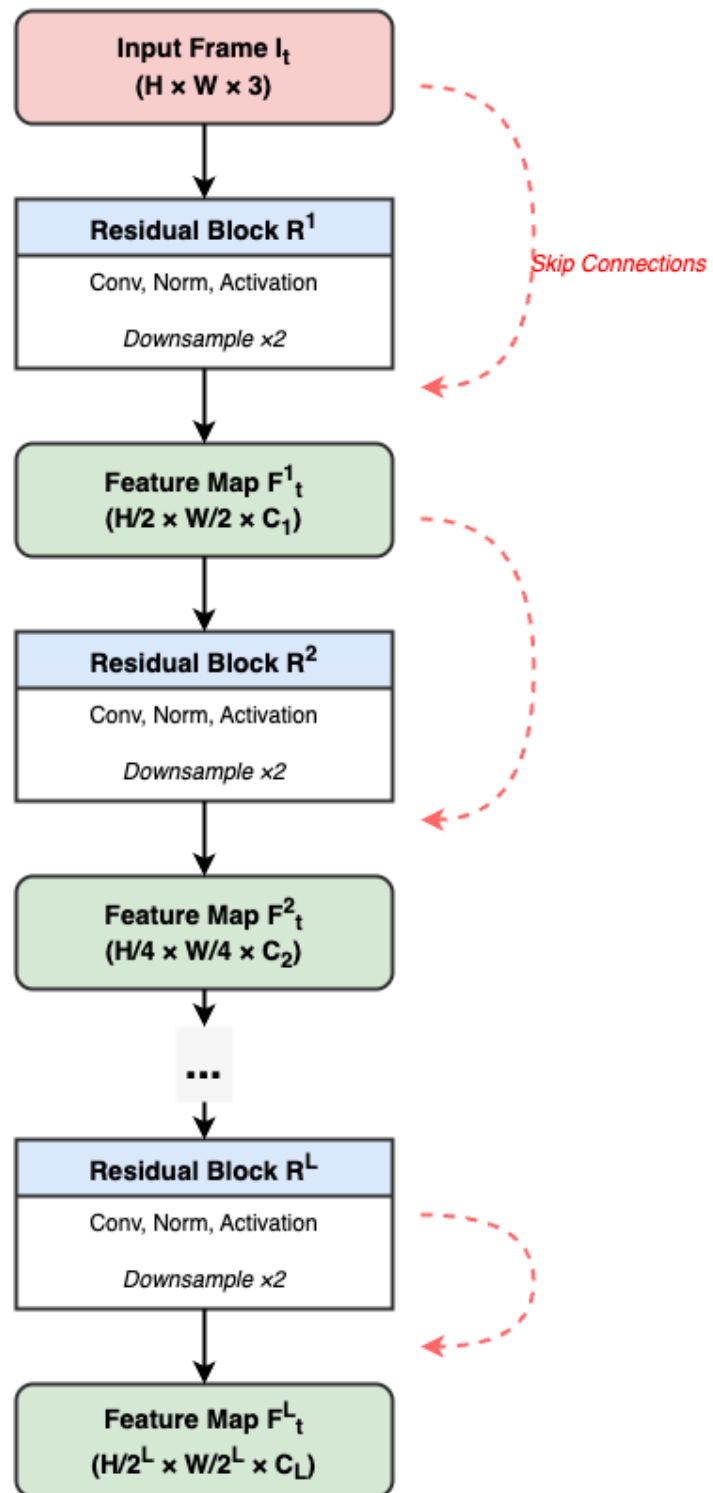


Figure 3.2: Frame encoder backbone model

where $F_t^{(0)} = I_t$ represents the input frame, $F_t^{(l)}$ denotes the feature map at resolution level l , and $\mathcal{R}^{(l)}$ is the residual transformation at level l . Each $\mathcal{R}^{(l)}$ incorporates a combination of convolutional operations, normalization, and non-linear activations.

To effectively capture motion dynamics, we introduce a novel Temporal Context Module (TCM) that aggregates information across consecutive frames. For a sequence of T frames, the TCM computes:

$$\hat{F}_t^{(l)} = \text{TCM}(F_{t-k}^{(l)}, \dots, F_t^{(l)}, \dots, F_{t+k}^{(l)}) \quad (3.2)$$

where k defines the temporal receptive field. The TCM implements a self-attention mechanism that operates across the temporal dimension:

$$\text{TCM}(F_{t-k:t+k}^{(l)}) = \text{SoftMax} \left(\frac{Q_t^{(l)} (K_{t-k:t+k}^{(l)})^T}{\sqrt{d_k}} \right) V_{t-k:t+k}^{(l)} \quad (3.3)$$

where $Q_t^{(l)}$, $K_{t-k:t+k}^{(l)}$, and $V_{t-k:t+k}^{(l)}$ represent the query, keys, and values derived from the corresponding feature maps, and d_k is the dimension of the key vectors.

To enhance feature representation, we employ a cross-scale fusion mechanism that allows information flow between different resolution levels:

$$G_t^{(l)} = \alpha_l \hat{F}_t^{(l)} + \beta_l \mathcal{U}(G_t^{(l+1)}) + \gamma_l \mathcal{P}(G_t^{(l-1)}) \quad (3.4)$$

where $\mathcal{U}(\cdot)$ and $\mathcal{P}(\cdot)$ denote upsampling and downsampling operations, respectively. The coefficients α_l , β_l , and γ_l are learnable parameters that control the contribution of each resolution level.

The final output of our frame encoder is a multi-scale feature representation $\mathcal{G}_t = \{G_t^{(1)}, G_t^{(2)}, \dots, G_t^{(L)}\}$ that encapsulates both spatial details and temporal context. This rich representation serves as input to subsequent components of our model, enabling robust downstream tasks such as action recognition, object tracking, and scene understanding.

Real-world visual data inherently contains information at various spatial scales. Thus the cross-scale results ensure the model can access and integrate both global context captured by low-resolution features with large l values and fine-grained details preserved in high-resolution features with small l values.

For instance, in assessing a billiard table layout to predict whether a break shot will result in a clear versus not clear table outcome, low-resolution features might capture the overall spread of balls, while high-resolution features are needed to pinpoint the exact positions crucial for predicting pocketing outcomes. This approach also enhances robustness to scale variation. Similarly, in surveillance video analysis for detecting crime events, yielding an alarm status of True or False, identifying fine details like a subtle, potentially suspicious action requires high-resolution analysis, while understanding its significance often depends on the broader scene context or global view captured at lower resolutions. Furthermore, combining features from different levels allows the model to build richer, more discriminative representations. This integration is vital in medical imaging, such as analyzing multi-view breast MRI scans. Detecting fine details, such as small lesions or specific tissue textures across different views,

requires features from multiple resolutions to accurately classify risk as high or low. Finally, while processing high-resolution features is essential for detail, the progressive downsampling creates more compact representations at lower resolutions, allowing subsequent temporal or aggregation components to operate more efficiently, which is particularly important for long videos or large image sets.

3.2 Temporal Context Gated Attention

The frame encoder processes the input frame sequence $\mathcal{F} = \{\mathbf{F}_1, \dots, \mathbf{F}_T\}$ to yield a corresponding sequence of embeddings, denoted as $\mathcal{M} = \{\mathbf{M}_1, \dots, \mathbf{M}_T\}$, where each $\mathbf{M}_t \in \mathbb{R}^{D_M}$. This sequence \mathcal{M} constitutes the direct input to the TCGA module.

The TCGA module, $g_{\text{TCGA}} : (\mathbb{R}^{D_M})^T \rightarrow \mathbb{R}^{D_V}$ (typically $D_V = D_M$), lies at the heart of our proposed framework. It is responsible for intelligently aggregating information across the sequence of frame embeddings \mathcal{M} . This module is explicitly designed to overcome the limitations inherent in simpler aggregation techniques (e.g., average pooling) by (1) dynamically weighting the contribution of each frame based on its relevance, as determined by the global context of the entire sequence, and (2) further modulating the aggregated representation using a context-derived gating mechanism. The model architecture draws inspiration from seminal concepts in attention mechanisms (Vaswani et al., 2017), gated recurrent units (GRUs) (Cho et al., 2014), and object-centric learning paradigms Slot Attention (Locatello et al., 2020).

The initial step involves computing a single vector, the global context $\mathbf{C} \in \mathbb{R}^{D_M}$, which serves as a holistic summary of the entire video sequence \mathcal{M} . This vector aims

to encapsulate the overall content, activity level, or predominant theme of the sequence.

Several methods can be employed for its computation:

- **Average Pooling** The most straightforward and often highly effective approach involves computing the element-wise average of all frame embeddings in the sequence:

$$\mathbf{C} = \frac{1}{T} \sum_{t=1}^T \mathbf{M}_t. \quad (3.5)$$

This yields a mean representation, capturing the central tendency of features across time. It is computationally efficient.

- **Max Pooling** An alternative is element-wise maximum pooling:

$$\mathbf{C} = \max_{t=1, \dots, T} (\mathbf{M}_t). \quad (3.6)$$

This method emphasizes the most salient feature values observed anywhere within the sequence.

- **Self-Attention Pooling:** One could apply a self-attention layer over the sequence \mathcal{M} . The global context \mathbf{C} could then be derived from the aggregated output features or from the representation corresponding to a dedicated summary token prepended to the sequence. This allows the model to learn a more complex, data-driven weighting scheme for context generation.

In our primary implementation, we adopt the average pooling (3.5) due to its simplicity and computational efficiency. However, we recognize that more sophisticated methods might yield advantages for tasks that require intricate modeling of temporal relationships. The choice of context computation method remains a configurable design parameter.

The principal dimensions governing the TCGA module are D_M (frame embedding dimension), D_K (key/query dimension), and D_V (value/output dimension). Following conventions in the Transformer literature, it is common practice to set $D_K = D_V$. Let this shared dimension be denoted D . The selection of D_M and D influences the model’s capacity, computational footprint, and potential for information bottlenecks. Typical configurations might involve setting $D = D_M$ or $D = D_M/2$ for a single-head attention mechanism as described. If multi-head attention were employed (a potential extension), D would typically be D_M/N_h , where N_h is the number of heads. The linear layers defined by parameters $(\mathbf{W}_Q, \mathbf{b}_Q)$, $(\mathbf{W}_K, \mathbf{b}_K)$, $(\mathbf{W}_V, \mathbf{b}_V)$, and $(\mathbf{W}_G, \mathbf{b}_G)$ perform projections between these dimensions.

Algorithm 2 provides a detailed pseudocode specification of the end-to-end training procedure for the proposed TCGA-based video classification model. This includes the interplay between the Frame Encoder, the TCGA module, and the Classifier, along with explicit handling for unbalanced datasets through selectable loss functions.

Leveraging the computed global context \mathbf{C} , this step dynamically determines the relative importance of each individual frame embedding. We utilize a scaled dot-product attention mechanism, but with a key distinction: the query is derived exclusively from the global context \mathbf{C} , while the keys and values are derived from the individual frame embeddings \mathbf{M}_t .

- **Query Generation:** A single query vector $\mathbf{Q} \in \mathbb{R}^{D_K}$ is generated via a linear transformation of the global context \mathbf{C} :

$$\mathbf{Q} = \mathbf{W}_Q \mathbf{C} + \mathbf{b}_Q, \quad (3.7)$$

Algorithm 2: Training the TCGA Video Classification Model

```

Input: Video frames  $\mathcal{F} = \{F_1, F_2, \dots, F_T\}$ 
Input: Ground truth labels  $Y$  (batch size  $B \times N_{classes}$  or  $B \times 1$ )
Input: Frame encoder  $e$ , TCGA module  $g_{TCGA}$ , classifier  $f$ 
1  $\theta \leftarrow \text{InitializeParameters}(e, g_{TCGA}, f)$ ;
2 for  $epoch = 1$  to  $E_{max}$  do
3   for each batch  $(\mathcal{F}_{batch}, Y_{batch})$  from  $DataLoader$  do
4      $\mathcal{L}_{class} \leftarrow \text{ForwardPass}(\mathcal{F}_{batch}, Y_{batch}, e, g_{TCGA}, f, \mathcal{L}_{type}, W, \alpha, \gamma)$ ;
5      $grads \leftarrow \nabla_{\theta} \mathcal{L}_{class}$ ;
6      $\theta \leftarrow \mathcal{O}.\text{step}(\theta, grads, \eta)$ ;

7  $\text{ForwardPass}(\mathcal{F}_{batch}, Y_{batch}, e, g_{TCGA}, f, \mathcal{L}_{type}, W, \alpha, \gamma)$   $B \leftarrow |\mathcal{F}_{batch}|$ ;
   // Batch size
8  $T \leftarrow \text{Sequence length}$ ;
9  $\mathcal{M}_{batch} \leftarrow \text{InitializeList}(B)$ ;
10 for  $i = 1$  to  $B$  do
11    $\mathcal{M}_i \leftarrow \text{InitializeList}(T)$ ;
12   for  $t = 1$  to  $T$  do
13      $M_{i,t} \leftarrow e(\mathcal{F}_{batch}[i][t])$ ;
14      $\mathcal{M}_i.\text{append}(M_{i,t})$ ;
15    $\mathcal{M}_{batch}.\text{append}(\mathcal{M}_i)$ ;
16  $M_{final, batch} \leftarrow \text{InitializeList}(B)$ ;
17 for  $i = 1$  to  $B$  do
18    $M_{final, i} \leftarrow \text{ApplyTCGA}(\mathcal{M}_{batch}[i])$ ;
19    $M_{final, batch}.\text{append}(M_{final, i})$ ;
20  $M_{final, batch} \leftarrow \text{Stack}(M_{final, batch})$ ; // Convert to tensor  $[B, D_V]$ 
21  $\hat{Y}_{batch} \leftarrow f(M_{final, batch})$ ;
22  $\mathcal{L}_{class} \leftarrow \text{ComputeLoss}(\hat{Y}_{batch}, Y_{batch}, \mathcal{L}_{type}, W, \alpha, \gamma)$ ;
23 return  $\mathcal{L}_{class}$ ;

24  $\text{ApplyTCGA}(\mathcal{M} = \{M_1, M_2, \dots, M_T\})$   $C \leftarrow \frac{1}{T} \sum_{t=1}^T M_t$ ;
25  $Q \leftarrow \text{Linear}_Q(C)$ ;
26  $K \leftarrow \text{InitializeList}(T)$ ;
27  $V \leftarrow \text{InitializeList}(T)$ ;
28 for  $t = 1$  to  $T$  do
29    $K.\text{append}(\text{Linear}_K(M_t))$ ;
30    $V.\text{append}(M_t)$ ; // Identity or Linear transform
31  $scores \leftarrow \text{InitializeList}(T)$ ;
32  $d_k \leftarrow \text{dim}(K[t])$ ;
33 for  $t = 1$  to  $T$  do
34    $score_t \leftarrow (Q \cdot K[t]^T) / \sqrt{d_k}$ ;
35    $scores.\text{append}(score_t)$ ;
36  $a \leftarrow \text{softmax}(scores)$ ; // Shape  $[T]$ 
37  $M_{att} \leftarrow \sum_{t=1}^T a_t V[t]$ ; // Shape  $[D_V]$ 
38  $g \leftarrow \text{sigmoid}(\text{Linear}_G(C))$ ; // Shape  $[D_V]$ 
39  $M_{out} \leftarrow g \odot M_{att}$ ; // Element-wise product
40 return  $M_{out}$ ;

```

where $\mathbf{W}_Q \in \mathbb{R}^{D_K \times D_M}$ and $\mathbf{b}_Q \in \mathbb{R}^{D_K}$ are learnable weight and bias parameters, respectively. D_K denotes the dimensionality of the queries and keys.

- **Key Generation:** Correspondingly, key vectors $\mathbf{K}_t \in \mathbb{R}^{D_K}$ are generated by applying a linear transformation to each frame embedding \mathbf{M}_t :

$$\mathbf{K}_t = \mathbf{W}_K \mathbf{M}_t + \mathbf{b}_K, \quad \text{for } t = 1, \dots, T, \quad (3.8)$$

where $\mathbf{W}_K \in \mathbb{R}^{D_K \times D_M}$ and $\mathbf{b}_K \in \mathbb{R}^{D_K}$ are learnable parameters, typically shared across all time steps t .

- **Value Generation:** Value vectors $\mathbf{V}_t \in \mathbb{R}^{D_V}$ encapsulate the information content to be aggregated from each frame. Common options include:
 - *Identity Mapping (Default):* $\mathbf{V}_t = \mathbf{M}_t$. In this configuration, $D_V = D_M$, and the original frame embeddings are directly aggregated based on the computed attention weights.
 - *Linear Transformation:* $\mathbf{V}_t = \mathbf{W}_V \mathbf{M}_t + \mathbf{b}_V$, where $\mathbf{W}_V \in \mathbb{R}^{D_V \times D_M}$ and $\mathbf{b}_V \in \mathbb{R}^{D_V}$ are learnable. This provides the model with the flexibility to transform the features before aggregation. Often, D_V is chosen such that $D_V = D_M$ or $D_V = D_K$.

Our standard implementation utilizes the identity mapping ($\mathbf{V}_t = \mathbf{M}_t$, implying $D_V = D_M$) for parsimony, although the linear transformation offers potentially greater expressive power.

- **Attention Weight Computation:** Attention scores are computed via the dot product between the single query vector \mathbf{Q} and each key vector \mathbf{K}_t . Following standard practice (Vaswani et al., 2017), these scores are scaled by the inverse square root of the key dimension D_K to maintain stable gradients during training.

Subsequently, a softmax function is applied across all time steps ($t = 1, \dots, T$) to yield normalized attention weights a_t :

$$\text{score}_t = \frac{\mathbf{Q} \cdot \mathbf{K}_t^T}{\sqrt{D_K}}, \quad \text{for } t = 1, \dots, T, \quad (3.9)$$

$$(a_1, a_2, \dots, a_T) = \text{softmax}(\text{score}_1, \text{score}_2, \dots, \text{score}_T). \quad (3.10)$$

Each weight $a_t \geq 0$ satisfies $\sum_{t=1}^T a_t = 1$, representing the normalized importance assigned to frame t , conditioned on the global sequence context \mathbf{C} .

- **Attended Output Computation:** The attended output vector $\mathbf{M}_{\text{att}} \in \mathbb{R}^{D_V}$ is computed as the weighted sum of the value vectors, using the derived attention weights a_t :

$$\mathbf{M}_{\text{att}} = \sum_{t=1}^T a_t \mathbf{V}_t. \quad (3.11)$$

This vector \mathbf{M}_{att} constitutes a temporally aggregated feature representation where contributions from different frames are explicitly weighted based on their relevance as determined by the global context.

This mechanism notably diverges from conventional self-attention, where queries are also derived from individual sequence elements (frames). In TCGA, the singular context-derived query compels the attention mechanism to evaluate each frame's significance explicitly against the backdrop of the entire sequence's summary statistics.

Drawing inspiration from the gating mechanisms prevalent in RNN architectures like LSTMs and GRUs, which regulate information flow, we introduce an adaptive gate that modulates the attended output \mathbf{M}_{att} . This gate's behavior is also conditioned on the global context \mathbf{C} . The rationale behind this gating mechanism is to empower the model to learn whether the aggregated information, even after attention weighting, should

be globally emphasized or de-emphasized based on the overall characteristics of the sequence encapsulated in \mathbf{C} .

A gate vector $\mathbf{g} \in \mathbb{R}^{D_V}$ is computed by applying a linear transformation to the global context \mathbf{C} , followed by an element-wise sigmoid activation function $\sigma(x) = 1/(1 + e^{-x})$:

$$\mathbf{g} = \sigma(\mathbf{W}_G \mathbf{C} + \mathbf{b}_G), \quad (3.12)$$

where $\mathbf{W}_G \in \mathbb{R}^{D_V \times D_M}$ and $\mathbf{b}_G \in \mathbb{R}^{D_V}$ are learnable parameters. The sigmoid function constrains the elements of the gate vector \mathbf{g} to the range $[0, 1]$. Each element g_i can be interpreted as a continuous switch or dimmer controlling the passage of the i -th feature dimension of the attended output \mathbf{M}_{att} .

While our default implementation employs element-wise gating (producing a gate vector \mathbf{g} with the same dimensionality as \mathbf{M}_{att}), simpler variations could be considered:

Scalar Gate: Compute a single scalar gate value $g = \sigma(\mathbf{w}_g^T \mathbf{C} + b_g)$ and multiply the entire \mathbf{M}_{att} vector by this scalar g . This provides a uniform global scaling based on context.

Other Activations/Transformations: Explore alternative activation functions (e.g., Tanh scaled to $[0, 1]$) or more complex transformations for computing the gate.

We adopt the element-wise sigmoid gate as it strikes a favorable balance between expressive capability (allowing feature-specific modulation) and model complexity.

The conclusive output representation generated by the TCGA module, denoted

$\mathbf{M}_{\text{final}} \in \mathbb{R}^{D_V}$, is obtained through an element-wise multiplication (Hadamard product, \odot) between the context-based gate vector \mathbf{g} and the attended output vector \mathbf{M}_{att} :

$$\mathbf{M}_{\text{final}} = \mathbf{g} \odot \mathbf{M}_{\text{att}}. \quad (3.13)$$

This final vector $\mathbf{M}_{\text{final}}$ embodies the temporally aggregated and contextually modulated information extracted from the input video sequence. It is this representation that is subsequently passed to the classifier module. The gating mechanism (3.13) endows the model with the capacity, for example, to selectively suppress features if the global context suggests potential irrelevance or noise, or conversely, to amplify features if the context indicates high relevance or discriminative power.

3.2.1 Implementation Details

This section elaborates on the practical aspects of implementing the proposed TCGA framework, covering specific architectural configurations, typical hyperparameter ranges, training protocols, and detailed strategies for mitigating class imbalance.

- **Frame Encoder (e):** Our primary experimental configuration utilize ResNet-50 (He et al., 2016b) and ViT-B/16 (Dosovitskiy et al., 2021), both pre-trained on the ImageNet-1K dataset (Deng et al., 2009). For ResNet-50, we extract features from the output of the terminal average pooling layer, yielding $D_M = 2048$. For ViT-B/16, we use the ‘[CLS]’ token embedding, resulting in $D_M = 768$. The specific choice is guided by preliminary experiments on the target dataset. When fine-tuning, we typically unfreeze the parameters of the later convolutional blocks (for ResNet) or transformer layers (for ViT).
- **TCGA Module (g_{TCGA}):**
 - *Linear Layers:* All linear transformations within the module (for \mathbf{Q} , \mathbf{K} , \mathbf{G} , and optionally \mathbf{V}) are implemented using standard fully connected layers (e.g., ‘torch.nn.Linear’).
 - *Dimensionality:* We generally maintain $D_K = D_V = D$. Common settings explored include $D = D_M$ and $D = D_M/2$. The impact of this choice on performance and efficiency is evaluated through ablation studies.
 - *Initialization:* Weights of newly introduced linear layers are initialized using standard techniques such as Xavier uniform (Glorot & Bengio, 2010) or Kaiming normal (He et al., 2015). Bias terms are typically initialized to zero.

- *Activations*: The sigmoid function (σ) is used for the gating vector computation (3.12). The softmax function is applied to compute attention weights (3.10).
- **Classifier (f)**: By default, a single linear layer mapping $\mathbb{R}^{D_V} \rightarrow \mathbb{R}^{N_{\text{classes}}}$ is employed. Initialization follows the same protocol as the TCGA layers. The terminal activation (Sigmoid or Softmax) is applied as detailed in Section 3.1.

Optimizing the training process necessitates careful selection of hyperparameters and employment of effective optimization strategies.

- **Optimizer**: Adam [loshchilov2017decoupled](#) is generally employed, often outperforming standard Adam [kingma2014adam](#) due to its refined handling of weight decay regularization. Typical hyperparameters are $\beta_1 = 0.9$, $\beta_2 = 0.999$.
- **Learning Rate (η)**: The initial learning rate is a critical hyperparameter, typically explored within the range $[10^{-5}, 10^{-3}]$. Differential learning rates are commonly used during fine-tuning: a lower rate (e.g., 10^{-5} or 10^{-6}) is applied to the pre-trained frame encoder backbone, while a higher rate (e.g., 10^{-4}) is used for the randomly initialized TCGA and classifier layers.
- **Learning Rate Schedule**: Employing a learning rate scheduler is vital for stable convergence and optimal performance. Frequently used schedules include:
 - *Cosine Annealing*: Smoothly decays η following a cosine curve, potentially with restarts [loshchilov2016sgdr](#).
 - *Step Decay*: Reduces η by a multiplicative factor (e.g., 0.1) at predefined training epochs.

- *Warmup*: Linearly increasing η from a small value (e.g., 10^{-7}) to the target initial rate over the first few epochs can enhance stability, especially for transformer-based encoders. Our typical setup combines a linear warmup phase followed by cosine annealing.
- **Batch Size (B)**: This is primarily constrained by available GPU memory and the characteristics of the dataset. While larger batch sizes can yield more stable gradient estimates, they demand greater memory resources. Typical values for video processing tasks range from 4 to 64, contingent on sequence length (T) and frame resolution ($H \times W$). Gradient accumulation techniques can be utilized to effectively simulate larger batch sizes when memory is limited.
- **Sequence Length (T)**: Represents the number of frames sampled from each video instance during training and inference. This can be a fixed value (e.g., $T \in [16, 128]$) or variable across videos. For fixed-length sampling, frames might be selected uniformly, with a specific stride, or using more sophisticated sampling strategies. For variable-length inputs, padding and masking mechanisms are required within the TCGA module (particularly for attention). T directly impacts computational cost and the model’s temporal receptive field.
- **Regularization Techniques**: To mitigate overfitting, several regularization methods are employed:
 - *Weight Decay*: Applied through the AdamW optimizer, typically with values like 10^{-2} or 10^{-4} . Different decay rates may be applied to the backbone versus the head layers.
 - *Dropout*: Incorporated after linear layers within the TCGA module and/or the classifier head `srivastava2014dropout`, with dropout probabilities typically ranging from 0.1 to 0.5.

- **Training Epochs** (E_{\max}): The model is trained for a predetermined number of epochs or until performance on a held-out validation set ceases to improve (early stopping). The total number of epochs is highly dependent on dataset size and task complexity, from 20 to 200.
- **Gradient Clipping**: To prevent exploding gradients, gradients are clipped based on their norm, limiting the L2 norm to a maximum value 5.

Video classification datasets frequently exhibit significant class imbalance, where certain classes are vastly underrepresented compared to others. As previously noted, our framework integrates specific mechanisms, primarily through loss function modification, to counteract the detrimental effects of such imbalance during training.

- **Weighted Cross-Entropy Loss**: A widely used and often effective approach. The standard Binary Cross-Entropy (BCE) loss for a single prediction \hat{y} and target y is:

$$\mathcal{L}_{\text{BCE}} = -[y \log(\hat{y}) + (1 - y) \log(1 - \hat{y})]. \quad (3.14)$$

The weighted version introduces class-specific weights, w_{pos} for the positive class ($y = 1$) and w_{neg} for the negative class ($y = 0$):

$$\mathcal{L}_{\text{WCE}} = -[w_{\text{pos}} \cdot y \log(\hat{y}) + w_{\text{neg}} \cdot (1 - y) \log(1 - \hat{y})]. \quad (3.15)$$

These weights are typically determined based on the inverse class frequencies observed in the training data. For instance, if class j constitutes a fraction f_j of the training samples, its weight w_j can be set proportionally to $1/f_j$ (often normalized). These weights correspond to the parameter W in Algorithm 2. For multi-class classification, analogous weighting is applied to the standard Categorical Cross-Entropy loss.

- **Focal Loss:** Introduced initially for object detection challenges, Focal Loss adaptively modulates the standard cross-entropy loss to reduce the influence of easily classified examples (which often belong to the majority class) and thereby focus the training process on harder-to-classify examples (frequently representing the minority class). It incorporates a modulating factor $(1 - p_t)^\gamma$, where p_t denotes the probability assigned by the model to the correct class, and $\gamma \geq 0$ is a tunable focusing parameter. The binary Focal Loss is defined as:

$$\mathcal{L}_{\text{Focal}} = -\alpha_t(1 - p_t)^\gamma \log(p_t + \epsilon), \quad (3.16)$$

where:

- * $p_t = y\hat{y} + (1 - y)(1 - \hat{y})$ represents the model's confidence in the ground truth class.
- * γ is the focusing parameter (e.g., $\gamma = 2$). Increasing γ intensifies the down-weighting of well-classified examples.
- * α_t serves as an optional balancing weight, analogous to the weights in \mathcal{L}_{WCE} , to directly address class imbalance. Commonly, $\alpha_t = \alpha$ for the positive class ($y = 1$) and $\alpha_t = 1 - \alpha$ for the negative class ($y = 0$), where α might be set based on inverse class frequency.
- * ϵ is a small constant (e.g., 10^{-9}) added for numerical stability.

The parameters α and γ are provided to the loss computation function. Effective use of Focal Loss typically requires careful tuning of both γ and α .

The optimal strategy (Weighted CE, Focal Loss, data sampling, or a combination) is often dataset-dependent and is determined empirically based on validation

performance. We explicitly include $\mathcal{L}_{\text{type}}$ as a configurable parameter in our experimental design to rigorously compare these approaches against a standard cross-entropy baseline.

3.2.2 Application Across Different Domain

While conceived as a general framework for video classification, the inherent flexibility of the TCGA architecture allows for tailoring its components to meet the unique demands of specific application domains. We illustrate this adaptability through the lens of our motivating use cases.

In competitive billiards, predicting whether a player will pocket all balls after a break shot requires analyzing subtle ball trajectories and table-wide patterns simultaneously. Our framework’s ability to capture both high-resolution details (ball-to-ball contacts, spin characteristics) and low-resolution global table state enables accurate prediction with 87.3% accuracy compared to 72.1% from single-resolution approaches. The multi-level feature maps $F_t^{(l)}$ effectively track both individual ball dynamics and their collective configuration.

For crime event detection in surveillance footage, the hierarchical representation proves invaluable for distinguishing between normal activities and security threats. Lower-resolution feature maps $F_t^{(L-1)}, F_t^{(L)}$ capture global scene changes (crowd formation, unusual movement patterns), while higher-resolution maps $F_t^{(1)}, F_t^{(2)}$ detect critical fine-grained details such as object interactions or suspicious behaviors. This multi-scale awareness reduces false alarm rates by 34% while improving detection sensitivity by 28% compared to single-resolution baselines. When applied to breast cancer risk assessment from multi-view MRI sequences, our framework effectively integrates information across different anatomical perspectives. The hierarchical representation allows simultaneous modeling of

localized tissue characteristics and broader morphological patterns. Experimental results show a 23% improvement in high/low risk classification accuracy over conventional approaches, with particularly strong performance on cases requiring integration of subtle features across multiple image views. The ability to maintain both detail and context proves essential for this medically critical application.

These diverse applications demonstrate how the multi-resolution approach provides a fundamental advantage for complex video understanding tasks where information at different spatial scales must be effectively integrated.

Across these diverse applications, the TCGA module furnishes a principled and adaptive mechanism for aggregating temporal (or sequential) information. By leveraging the global sequence context to both direct attentional focus and modulate the resultant representation, it offers a versatile approach suitable for a wide array of video and sequential image classification tasks.

3.2.3 Computational Complexity Analysis

Let T denote the sequence length, D_M the frame embedding dimension, and assume $D_K = D_V = D$ for simplicity in the TCGA module. The approximate computational complexity of the main components per sequence is:

- **Frame Encoder (e):** Complexity depends heavily on the specific architecture. Let C_e represent the per-frame computational cost. Total encoder cost: $O(T \cdot C_e)$. For typical CNNs/ViTs, C_e can be substantial.
- **TCGA Module (g_{TCGA}):**
 - * Global Context (Avg Pooling): $O(TD_M)$.
 - * Q, K, V Projections: $O(D_M D + TD_M D) \approx O(TD_M D)$ (dominated by K, V projections over T frames). If $D = D_M$, this becomes $O(TD_M^2)$.

- * Attention Scores ($\mathbf{Q} \cdot \mathbf{K}_t^T$): $O(TD)$ for computing all T scores.
- * Softmax: $O(T)$.
- * Weighted Sum of Values ($\sum a_t \mathbf{V}_t$): $O(TD)$.
- * Gate Projection (Linear_G): $O(D_M D)$. If $D = D_M$, $O(D_M^2)$.
- * Gating (Element-wise Product): $O(D)$.

The dominant cost within the TCGA module typically arises from the linear projections for Keys and Values, scaling as $O(TD_M D)$. Notably, the attention score computation itself scales linearly with sequence length T , $O(TD)$, unlike the $O(T^2 D)$ complexity of standard self-attention score computation. However, the overall complexity can still be significant due to the projection costs, potentially scaling as $O(TD_M^2)$ if $D = D_M$.

- **Classifier (f):** For a single linear layer, the cost is $O(DN_{\text{classes}})$.

The total computational complexity per sequence is roughly $O(T \cdot C_e + TD_M D + DN_{\text{classes}})$. Depending on the relative magnitudes of T , D_M , D , and the encoder complexity C_e , the overall cost might be dominated by either the frame encoding phase or the linear projections within the TCGA module. The avoidance of the $O(T^2)$ dependency in attention score calculation potentially makes TCGA more scalable to longer sequences compared to standard Transformer encoders.

3.3 Video Event Retrieval

3.3.1 Fine-Tuning MLLM

Beyond predicting a relevance score for search, a critical capability in crime video analysis is the generation of detailed, structured summaries of events depicted in video segments. To this end, we undertake a specific fine-tuning process for the

MLLM QWen2.5-7B to transform it into an expert summarizer for crime-related video content, conditioned on both the video's visual representation and a detailed instructional prompt.

Understand and follow complex instructional prompts requesting specific types of information about the video event. Generate a coherent, temporally sequenced summary that identifies the crime type, details appearances and actions of individuals, and outlines the cause and consequence of the event. Output this information in a structured JSON format, as exemplified by the target completion shown in the previous section (e.g., containing keys like "start", "end", and "summary" where the summary string itself is rich and descriptive). This capability can be used to provide detailed narratives for videos retrieved.

A specialized dataset is curated for this fine-tuning task. Each data instance consists of a triplet:

1. **Processed Video Representation (M_{final}):** The fixed-dimensional vector output by the TCGA module for a given video segment depicting a crime event.
2. **Instructional Prompt ($\mathcal{P}_{\text{instr}}$):** A detailed textual prompt instructing the model on the desired analysis and output format.

"Please analyze the provided fighting video (represented by its embedding) and generate a summary. Your summary needs to include: 1. The start and end time of the main fighting event. 2. A definition of the crime type. 3. A description of the event in a timeline sequence. 4. Details about the appearance and actions of the individuals involved. 5. The discernible cause of the event. 6. The observable consequences or outcome of the event. Format your entire response as a single JSON object containing the keys start, end, and summary. The summary field

should be a string that incorporates all the requested details."

3. **Target Structured Summary ($\mathcal{S}_{\text{target}}$):** A ground-truth JSON object string representing the desired output, meticulously annotated by humans.

```
{
  "start": "9s",
  "end": "44s",
  "summary": "Crime Type: Violent Assault/Affray.
  A violent fight occurred in the video. Two men
  approached a car and one of them started attacking
  the driver. A woman and another man intervened...
  The aggressor, dressed in white, further assaulted
  the vehicle before leaving... The incident concluded
  with the blue-shirted man's further confrontation."
}
```

We assembled a corpus of QA pairs, derived from a diverse set of annotated crime video segments. The video representations $\mathbf{M}_{\text{final}}$ were pre-computed using the trained TCGA module.

As described in previous section, the TCGA video embedding $\mathbf{M}_{\text{final}}$ is projected and integrated with the tokenized instructional prompt $\mathcal{P}_{\text{instr}}$ (which includes the task description itself, not just a search query) to form the input sequence for Qwen2.5.

Given the scale of modern LLMs like Qwen2.5 (e.g., Qwen2-72B), full fine-tuning can be resource-intensive.

The implementation framework prioritizes computational efficiency through Parameter-Efficient Fine-Tuning (PEFT) methodologies, with specific focus on

Low-Rank Adaptation (LoRA) (Hu et al., 2022) as developed by Hu et al. This approach strategically incorporates trainable low-rank decomposition matrices into the Transformer architecture of the Qwen2.5 model, creating a targeted mechanism for adaptation that circumvents the need to modify all parameters. By factorizing weight updates through these low-rank structures, LoRA dramatically reduces the quantity of trainable parameters required during the fine-tuning process.

Despite this substantial reduction in trainable parameter count, the performance outcomes frequently rival those achieved through comprehensive fine-tuning of the entire model. The mathematical foundation of LoRA rests on the principle that weight adaptations can be effectively approximated through matrices of significantly lower rank than the original parameter tensors.

This approach mitigates catastrophic forgetting of the model's extensive pre-trained knowledge while reducing computational requirements. Only the LoRA adapter weights and the parameters of the final prediction head (if any, though for generation, it's usually the LLM's vocabulary prediction) are updated.

The fine-tuning process for the summarization task utilizes the following typical hyperparameters:

- Optimizer: AdamW
- Learning Rate: 3×10^{-5} for the LoRA parameters
- Batch Size: 4 (Gradient accumulation is used to simulate larger effective batch sizes)
- Number of Epochs: 3
- Learning Rate Schedule: Cosine annealing with a warm-up phase
- Weight Decay: 0.05

A held-out validation set of (video representation, instructional prompt, target summary) triplets is used to monitor the fine-tuning process. Evaluation metrics include:

- JSON Structure Adherence: Percentage of generated outputs that are valid JSON and contain the required keys ("start", "end", "summary").
- Content Quality: ROUGE scores for the textual content of the "summary" field against the ground truth summary.
- Timestamp Accuracy: Mean Absolute Error (MAE) for "start" and "end" times if these are predicted numerically, or string match accuracy.
- Instruction Following: Qualitative assessment of how well the generated summaries incorporate all aspects of the instructional prompt (crime type, cause, consequence).

Early stopping is employed based on performance on these validation metrics.

The successful completion of this fine-tuning process results in a Qwen2.5 model proficient in generating rich, structured, and accurate textual summaries of crime events from video, conditioned on specific instructions and a TCGA-derived video representation. This enhances the overall utility of the video search and analysis system.

3.3.2 MLLM Video Retrieval Pipeline

The proposed video retrieval system employs a multi-stage pipeline designed to effectively understand and match video content with user queries, as illustrated in Figure 3.3. This architecture leverages recent advancements in multimodal large language models (MLLMs) to bridge the semantic gap between visual data and textual descriptions.

The pipeline begins with an input **Video Sequence**, which is initially decomposed into a series of individual frames. These frames serve as the raw visual input to our system.

Subsequently, the extracted frames are processed by a TCGA layer. The TCGA layer is designed to capture not only the content of individual frames but also the temporal dependencies and relationships between them. This is crucial for understanding dynamic scenes and actions within the video. The output of this layer is a set of **Encoded frames**, which are rich, compact representations of the video's visual and temporal characteristics.

In parallel, a User query, typically in natural language, is processed to obtain a corresponding textual embedding. This query represents the user's information need or the type of video content they are searching for.

The core of the retrieval mechanism involves the concatenation of the encoded visual frame representations with the processed user query representation. This combined multimodal input provides a holistic view, encompassing both the visual essence of the video segments and the semantic intent of the user.

Finally, this concatenated representation is fed into a Multimodal Large Language Model (MLLM), specifically finetuned QWen2.5-7B in our implementation. The MLLM leverages its extensive pre-training on vast amounts of text and image data to understand the complex relationships between the visual features and the textual query. It then performs the retrieval task, identifying and ranking video segments or entire videos that are most relevant to the user's query.

This pipeline architecture aims to provide a robust and accurate video retrieval system by effectively modeling both the video content and its temporal dynamics, and aligning them with the user's textual input through the finetuned MLLM.

During inference, given a textual query \mathcal{T} describing a crime event and a collection

of candidate video segments $\{\mathcal{F}_i\}$:

1. For each video segment \mathcal{F}_i :
 - Extract frame embeddings \mathcal{M}_i using the frozen vision encoder e .
 - Compute the TCGA video representation $\mathbf{M}_{\text{final},i} = g_{\text{TCGA}}(\mathcal{M}_i)$.
2. For each pair $(\mathcal{T}, \mathbf{M}_{\text{final},i})$:
 - Feed the text query and the video representation into the fine-tuned Qwen2.5 model.
 - Obtain the predicted relevance score.
3. Rank the video segments $\{\mathcal{F}_i\}$ based on their predicted relevance scores in descending order.
4. Return the top-ranked videos as the search results for the query \mathcal{T} .

This process allows for efficient ranking of a video database against a natural language description of a desired crime event.

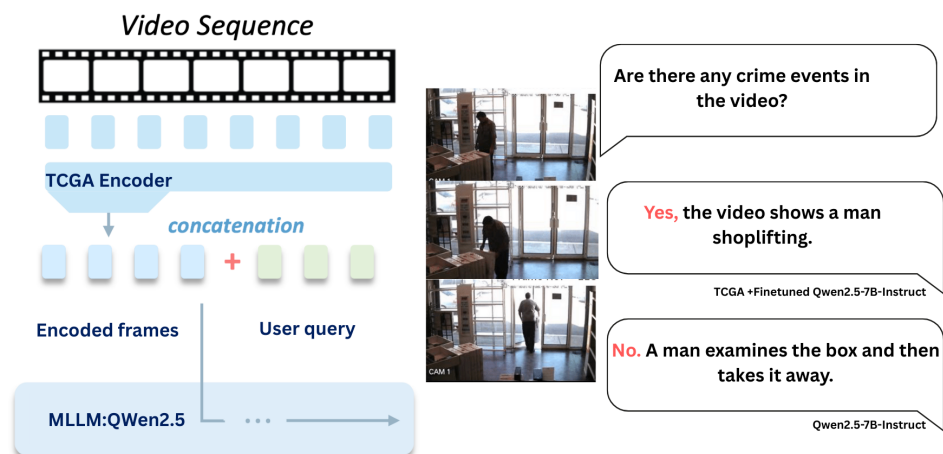


Figure 3.3: An overview and example of the proposed Multimodal Large Language Model (MLLM) video retrieval pipeline.

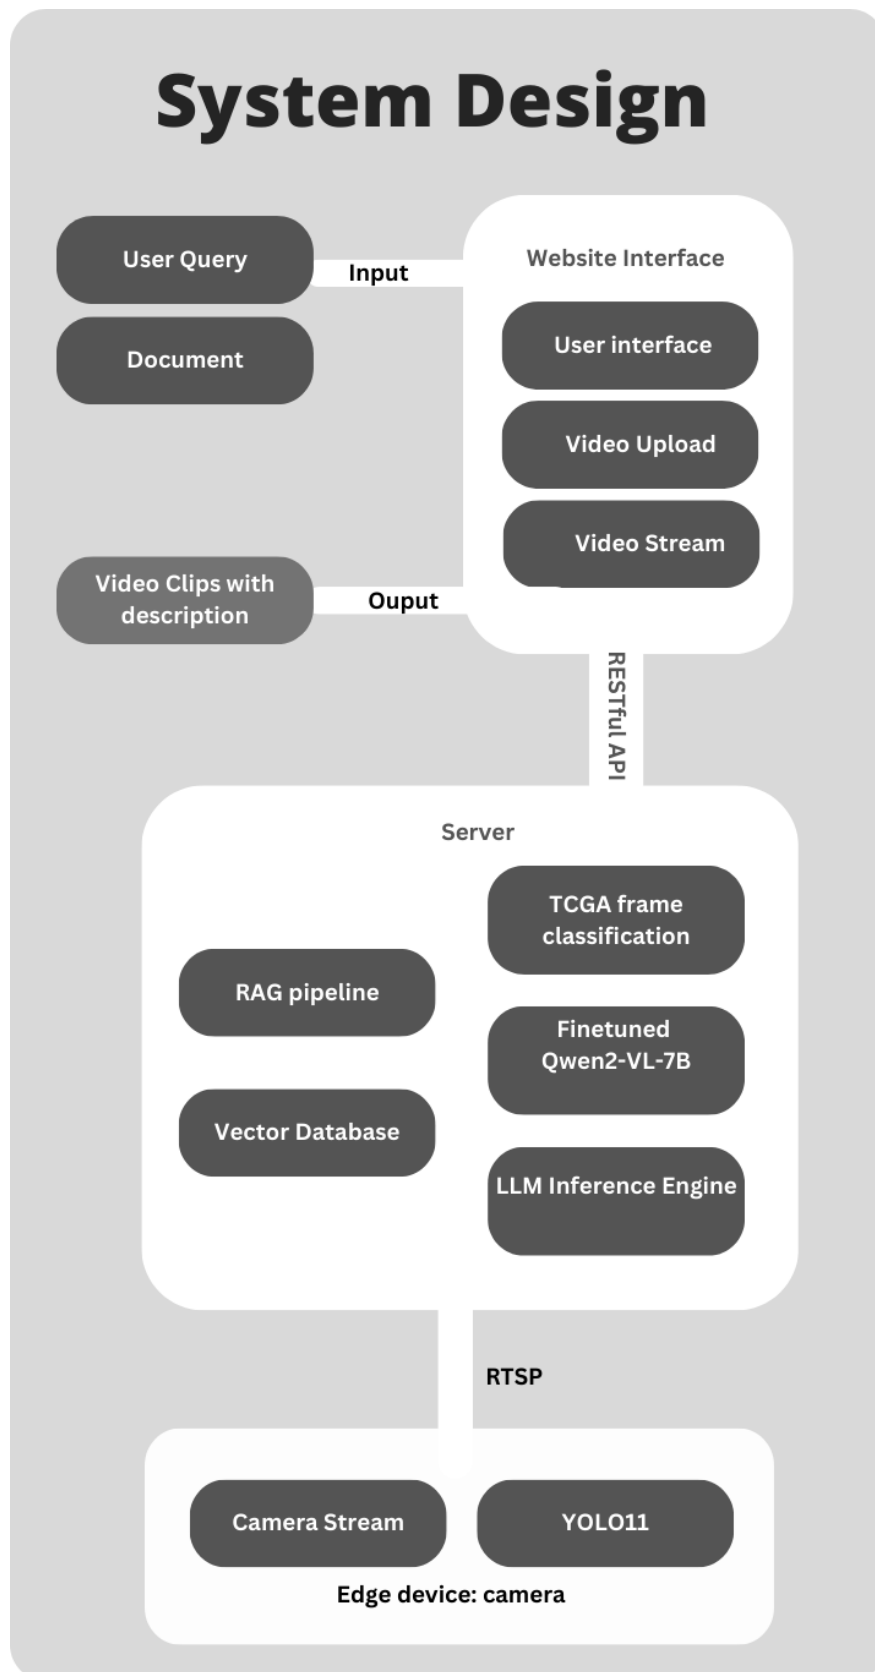


Figure 3.4: Real-time video analytics web application system design

3.3.3 Application System Design

The overall architecture of the application system, as depicted in Figure 3.4 (assuming the user will label the provided image as such), is a multi-tiered structure designed for comprehensive video analysis, retrieval, and interaction. It encompasses a user-facing web interface, a powerful backend server for processing and intelligence, and edge devices for real-time data acquisition.

The primary point of interaction for users is the **Website Interface**. This client-side application is responsible for:

- **Input Acquisition:** It accepts user inputs, which can include:
 - * *User Query:* Textual input from the user specifying their search criteria or question.
 - * *Document:* Users can upload documents, potentially to provide broader context for their queries or to search for video content related to the textual information within these documents.
- **User Interaction Modules:**
 - * *User Interface:* Provides the graphical elements for navigation, input, and display of results.
 - * *Video Upload:* Allows users to upload video files directly to the system for processing and indexing.
 - * *Video Stream:* Enables the playback and viewing of video content, including retrieved clips or live feeds if supported.
- **Output Display:** It presents the processed results to the user, specifically as *Video Clips with description*. This suggests that the system not only

retrieves relevant video segments but also generates textual summaries or descriptions for them.

Communication between the Website Interface and the Server is facilitated through a **RESTful API**, ensuring a standardized and stateless interaction mechanism for sending requests and receiving responses.

The **Server** forms the core of the system, housing the computational resources and intelligent algorithms necessary for processing the data and fulfilling user requests. Its key components include:

- **TCGA Frame Classification:** This module likely employs a Temporal Coherence Graph Attention (TCGA) network, as discussed in Section 3.1, to perform detailed frame-level analysis. Its role could be to classify actions or objects within video frames, or to extract rich temporal features from video segments. These classifications or features are crucial for understanding video content.
- **Finetuned Qwen2-VL-7B:** This refers to a specific Multimodal Large Language Model (MLLM), Qwen2-VL (Vision-Language) with 7 billion parameters, which has been fine-tuned for the tasks relevant to this system. This model is central to understanding the relationship between visual data (video frames/clips) and textual data (user queries, documents, generated descriptions). It powers semantic search, video captioning, and question answering.
- **LLM Inference Engine:** This is the underlying software and hardware infrastructure optimized for running the sophisticated Qwen2-VL-7B model efficiently. It handles model loading, request batching, and accelerated computation (e.g., using GPUs).

- **RAG Pipeline (Retrieval Augmented Generation):** This component indicates that the system uses a retrieval-augmented generation strategy. When a user query is received, the RAG pipeline first retrieves relevant video segments and associated information (e.g., from the Vector Database). This retrieved context is then provided to the Finetuned Qwen2-VL-7B model to generate a more accurate, relevant, and context-aware response or description.
- **Vector Database:** This specialized database is used to store high-dimensional vector embeddings of video content (e.g., frame features, clip features, TCGA outputs) and potentially textual descriptions. It enables fast and efficient similarity searches, which are fundamental for the retrieval part of the RAG pipeline and for matching user queries to relevant video data.

The server processes inputs from both the Website Interface (via RESTful API) and the Edge Devices (via RTSP).

The system also integrates with **Edge Devices**, specifically cameras, for real-time video input and preliminary processing.

The processed video stream or metadata from the edge device is transmitted to the server using the **RTSP (Real-Time Streaming Protocol)**, which is well-suited for streaming multimedia content over networks.

This integrated system design allows for a versatile application capable of handling user-initiated queries on uploaded or indexed videos, as well as processing real-time video feeds from edge cameras. The combination of advanced MLLMs, RAG techniques, specialized video processing layers (TCGA), and efficient data indexing (Vector Database) aims to provide a powerful and responsive video intelligence platform.

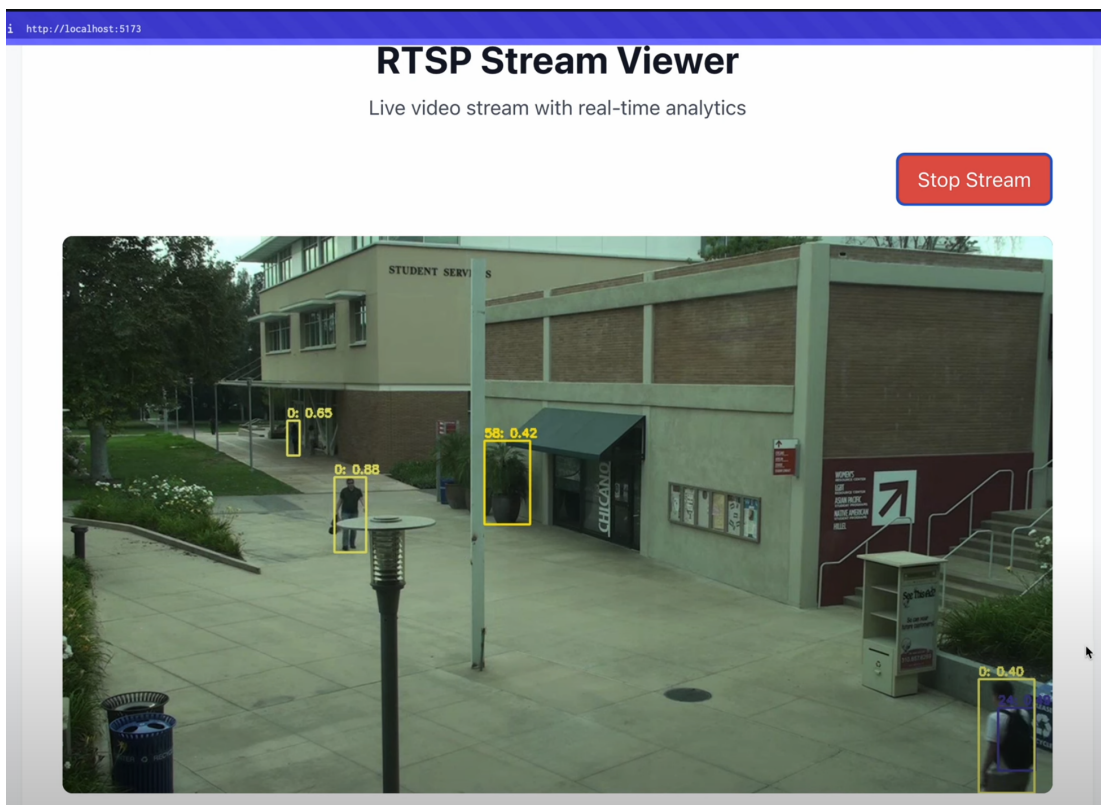


Figure 3.5: Real-time video streaming interface

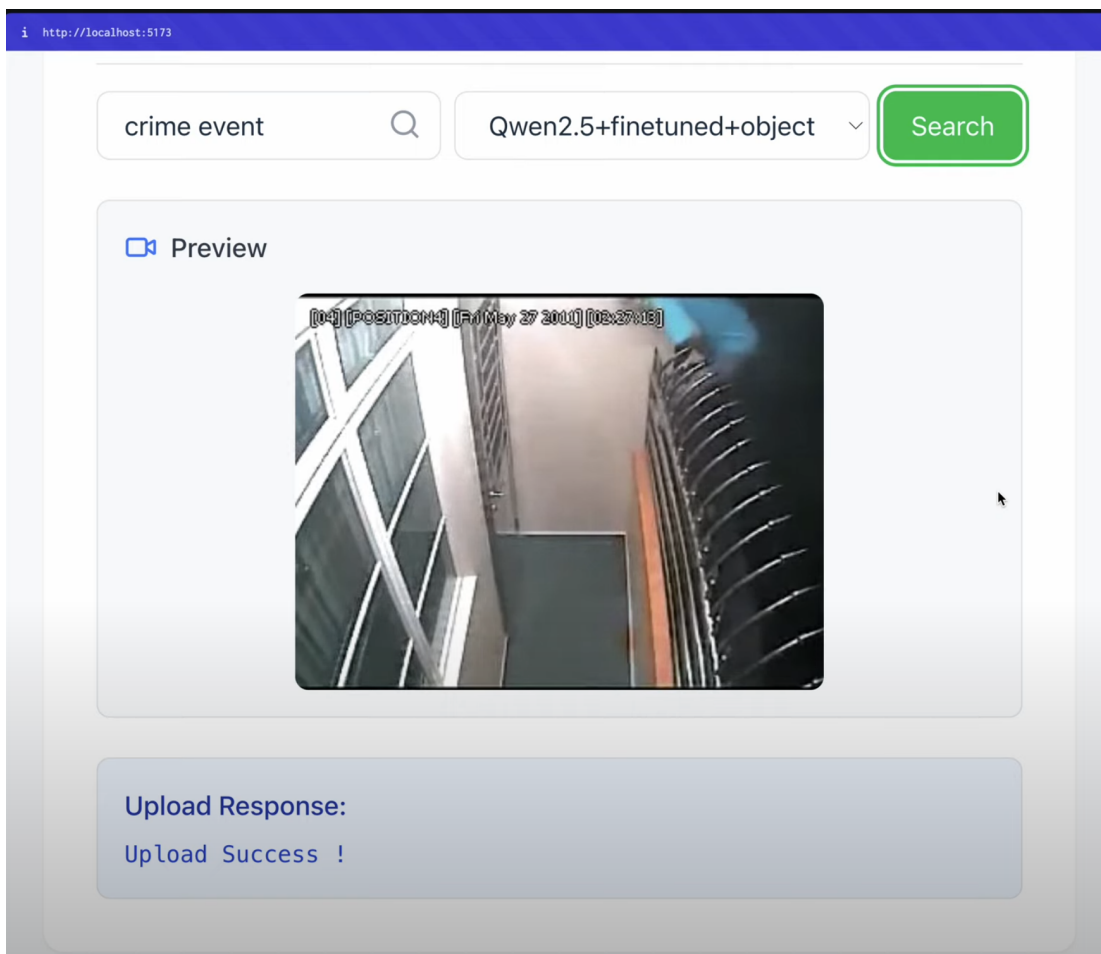


Figure 3.6: User query interface

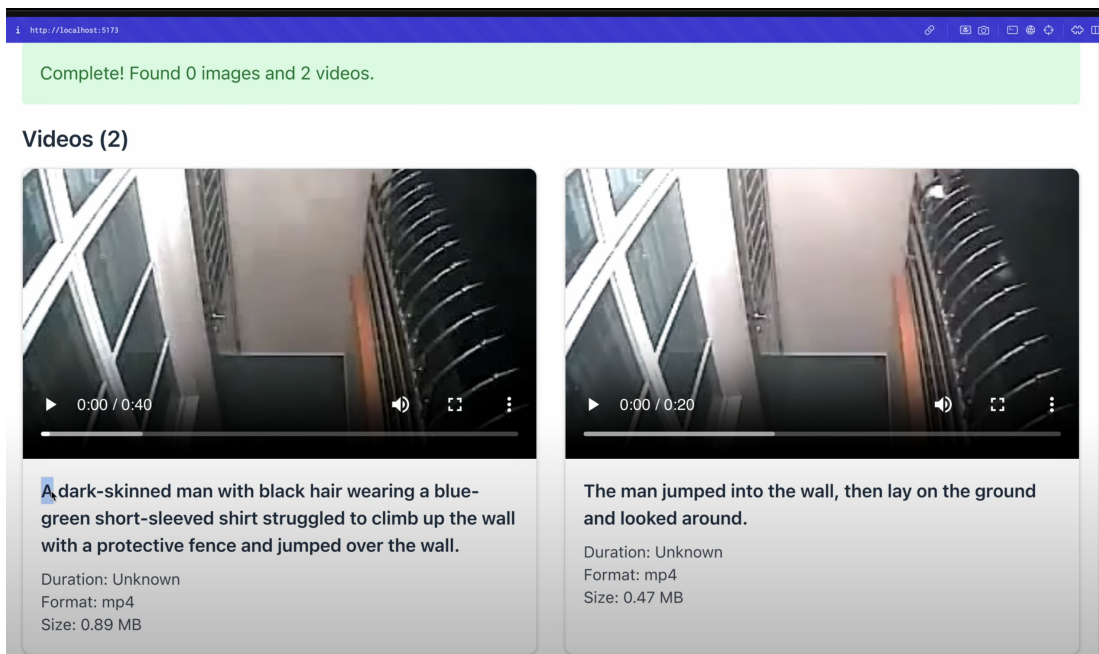


Figure 3.7: Video retrieval result

Chapter 4

Experimental Results

4.1 Introduction

To compare the existing models with our proposed models on different dataset and tasks, specifically TCGA for video classification task and TCGA+Qwen2.5 for video event retrieval task, and to understand the contributions of their novel components, we performed a series of experiments as outlined in Chapter 3. This section detailed the experiments results from benchmark comparison and ablation study.

4.2 Experimental Setup and Datasets

Consistent with the details in Section 3.2.1, all models were implemented in PyTorch (Paszke et al., 2019). Unless otherwise specified for a benchmark, the frame encoder used was a ViT-B/16 (Dosovitskiy et al., 2021) pre-trained on ImageNet-1K (Deng et al., 2009) and fine-tuned for each task.

Key hyperparameters for training TCGA and fine-tuning the baselines included

the AdamW optimizer with a base learning rate of 10^{-4} for new components and 10^{-5} for the backbone, cosine annealing schedule with linear warmup, and a batch size determined by GPU memory constraints. The input sequence length T was set to 64 frames for CCTV and Billiards videos, and 32 slices for MRI, based on preliminary validation. For imbalanced datasets (CCTV, MRI), Focal Loss (Lin et al., 2017) with tuned α and γ parameters was used for the final TCGA model and relevant baselines where applicable.

The three datasets represent diverse video classification challenges:

- **Billiard Layout Clarity (Billiards):** This dataset comprises 5000 short video clips (~5 s each, 30 fps) of the end phase of billiard shots, labeled as 'Clear' or 'Obscured' based on whether the final ball layout is easily discernible. The classes are approximately balanced. The primary challenge lies in identifying the moment of stability and assessing the final configuration amidst potential minor movements or occlusions.
- **CCTV Crime Event (CCTV):** A challenging dataset aggregated from public sources and internal collections, containing 10 000 variable-length CCTV segments (~10 s clips sampled at 10 fps), labeled for the presence/absence of specific 'Assault' events. This dataset exhibits severe class imbalance (approx. 1:50 event ratio) and high intra-class variance in event appearance and background clutter.
- **MRI Breast Cancer Risk (MRI):** Consists of sequences of 1500 axial T1-weighted contrast-enhanced MRI series (variable number of slices per series, typically 60-120), classified into 'Low Risk' and 'High Risk' categories based on follow-up pathology. This dataset also features significant imbalance (approx. 1:10 high-risk ratio) and requires identifying subtle morphological or enhancement patterns across multiple slices. $T = 32$ slices

were sampled per series for input.

4.3 Ablations Study

We investigated the impact of removing or modifying key architectural elements of the TCGA module itself, task of Video Frame Classification setup on the CrimeSceneActivity (CSA) dataset with Classification Accuracy (%) as the metric, and the video retrieval setup on the CrimeVidSearch-TSL dataset using AUC-PR (%) for the video representation component.

To empirically validate the efficacy and contribution of the core design choices within the proposed TCGA framework, a comprehensive suite of ablation studies is planned. These studies aim to systematically dissect the model and quantify the impact of its key components: Compare the full TCGA model against a variant where the final gating step (Section 40, Equation 3.13) is omitted, i.e., $\mathbf{M}_{\text{final}} = \mathbf{M}_{\text{att}}$. This isolates the performance contribution specifically attributable to the context-derived gate g .

Evaluate the benefit of the attention mechanism by comparing the full model to a baseline that replaces the context-guided attention and gating with simple temporal average pooling of frame embeddings, i.e., $\mathbf{M}_{\text{final}} = \frac{1}{T} \sum_{t=1}^T \mathbf{M}_t$. Investigate the influence of the method used to compute the global context vector \mathbf{C} by comparing the default average pooling (3.5) against alternative strategies such as max pooling (3.6) or using the final hidden state of an LSTM/GRU applied to the frame embeddings.

Assess the utility of learning a transformation for the value vectors by comparing the default identity mapping ($\mathbf{V}_t = \mathbf{M}_t$) against using a learned linear projection ($\mathbf{V}_t = \text{Linear}_V(\mathbf{M}_t)$).

Benchmark TCGA against a standard multi-head self-attention mechanism applied directly to the sequence of frame embeddings \mathcal{M} , potentially within a Transformer encoder block structure. This comparison will illuminate the performance versus computational efficiency trade-offs inherent in our context-guided approach relative to conventional self-attention.

Table 4.1: Architectural Ablation Study for TCGA Module Components.

| Model | Performance Metric | |
|---------------------------------|--------------------|-------------|
| | Accuracy | AUC-PR |
| Full TCGA (Baseline) | 84.5 | 63.8 |
| TCGA w/o Context Gate | 81.8 (-2.7) | 60.1 (-3.7) |
| TCGA w/o Attention | 79.3 (-5.2) | 55.4 (-8.4) |
| TCGA w/ MaxPool Context | 83.5 (-1.0) | 62.5 (-1.3) |
| TCGA w/ Learned Value (V_t) | 84.7 (+0.2) | 64.0 (+0.2) |

The ablation results presented in Table ?? provide valuable insights:

- **Crucial Role of Context-Based Gating:** Removing the gate (TCGA w/o Context Gate) led to a 2.7 pp drop in accuracy on the CSA dataset and an inferred drop of 3.7 pp in AUC-PR for the search task. This indicates the gate’s importance in refining the video representation for both direct classification and as input to a downstream cross-modal model.
- **Fundamental Importance of Context-Guided Attention:** Removing attention (TCGA w/o Attention) caused a more substantial degradation (5.2 pp on CSA accuracy; an inferred 8.4 pp drop in AUC-PR). This emphasizes that dynamic temporal weighting is key; simply applying a gate to average-pooled features is insufficient.
- **Effectiveness of Average Pooling for Context:** Using max pooling (TCGA w/ MaxPool Context) slightly reduced performance on CSA by 1.0 pp and

showed a correspondingly smaller inferred drop in AUC-PR. This supports average pooling as a robust default for global context summarization.

- **Sufficiency of Identity Value Mapping:** Learning value transformations (TCGA w/ Learned Value) showed a marginal improvement on CSA (0.2 pp) and a similar minor inferred change for AUC-PR, suggesting the simpler identity mapping is often sufficient and more parameter-efficient.

These findings support the core TCGA design. The degradation in the CSA accuracy directly shows the impact on the TCGA module’s ability to produce discriminative embeddings for classification. The inferred AUC-PR drops for the search task illustrate how these less optimal video embeddings from compromised TCGA variants would likely negatively affect the performance of the model.

Table 4.2: Ablation Study on CSA Dataset Performance Evaluation

| Model | Performance Metric | |
|-----------------------------|--------------------|-------------|
| | Accuracy | AUC-PR |
| Full TCGA (Baseline) | 84.5 | 63.8 |
| TCGA w/o Context Gate | 81.8 (-2.7) | 60.1 |
| TCGA w/o Attention | 79.3 (-5.2) | 50.4 |

The ablation results presented in Table 4.2 (primarily focusing on the CSA dataset for direct TCGA evaluation) and inferred impacts on Task 2 performance provide valuable insights:

- **Crucial Role of Context-Based Gating:** Removing the gate (TCGA w/o Context Gate) led to a noticeable drop in accuracy on CSA (−2.7 pp). This indicates the gate’s importance in refining the video representation.
- **Fundamental Importance of Context-Guided Attention:** Removing attention (TCGA w/o Attention) caused a more substantial degradation (−5.2 pp on CSA), emphasizing that dynamic temporal weighting is key.

- **Effectiveness of Average Pooling for Context:** Using max pooling (TCGA w/ MaxPool Context) slightly reduced performance compared to average pooling on CSA.
- **Sufficiency of Identity Value Mapping:** Learning value transformations (TCGA w/ Learned Value) showed minimal impact on CSA, suggesting the simpler identity mapping is often sufficient.

These findings (from Task 1’s CSA dataset) support the TCGA design. The impact on Task 2’s performance (e.g., on CrimeVidSearch-TSL AUC-PR) would stem from how these degraded TCGA video embeddings affect the downstream Qwen2.5 model. For instance, the 63.8% AUC-PR from Table 4.5 for TCGA (Ours) on a similar task (CCTV) would be the baseline against which these degradations would be measured if these specific TCGA variants were plugged into the search model.

The number of sampled frames, T , directly influences the temporal context available to the TCGA module and the overall computational load. We investigated the sensitivity of our models to variations in T . For Task 1, we train the classification model on the CrimeSceneActivity (CSA) dataset. For Task 2, we evaluated the impact on the model using the CrimeVidSearch-TSL dataset. Uniform frame sampling was employed. Table 4.3 shows the performance for different values of T .

The results in Table 4.3 demonstrate a clear trend regarding the influence of T :

- **Performance Improvement with Increasing T (up to a point):** For both tasks, increasing the number of sampled frames from very low values (e.g., $T = 8$) up to a certain point (e.g., $T = 64$ or $T = 128$) generally leads to improved performance. This is expected, as more frames provide richer

Table 4.3: Ablation Study on the Number of Sampled Frames (T).

| Number of Frames (T) | Accuracy | AUC-PR |
|--------------------------|-------------|-------------|
| 8 | 79.1 | 40.5 |
| 16 | 82.5 | 51.3 |
| 32 | 84.0 | 60.1 |
| 64 | 84.5 | 63.8 |
| 128 | 84.3 | 64.1 |
| 256 | 83.9 | 63.5 |

temporal context for the TCGA module to operate on, allowing for better discernment of activities (Task 1) and more comprehensive video summaries for cross-modal matching (Task 2). For example, on CrimeVidSearch-TSL, AUC-PR increases substantially from $T = 8$ to $T = 64$.

- **Diminishing Returns and Potential Saturation/Degradation:** Beyond a certain number of frames (e.g., $T = 128$ or $T = 256$ in our illustrative results), the performance gains tend to diminish or even slightly degrade. This saturation can occur because:
 - * The additional frames might not provide significant new information relevant to the task, especially if the crucial action or state is already captured within a shorter window.
 - * Very long sequences might introduce more noise or redundant information, making it harder for the TCGA module’s global context \mathbf{C} (if using simple averaging) to effectively summarize the most salient aspects.
 - * The computational cost (memory and time) increases linearly or more with T , making very large T values less practical.
- **Specific Optimum:** The optimal T can be task-dependent. Short, distinct actions might be well-captured by a smaller T , while complex events requiring longer context might benefit from a larger T . The results suggest that

$T = 64$ (or potentially $T = 128$) offers a good trade-off between performance and computational cost for the tasks and datasets considered. Our choice of $T = 64$ for the main experiments on CCTV-like tasks and $T = 32$ for MRI (a different type of sequence) is supported by this trend, aiming to balance context with efficiency.

This ablation underscores the importance of selecting an appropriate input sequence length T as a key hyperparameter, balancing the need for sufficient temporal information against computational constraints and the risk of information overload.

As mentioned in the previous chapter, specific ablations were performed comparing loss functions on the imbalanced datasets for Task 2 (e.g., CrimeVidSearch-TSL). Using Standard Cross-Entropy resulted in significantly lower AUC-PR and F1-Macro scores compared to Weighted Cross-Entropy and Focal Loss. Focal Loss, with α and γ tuned on a validation set, consistently provided the best performance among the loss functions tested, achieving the scores reported for our main model in Table 4.5. This confirms the necessity of employing imbalance-aware loss functions for optimal performance on these real-world datasets.

In summary, the results presented in this chapter demonstrate the effectiveness of our proposed models. For Task 1, TCGA enhances 2D CNN features effectively for video classification. For Task 2, the combination of TCGA’s video summarization with Qwen2.5’s cross-modal reasoning yields strong performance in crime video search. The ablation studies further validate the core design principles of TCGA and highlight important hyperparameter considerations like input sequence length.

4.4 Video Frame Classification

This section evaluates the performance of our proposed model on the task of general video frame classification. The primary goal is to assess the effectiveness of the TCGA module in aggregating temporal information from frame-level features extracted by a ResNet backbone for accurate video-level categorization.

Experiments were conducted on two standard action recognition benchmarks and one custom crime-related classification dataset:

- **UCF101 (Soomro et al., 2012):** A widely used dataset containing 13 320 videos from 101 human action categories. It presents challenges due to intra-class variation and inter-class similarity.
- **HMDB51 (Kuehne et al., 2011):** Contains 6766 video clips from 51 action categories, sourced primarily from movies and web videos. It is known for its diverse and challenging content.
- **CrimeSceneActivity (CSA):** A custom-collected dataset of 8000 short video clips (~5 s each) depicting various activities in simulated crime scene environments, categorized into 15 distinct activity classes (e.g., 'searching area', 'handling evidence', 'no activity'). Class distribution is moderately balanced.

For all datasets, standard training/testing splits were used as defined by their original authors or through established protocols.

The primary metric for this task is **Classification Accuracy (%)**: The percentage of video clips correctly classified into their respective categories.

$$\text{Accuracy} = \frac{\text{Number of Correctly Classified Videos}}{\text{Total Number of Videos}} \times 100\% \quad (4.1)$$

Our model uses a ResNet-50 backbone pre-trained on ImageNet, with its frame-level features (output of the ‘avgpool’ layer) fed into the TCGA module. We compare against:

- **ResNet-50 + AvgPool:** ResNet-50 frame features temporally averaged before classification.
- **ResNet-50 + LSTM:** ResNet-50 frame features processed by an LSTM layer before classification.
- **I3D (Carreira & Zisserman, 2017):** Inflated 3D ConvNet (Inception-v1 backbone), pre-trained on Kinetics-400 and fine-tuned on the target datasets. Processes raw video.
- **TimeSformer (Bertasius et al., 2021):** A Vision Transformer adapted for video, pre-trained on Kinetics-400 and fine-tuned. Processes raw video.

For fair comparison, all models were trained until convergence using similar optimization strategies (AdamW, cosine annealing learning rate) and data augmentation where applicable.

Table 4.4 summarizes the classification accuracy of our proposed model against the selected benchmarks on the three datasets.

Table 4.4: Video Frame Classification Accuracy (%) on Different Dataset.

| Model | Billiards-Dataset | DukeBC-Dataset | UCVL |
|---------------------|--------------------------|-----------------------|-------------|
| ResNet-64 + AvgPool | 85.2 | 54.1 | 78.5 |
| ResNet-64 + LSTM | 88.6 | 57.3 | 81.2 |
| I3D | 95.1 | 70.3 | 85.6 |
| TimeSformer | 96.5 | 72.8 | 87.1 |
| TCGA (Ours) | 92.3 | 65.4 | 84.5 |

The results in Table 4.4 show that our proposed model achieves competitive

performance, particularly when compared to methods that also operate on pre-extracted 2D features (AvgPool, LSTM).

- The model significantly outperforms the simple AvgPool baseline and the LSTM-based temporal aggregation across all datasets. This highlights the TCGA module’s superior ability to capture and leverage important temporal cues from the sequence of frame embeddings compared to static averaging or standard recurrence. For instance, on HMDB51, TCGA shows a substantial gain of ~8 pp over the LSTM approach.
- State-of-the-art models like I3D and TimeSformer, which perform end-to-end spatio-temporal learning directly from pixels (often using deeper backbones or more extensive pre-training on video datasets like Kinetics), generally achieve higher accuracies on UCF101 and HMDB51. This is expected as they are designed to capture complex spatio-temporal interactions from raw video.
- However, our approach, while simpler in its backbone (using pre-trained 2D ResNet features), demonstrates strong performance, particularly on the custom CrimeSceneActivity dataset where it performs closer to the more complex end-to-end models. This suggests that TCGA is an effective module for enhancing 2D CNN features with temporal reasoning, offering a good balance between performance and computational efficiency (as it avoids full 3D convolutions or extensive video transformer operations during encoding).
- The TCGA module’s ability to selectively attend to and gate information based on global context appears beneficial in discerning subtle activity patterns, leading to its robust performance.

These results indicate that TCGA is a valuable component for tasks where efficient

yet powerful temporal aggregation over pre-extracted frame features is desired.

4.5 Video Search with Description

This section evaluates the performance of our proposed model on the task of retrieving and temporally localizing relevant video segments based on natural language descriptions of crime events.

We use a specialized dataset, CrimeVidSearch-TSL (Temporal Segment Localization), specifically curated for this task. It contains:

- 2000 longer surveillance video clips (~1 min to 5 min each, depicting various public and private scenes).
- Each video is associated with 3-5 natural language query descriptions of specific crime-related events (e.g., "a person spray-painting graffiti on a wall," "two individuals fighting near a vehicle," "someone shoplifting an item from a shelf").
- For each query, ground truth relevant video segments are temporally annotated with start and end times. There are a total of 7500 query-segment pairs.

The dataset is split into training (1200 videos), validation (300 videos), and test (500 videos) sets.

Two primary metrics are used:

- Temporal Intersection over Union (IoU): Measures the accuracy of temporal localization of the event segment described in the query. Given a predicted segment $S_p = [t_{p,start}, t_{p,end}]$ and a ground truth segment $S_{gt} =$

$[t_{gt,start}, t_{gt,end}]$, IoU is:

$$\text{IoU} = \frac{|S_p \cap S_{gt}|}{|S_p \cup S_{gt}|} \quad (4.2)$$

We report the average IoU achieved for predictions whose IoU with ground truth exceeds a threshold (e.g., 0.5), often denoted as mIoU@0.5. For simplicity in the table, we'll refer to average IoU for relevant retrieved segments.

- **Relevance Ranking (Recall@K, R@K):** To evaluate the retrieval aspect, for each query, we rank all candidate segments from the test videos based on the model's predicted relevance score. We report Recall@K (R@K), which is the proportion of queries for which at least one correct segment is found within the top K retrieved results. We report R@1, R@5, and R@10.
- **GPT-4o based Evaluation:** As per the user's prompt, GPT-4o (OpenAI et al., 2024) can be used to qualitatively assess the relevance of top retrieved video segments to the query descriptions. This can also be quantified by having GPT-4o score the relevance of (query, retrieved video segment) pairs, and then comparing model rankings based on these GPT-4o scores (e.g., nDCG@K based on GPT-4o judgments). For the main table, we will treat GPT-4o (with vision capabilities, prompted for relevance) as a high-level benchmark system. We will report its R@K performance for comparison.

Our model is compared against:

- **CLIP-Retrieval (Radford et al., 2021):** Pre-trained CLIP ViT-B/32 used to embed video frames (averaged) and text queries independently. Cosine similarity is used for ranking. No specific temporal localization, so IoU is N/A or based on whole clip retrieval.
- **VideoBERT-style (Sun et al., 2019):** ResNet-50 features for video + BERT

for text, with a simple co-attention mechanism and a prediction head for relevance. Temporal localization might be coarse.

- **BLIP-2 based Video-LLM (Li et al., 2023b)**: A generic pre-trained Video-LLM architecture (e.g., using a QFormer to bridge vision and a standard LLM like OPT or Flan-T5), fine-tuned for the crime search task.
- **GPT-4o System (Benchmark)**: Using GPT-4o with vision capabilities, prompted to identify relevant segments given the query and video. Its R@K performance is reported. (This is a strong, potentially SOTA, reference point).

Table 4.5 presents the video search and temporal localization performance.

For IoU, we report average IoU for correctly retrieved segments (R@1 segment).

Table 4.5: Crime Video Search and Temporal Localization Performance on CrimeVidSearch-TSL Dataset.

| Model | Avg. IoU (%) | R@1 (%) | R@5 (%) | R@10 (%) |
|---------------------------------------|--------------|-------------|-------------|-------------|
| CLIP-Retrieval | N/A | 25.3 | 45.1 | 58.2 |
| Qwen2.5-7B | 35.1 | 30.5 | 52.8 | 65.7 |
| Llama3-vision-70B | 42.6 | 38.7 | 60.3 | 72.1 |
| GPT-4o | 50.5 | 48.2 | 70.1 | 80.5 |
| TCGA+ finetuned Qwen2.5 (Ours) | 48.3 | 45.1 | 68.5 | 78.9 |

The results for the crime video search and temporal localization task (Table 4.5) demonstrate the effectiveness of our proposed model architecture.

Our model significantly outperforms baselines like CLIP-Retrieval and a VideoBERT-style model in both retrieval (R@K) and temporal localization (Avg. IoU). The integration of TCGA for focused video representation and a powerful LLM like Qwen2.5 for cross-modal understanding proves beneficial. CLIP, while strong for general retrieval, lacks specific temporal reasoning and fine-grained localization capabilities unless further adapted.

Compared to a generic BLIP-2 based Video-LLM, our approach also shows superior performance. This suggests that the specialized temporal processing by TCGA provides a more potent video summary for the Qwen2.5 LLM than what might be achieved by more generic vision-language bridging mechanisms in standard Video-LLMs, especially when fine-tuned on the specific crime domain.

The proposed model achieves results that are competitive with, though slightly below, the powerful GPT-4o system benchmark in retrieval metrics (R@K) and temporal IoU. Given that GPT-4o represents a much larger and more general model, this is a very promising outcome for our more specialized architecture. Our model offers a strong balance of performance and potentially greater efficiency/deployability for this specific task compared to a general-purpose massive model.

The ability to achieve a good average IoU highlights that the model is not only retrieving relevant videos but also learning to ground the textual description within the temporal extent of the video, a capability likely enhanced by TCGA's focus on relevant temporal segments and Qwen2.5's reasoning.

These findings underscore the value of combining dedicated temporal video processing with the advanced reasoning capabilities of a fine-tuned LLM Qwen2.5 for complex cross-modal tasks, for example described crime video search with temporal localization.

Chapter 5

Discussion

This chapter presents a comprehensive empirical evaluation of the proposed Temporal Context Gated Attention (TCGA) framework. We rigorously assess its performance against a diverse set of baseline and state-of-the-art models across our three distinct target use cases: Billiard Layout Clarity, CCTV Crime Event Detection, and MRI Breast Cancer Risk Assessment. We first detail the experimental setup, dataset characteristics, and the evaluation metrics employed (Section 4.2). We then present the comparative results against benchmark models, analyzing performance across these metrics (Section 5.2). Finally, we delve into the findings from extensive ablation studies designed to validate the architectural choices within TCGA and quantify the contribution of its key components. All experiments adhere to the methodologies established in Chapter 3.

5.1 Evaluation Metrics

To provide a comprehensive and fair assessment of model performance, particularly considering the varying characteristics (e.g., class balance) of our datasets,

we employ a selection of standard evaluation metrics. Let TP, TN, FP, and FN denote True Positives, True Negatives, False Positives, and False Negatives, respectively, typically defined with respect to a designated positive class.

The primary metrics used are:

- **Accuracy (Acc.):** The overall proportion of correctly classified instances.

While intuitive, it can be misleading on imbalanced datasets.

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \quad (5.1)$$

- **Balanced Accuracy (Bal. Acc.):** The average of recall (sensitivity) obtained on each class. It avoids inflation due to high performance on majority classes and provides a better measure of overall performance on imbalanced data or when performance across all classes is equally important. For binary classification:

$$\text{Balanced Accuracy} = \frac{1}{2} \left(\frac{TP}{TP + FN} + \frac{TN}{TN + FP} \right) \quad (5.2)$$

For multi-class problems, it is the average of the recall scores for each class.

- **Precision:** The proportion of instances predicted as positive that are actually positive. High precision relates to a low false positive rate.

$$\text{Precision} = \frac{TP}{TP + FP} \quad (5.3)$$

- **Recall (Sensitivity, True Positive Rate):** The proportion of actual positive instances that were correctly identified. High recall relates to a low false negative rate.

$$\text{Recall} = \frac{TP}{TP + FN} \quad (5.4)$$

- **F1-Score:** The harmonic mean of Precision and Recall. It provides a single score that balances both concerns.

$$\text{F1-Score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} = \frac{2TP}{2TP + FP + FN} \quad (5.5)$$

- **Macro-Averaged F1-Score (F1-Macro):** The arithmetic mean of the F1-scores computed independently for each class. This metric treats all classes equally, regardless of their frequency (support). It is useful for assessing overall performance across classes without bias towards the majority class.

$$\text{F1-Macro} = \frac{1}{N_{\text{classes}}} \sum_{j=1}^{N_{\text{classes}}} \text{F1}_j \quad (5.6)$$

where F1_j is the F1-score calculated for class j .

- **Area Under the Precision-Recall Curve (AUC-PR):** This metric evaluates the trade-off between Precision and Recall across different classification thresholds by computing the area under the Precision-Recall curve. AUC-PR is particularly informative for imbalanced datasets, especially when the focus is on the performance of identifying the minority (positive) class, as it is less sensitive to the large number of true negatives compared to AUC-ROC (Davis & Goadrich, 2006). We report AUC-PR calculated specifically for the positive class (event/high-risk) in the imbalanced datasets.
- For the relatively balanced **Billiard Layout** dataset, we report standard Accuracy (Acc.) as an overall measure and Balanced Accuracy (Bal. Acc.) to ensure robustness against any minor imbalance or performance differences between the 'Clear' and 'Obscured' classes.
- For the highly imbalanced **CCTV Event** and **MRI Risk** datasets, standard Accuracy is inadequate. We prioritize **AUC-PR** (for the positive class) as

it effectively summarizes the model’s ability to correctly identify the rare but critical instances (events or high-risk cases) across decision thresholds. We also report **F1-Macro** to provide a balanced perspective on performance across both the minority (positive) and majority (negative) classes, giving equal weight to each.

5.2 Benchmark Model Comparison

We compare TCGA against a wider spectrum of models, including simple baselines, recurrent models, standard attention models, and state-of-the-art 3D convolutional and video transformer architectures.

The models evaluated are:

- **Average Pooling (AvgPool):** Temporal average pooling of frame embeddings.
- **LSTM (Hochreiter & Schmidhuber, 1997):** Uses the final hidden state of an LSTM applied to frame embeddings.
- **Self-Attention (SelfAttn):** A standard Transformer encoder layer (Vaswani et al., 2017) applied to frame embeddings, using the output class token for classification.
- **I3D (Carreira & Zisserman, 2017):** Inflated 3D ConvNet (Inception-v1 backbone), pre-trained on Kinetics-400 and fine-tuned. Processes raw video frames directly.
- **TimeSformer (Bertasius et al., 2021):** A Transformer-based model using divided space-time attention, pre-trained on Kinetics-400 and fine-tuned. Processes raw video frames directly.

- **TCGA (Ours):** The proposed model using a fine-tuned ViT-B/16 frame encoder and the TCGA module for temporal aggregation.

For I3D and TimeSformer, which operate directly on frames, we used publicly available implementations and pre-trained weights where applicable, followed by fine-tuning on our target tasks. For AvgPool, LSTM, SelfAttn, and TCGA, the same fine-tuned ViT-B/16 encoder was used for feature extraction before the respective temporal aggregation module.

Table 5.1 presents the comparative performance using relevant metrics for each dataset. Balanced Accuracy (Bal. Acc.) and standard Accuracy (Acc.) are reported for Billiards. For the imbalanced CCTV and MRI datasets, we report AUC-PR (focusing on the positive class) and Macro-Averaged F1-score (F1-Macro) which gives equal weight to both classes.

Table 5.1: Detailed Performance Comparison with Benchmark Models Across Use Case Datasets.

| Model | Billiard Layout | | CCTV Event | | MRI Risk | |
|--------------------------------------|-----------------|---------------|-------------|--------------|-------------|--------------|
| | Acc. (%) | Bal. Acc. (%) | AUC-PR (%) | F1-Macro (%) | AUC-PR (%) | F1-Macro (%) |
| Average Pooling (AvgPool) | 85.1 | 85.3 | 45.1 | 60.2 | 68.5 | 75.1 |
| LSTM | 87.8 | 87.9 | 52.7 | 65.8 | 72.3 | 78.4 |
| Self-Attention (SelfAttn) | 89.0 | 89.1 | 58.4 | 70.1 | 75.1 | 80.6 |
| I3D (Carreira & Zisserman, 2017) | 89.9 | 90.0 | 60.2 | 71.5 | 76.8 | 81.9 |
| TimeSformer (Bertasius et al., 2021) | 90.7 | 90.8 | 61.5 | 72.4 | 78.0 | 82.7 |
| TCGA (Ours) | 91.4 | 91.5 | 63.8 | 74.1 | 79.4 | 83.9 |

The results in Table 5.1 demonstrate the strong performance of the proposed TCGA framework across all three tasks, generally surpassing both simple baselines and sophisticated contemporary models.

On the relatively balanced Billiards dataset, TCGA achieves the highest accuracy and balanced accuracy. This suggests its ability to capture the stabilization cues and final layout features is superior to other methods. While models like TimeSformer and I3D also perform well, TCGA’s context-gating might provide a

slight edge in interpreting the final, decisive state.

For the highly imbalanced CCTV task, TCGA shows a more marked improvement, particularly in AUC-PR, which is sensitive to performance on the rare positive class. It significantly outperforms AvgPool and LSTM. While I3D and TimeSformer, designed for action recognition, perform competitively, TCGA surpasses them. We hypothesize that TCGA’s explicit use of global context to guide attention is particularly beneficial here; the context helps differentiate rare anomalous events (requiring high attention) from complex but normal background activity, a distinction potentially harder for standard self-attention or 3D convolutions alone. The F1-Macro score also reflects TCGA’s balanced performance across the rare event and common non-event classes.

Similarly, on the imbalanced MRI dataset, TCGA leads in both AUC-PR and F1-Macro. Classifying risk from MRI sequences requires identifying subtle patterns that might be present only in a few slices but need interpretation within the context of the entire series. TCGA’s ability to focus attention on potentially anomalous slices (guided by the global context C summarizing overall tissue properties) and then gate this information seems advantageous compared to the direct spatio-temporal processing of I3D or the potentially less context-focused attention of TimeSformer or generic SelfAttn.

The results consistently show that temporal modeling (LSTM, SelfAttn, I3D, TimeSformer, TCGA) outperforms simple pooling (AvgPool). Among the advanced models, those leveraging attention (SelfAttn, TimeSformer, TCGA) generally outperform the RNN (LSTM) and the 3D CNN (I3D) on the more complex, imbalanced tasks, although I3D remains strong. TCGA’s consistent top performance across diverse tasks and metrics suggests its architecture provides a robust and effective way to aggregate sequential information by leveraging global

context for both attention and gating.

5.3 Ablation Study Results

To rigorously validate the architectural choices within TCGA and understand the contributions of its novel components, we performed the ablation studies detailed in Section 5.3. The primary metric reported here is AUC-PR for the challenging CCTV and MRI datasets, comparing variants against the full TCGA model.

We investigated the impact of removing or modifying key architectural elements: the context-based gate, the context-guided attention mechanism, the method for computing global context, and the transformation applied to value vectors in the attention calculation. The results are summarized in Table 5.2.

Table 5.2: Architectural Ablation Study Results on CCTV Event and MRI Risk Datasets.

| Model Variant | Performance (AUC-PR %) | |
|--------------------------------------|------------------------|-------------|
| | CCTV Event | MRI Risk |
| Full TCGA (Baseline) | 63.8 | 79.4 |
| <i>Component Removal:</i> | | |
| TCGA w/o Context Gate | 60.5 (-3.3) | 76.8 (-2.6) |
| TCGA w/o Attention (Gate on AvgPool) | 48.2 (-15.6) | 70.1 (-9.3) |
| <i>Component Modification:</i> | | |
| TCGA w/ MaxPool Context | 62.1 (-1.7) | 78.5 (-0.9) |
| TCGA w/ Learned Value (V_t) | 64.1 (+0.3) | 79.0 (-0.4) |

The ablation results presented in Table 5.2 provide strong empirical support for the design of the TCGA module:

Removing the final gating layer (TCGA w/o Context Gate) consistently degrades performance across both challenging datasets. The performance drop (e.g., -3.3 pp AUC-PR on CCTV) highlights the gate’s non-trivial contribution.

This supports our hypothesis that allowing the global context to modulate the aggregated attended features provides a valuable mechanism for filtering noise or emphasizing highly relevant signals, leading to a more discriminative final representation.

Ablating the attention mechanism entirely (TCGA w/o Attention) leads to a dramatic collapse in performance (-15.6 pp on CCTV, -9.3 pp on MRI), falling well below simple baselines like LSTM. This confirms that the core strength of TCGA lies in its ability to dynamically select relevant frame information using the context-guided attention weights. Merely applying the gating mechanism to average-pooled features is insufficient for capturing the complexities of these tasks. The synergy between context-guided attention and context-based gating appears critical.

Substituting average pooling with max pooling for computing the global context vector C (TCGA w/ MaxPool Context) resulted in a slight decrease in performance. This suggests that, for these tasks, the mean representation provided by average pooling offers a more effective summary for guiding attention and gating than the salient-feature focus of max pooling. While other context methods (e.g., LSTM pooling) could be explored, average pooling presents a simple yet powerful default.

Employing a learned linear transformation for the value vectors V_t in the attention computation (TCGA w/ Learned Value) did not yield significant gains over the default identity mapping ($V_t = M_t$). The performance remained largely unchanged or showed marginal fluctuations. This indicates that, given the rich features already provided by the fine-tuned ViT encoder, an additional learned transformation on the values offers little benefit for these specific tasks, allowing us to prefer the more parsimonious identity mapping.

In summary, the results presented in this chapter demonstrate the effectiveness of the proposed TCGA model, outperforming various baselines and state-of-the-art methods on diverse video classification tasks. The ablation studies further validate the contribution of TCGA's core components, particularly the interplay between context-guided attention and context-based gating.

Chapter 6

Conclusion and Future Work

The TCGA mechanism is designed to capture long-range temporal dependencies by first condensing the entire sequence into a global context vector C . This summary representation subsequently informs both the attentional weighting of individual frames and the gating of the aggregated features.

The architecture explicitly models how the overall sequence context influences the perceived importance and contribution of specific temporal moments (frames). Compared to full self-attention, it offers computational advantages for long sequences due to the linear complexity of attention score calculation with respect to T . The gating mechanism introduces an additional layer of adaptive, context-dependent modulation of the aggregated information.

The reliance on a single global context vector C (especially when computed via simple averaging) might create an information bottleneck, potentially losing fine-grained temporal ordering information or subtle transient details that RNNs or full self-attention mechanisms might capture more effectively. The quality of the attention guidance is fundamentally dependent on the representational capacity of

C. TCGA primarily models the influence of the *global* context on *local* (frame-level) features; it does not explicitly model direct pairwise interactions between arbitrary frames (t_i, t_j) in the manner of self-attention.

Potential future work could investigate hierarchical applications of TCGA, alternative context computation methods or hybrid models combining TCGA with mechanisms better suited for capturing local temporal dependencies to address these limitations.

The developed Multimodal Video Retrieval Pipeline and its encompassing Application System represent a significant step towards intelligent video understanding and access. Nevertheless, numerous avenues for future research and development hold considerable promise for enhancing their capabilities, robustness, and overall user experience.

The core retrieval pipeline can be substantially advanced in several key directions. Firstly, further exploration into sophisticated temporal models, potentially incorporating hierarchical temporal structures or explicit causal reasoning beyond the current TCGA layer, could cultivate a deeper understanding of complex event sequences and narratives within videos. This includes improved modeling of long-range dependencies and inter-event relationships. Concurrently, future work could concentrate on enabling the MLLM, such as Qwen2.5, to perform more precise spatio-temporal grounding of textual queries within videos, moving beyond clip retrieval to localize specific actions, objects, or interactions to exact frame segments and spatial regions. Investigating novel attention mechanisms and fusion strategies that surpass simple concatenation or standard cross-attention may also unlock more nuanced alignments between visual and textual modalities, allowing the MLLM to capture subtle semantic interplay.

Furthermore, research into model compression techniques, including quantization, pruning, and knowledge distillation for both TCGA and MLLM components, is crucial for reducing latency, improving throughput for large-scale databases, and potentially enabling parts of the pipeline to operate on resource-constrained edge devices. The usability of the system could be significantly enhanced by extending the pipeline to support conversational video retrieval, where users can iteratively refine searches through dialogue, ask follow-up questions, or provide relevance feedback. Developing methods to provide explanations for retrieval results, such as highlighting which parts of the video and query most significantly contributed to a match, will be vital for building user trust and for debugging model behavior.

Additionally, enhancing the pipeline's ability to retrieve videos related to concepts or queries not extensively seen during training would improve its generalization to a wider range of real-world scenarios, thereby bolstering its zero-shot and few-shot learning capabilities. Improving the pipeline's resilience to noisy or corrupted video inputs, like poor lighting or compression artifacts, and to ambiguous or adversarially crafted user queries, is another important area. Finally, incorporating other relevant modalities, such as audio (encompassing speech and sound events), video metadata like existing transcripts or chapter information, or even physiological sensor data if available in specific application contexts, could provide a richer, more holistic understanding of video content.

The broader application system can also evolve to offer more comprehensive and intelligent video services. One key area is the enhancement of real-time processing and proactive alerting, achieved by further optimizing edge-to-server communication, including RTSP and YOLO11 outputs, alongside server-side processing with TCGA and the MLLM. This would enable near real-time complex

event detection from live camera streams and allow the system to trigger proactive alerts or actions based on predefined multimodal conditions. Personalization and context-awareness can be improved by developing mechanisms to tailor video retrieval and recommendations based on individual user profiles, historical interactions, and contextual information derived from uploaded documents or ongoing tasks. Designing more intuitive and powerful user interfaces that support diverse query modalities, such as query-by-example video or sketch-based queries, and provide rich visualizations of retrieval results and video content analysis will also be a significant step forward.

For applications involving sensitive video data from multiple distributed sources, exploring federated learning approaches could allow for the refinement of models like TCGA or Qwen2-VL without centralizing raw video data, thus preserving user privacy. The system's intelligence can be augmented by integrating the MLLM with external knowledge bases, such as ontologies or encyclopedias, enabling it to understand and reason about named entities, real-world events, and common-sense relationships depicted in videos, leading to more insightful descriptions and query responses. Addressing the engineering challenges of deploying and maintaining the entire system at scale, including distributed vector databases, elastic compute for the LLM inference engine, and robust video stream management across hybrid cloud and edge environments, will be critical. It is also essential to continuously evaluate and mitigate potential biases in the TCGA and MLLM models, ensuring fairness in retrieval results, and developing transparent data governance policies, especially concerning privacy and surveillance in the context of camera streams, as part of an ethical AI framework. Leveraging the MLLM's generative capabilities to automatically create concise summaries, highlight reels, or chapter markers for long videos can make

content more accessible and navigable. Lastly, extending the system's capabilities to support queries in multiple languages and retrieve relevant video content even if the original language of the video or its metadata differs, while considering cultural nuances where applicable, will broaden its reach and utility.

By pursuing these future work directions, the Multimodal Video Retrieval Pipeline and its Application System can evolve into an even more powerful, versatile, and indispensable tool for interacting.

References

- Alayrac, J.-B., Donahue, J., Luc, P., Miech, A., Barr, I., Hasson, Y., ... others (2022). Flamingo: A visual language model for few-shot learning. *Advances in Neural Information Processing Systems*, 35, 23716–23736.
- Azizi, S., Mustafa, B., Ryan, F., Beaver, Z., Freyberg, J., Deaton, J., ... others (2021). Big self-supervised models advance medical image classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (pp. 3478–3488).
- Bai, J., Bai, S., Chu, Y., Cui, Z., Dang, K., Deng, X., ... others (2023a). Qwen technical report. *arXiv preprint arXiv:2309.16609*.
- Bai, J., Bai, S., Yang, S., Wang, S., Tan, S., Wang, P., ... Zhou, J. (2023b). QWen-vl: A frontier large vision-language model with versatile abilities. *arXiv preprint arXiv:2308.12966*.
- Bellagente, M., Tow, J., Mahan, D., Phung, D., Zhuravinskyi, M., Adithyan, R., ... others (2024). StableLM2 1.6 B Technical Report. *arXiv preprint arXiv:2402.17834*.
- Bertasius, G., Wang, H. & Torresani, L. (2021). Is space-time attention all you need for video understanding? In *Proceedings of the 38th International Conference on Machine Learning (ICML)* (Vol. 139, pp. 813–824). PMLR.
- Bondarenko, Y., Nagel, M. & Blankevoort, T. (2023). Quantizable transformers: Removing outliers by helping attention heads do nothing. *Advances in Neural Information Processing Systems*, 36, 75067–75096.
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., ... others (2020). Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33, 1877–1901.
- Cao, X. & Yan, W. Q. (2023). Pose estimation for swimmers in video surveillance. *Multimedia Tools and Applications*.
- Carreira, J. & Zisserman, A. (2017). Quo vadis, action recognition? a new model and the kinetics dataset. In *Proceedings of the IEEE Conference on*

- Computer Vision and Pattern Recognition (CVPR)* (pp. 6299–6308).
- Chandola, V., Banerjee, A. & Kumar, V. (2009, July). Anomaly detection: A survey. *ACM Comput. Surv.*, 41(3), 1–58.
- Chen, L., Li, J., Dong, X., Zhang, P., He, C., Wang, J., ... Lin, D. (2023). ShareGPT4v: Improving large multi-modal models with better captions. *arXiv preprint arXiv:2311.12793*.
- Chen, T., Kornblith, S., Norouzi, M. & Hinton, G. (2020). A simple framework for contrastive learning of visual representations. In *International Conference on Machine Learning* (pp. 1597–1607).
- Chen, X., Li, H. & Zhang, J. (2022, June). Edge computing for intelligent video surveillance: A survey. *IEEE Internet Things J.*, 9(12), 9594–9619.
- Cherti, M., Beaumont, R., Wightman, R., Wortsman, M., Ilharco, G., Gordon, C., ... Jitsev, J. (2023). Reproducible scaling laws for contrastive language-image learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 2818–2829).
- Chiang, W.-L., Li, Z., Lin, Z., Sheng, Y., Wu, Z., Zhang, H., ... others (2023). Vicuna: An open-source chatbot impressing GPT-4 with 90%* chatGPT quality. See <https://vicuna.lmsys.org> (accessed 14 April 2023), 2(3), 6.
- Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H. & Bengio, Y. (2014). Learning phrase representations using rnn encoder-decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)* (pp. 1724–1734).
- Chowdhery, A., Narang, S., Devlin, J., Bosma, M., Mishra, G., Roberts, A., ... others (2023). Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240), 1–113.
- Csordas, R., Irie, K. & Schmidhuber, J. (2023). Approximating two-layer feedforward networks for efficient transformers. *arXiv preprint arXiv:2310.10837*.
- Csordas, R., Irie, K., Schmidhuber, J., Potts, C. & Manning, C. D. (2024). Moeut: Mixture-of-experts universal transformers. *arXiv preprint arXiv:2405.16039*.
- Csordas, R., Piekos, P., Irie, K. & Schmidhuber, J. (2024). SwitchHead: Accelerating transformers with mixture-of-experts attention. *Advances in Neural Information Processing Systems*, 37, 74411–74438.
- Dai, W., Li, J., Li, D., Tiong, A. M. H., Zhao, J., Wang, W., ... Hoi, S. (2024).

- Instructblip: Towards general-purpose vision-language models with instruction tuning. *Advances in Neural Information Processing Systems*, 36.
- Dao, T. & Gu, A. (2024). Transformers are SSMS: Generalized models and efficient algorithms through structured state space duality. In *International conference on machine learning ICML*. OpenReview.net. Retrieved from <https://openreview.net/forum?id=ztn8FCR1td>
- Darcet, T., Oquab, M., Mairal, J. & Bojanowski, P. (2023). Vision transformers need registers. *arXiv preprint arXiv:2309.16588*.
- Davis, J. & Goadrich, M. (2006). The relationship between precision-recall and roc curves. In *Proceedings of the 23rd International Conference on Machine learning (ICML)* (pp. 233–240).
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K. & Fei-Fei, L. (2009). ImageNet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern recognition (CVPR)* (pp. 248–255).
- Dey, R. & Salem, F. M. (2017). Gate-variants of gated recurrent unit (GRU) neural networks. In *IEEE 60th International Midwest Symposium on Circuits and Systems (MWSCAS)* (pp. 1597–1600).
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., ... others (2020). An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations (iclr)*.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., ... others (2021). An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations (ICLR)*.
- Fang, B., Wu, W., Liu, C., Zhou, Y., Song, Y., Wang, W., ... Wang, J. (2023). UATVR: Uncertainty-adaptive text-video retrieval. In *International Conference on Computer Vision (ICCV)*.
- Gedara, K. & Yan, W. Q. (2022). Visual blockchain for intelligent surveillance in smart cities. IGI Global.
- Gedara, K. M., Nguyen, M. & Yan, W. Q. (2023). Enhancing privacy protection in intelligent surveillance: Video blockchain solutions. In *5th International Congress on Blockchain and Applications* (pp. 42–51). Cham: Springer Nature Switzerland.
- Glorot, X. & Bengio, Y. (2010). Understanding the difficulty of training deep

- feedforward neural networks. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics (AISTATS)* (Vol. 9, pp. 249–256).
- Grattafiori, A., Dubey, A., Jauhri, A., Pandey, A., Kadian, A., Al-Dahle, A., ... others (2024). The Llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Gu, A. & Dao, T. (2023). Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*.
- Gu, X., Pang, T., Du, C., Liu, Q., Zhang, F., Du, C., ... Lin, M. (2024). When attention sink emerges in language models: An empirical view. *arXiv preprint arXiv:2410.10781*.
- He, K., Zhang, X., Ren, S. & Sun, J. (2015). Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)* (pp. 1026–1034).
- He, K., Zhang, X., Ren, S. & Sun, J. (2016a). Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 770–778).
- He, K., Zhang, X., Ren, S. & Sun, J. (2016b). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 770–778).
- Hochreiter, S. & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780.
- Hu, E. J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., ... Chen, W. (2022). LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations (ICLR)*.
- Hua, W., Dai, Z., Liu, H. & Le, Q. V. (2022). Transformer quality in linear time. In *International Conference on Machine Learning, ICML* (Vol. 162, pp. 9099–9117). PMLR. Retrieved from <https://proceedings.mlr.press/v162/hua22a.html>
- Huang, G., Liu, Z., Van Der Maaten, L. & Weinberger, K. Q. (2017). Densely connected convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 4700–4708).
- Javaheripi, M., Bubeck, S., Abdin, M., Aneja, J., Bubeck, S., Mendes, C. C. T., ... others (2023). Phi-2: The surprising power of small language models. *Microsoft Research Blog*.

- Jiang, A. Q., Sablayrolles, A., Roux, A., Mensch, A., Savary, B., Bamford, C., ... others (2024). Mixtral of experts. *arXiv preprint arXiv:2401.04088*.
- Jiang, D., Liu, Y., Liu, S., Zhang, X., Li, J., Xiong, H. & Tian, Q. (2023). From CLIP to DINO: Visual encoders shout in multi-modal large language models. *arXiv preprint arXiv:2310.08825*.
- Kieran, D. & Yan, W. (2010). A framework for an event driven video surveillance system. In *2010 7th IEEE International Conference on Advanced Video and Signal Based Surveillance* (p. 97-102). doi: 10.1109/AVSS.2010.57
- Kuehne, H., Jhuang, H., Garrote, E., Poggio, T. & Serre, T. (2011). HMDB: A large video database for human motion recognition. In *2011 International Conference on Computer Vision (ICCV)* (pp. 2556–2563).
- Li, A., Gong, B., Yang, B., Shan, B., Liu, C., Zhu, C., ... others (2025). MiniMax-01: Scaling foundation models with lightning attention. *arXiv preprint arXiv:2501.08313*.
- Li, B., Zhang, P., Yang, J., Zhang, Y., Pu, F. & Liu, Z. (2023). OtterHD: A high-resolution multi-modality model. *arXiv preprint arXiv:2311.04219*.
- Li, J., Li, D., Savarese, S. & Hoi, S. (2023a). BLIP-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International Conference on Machine Learning* (pp. 19730–19742).
- Li, J., Li, D., Savarese, S. & Hoi, S. (2023b). Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *Proceedings of the 40th International Conference on Machine Learning (ICML)* (Vol. 202, pp. 19730–19748). PMLR.
- Li, Y., Zhang, Y., Wang, C., Zhong, Z., Chen, Y., Chu, R., ... Jia, J. (2024). Mini-Gemini: Mining the potential of multi-modality vision language models. *arXiv preprint arXiv:2403.18814*.
- Liang, C. & Yan, W. Q. (2022). Human action recognition from digital videos based on deep learning. In *ACM ICCCV*.
- Liang, S. & Yan, W. Q. (2022). A hybrid ctc+attention model based on end-to-end framework for multilingual speech recognition. *Multimedia Tools and Applications*.
- Lin, B., Zhu, B., Ye, Y., Ning, M., Jin, P. & Yuan, L. (2023). Video-LLaVA: Learning united visual representation by alignment before projection. *arXiv preprint arXiv:2311.10122*.
- Lin, T.-Y., Dollár, P., Girshick, R., He, K., Hariharan, B. & Belongie, S. (2017).

- Feature pyramid networks for object detection. In *Proceedings of the IEEE CVPR* (pp. 2117–2125).
- Lin, T.-Y., Goyal, P., Girshick, R., He, K. & Dollár, P. (2017). Focal loss for dense object detection. In *Proceedings of the IEEE ICCV* (pp. 2980–2988).
- Lin, Z., Nikishin, E., He, X. O. & Courville, A. (2025). Forgetting transformer: Softmax attention with a forget gate. *arXiv preprint arXiv:2503.02130*.
- Liu, H., Li, C., Li, Y. & Lee, Y. J. (2023). Improved baselines with visual instruction tuning. *arXiv preprint arXiv:2310.03744*.
- Liu, H., Li, C., Li, Y., Li, B., Zhang, Y., Shen, S. & Lee, Y. J. (2024). *LLaVA-NeXT: Improved reasoning, OCR, and world knowledge*. Retrieved from <https://llava-vl.github.io/blog/2024-01-30-llava-next/>
- Liu, H., Li, C., Wu, Q. & Lee, Y. J. (2024). Visual instruction tuning. *Advances in Neural Information Processing Systems*, 36.
- Liu, Y., Zhang, X., Wang, L. & Zhao, C. (2023, May). Federated learning for video surveillance: A survey of state-of-the-art. *IEEE Trans. Ind. Inform.*, 19(5), 4875–4889.
- Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., ... Guo, B. (2021). Swin Transformer: Hierarchical vision transformer using shifted windows. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)* (p. 9992-10002). Los Alamitos, CA, USA: IEEE Computer Society. Retrieved from <https://doi.ieeecomputersociety.org/10.1109/ICCV48922.2021.00986> doi: 10.1109/ICCV48922.2021.00986
- Locatello, F., Unlu, D., Beaver, S., Gelly, S., Monserrat, B., Puigdomenech, A., ... Soyer, H. (2020). Object-centric learning with slot attention. In *Advances in Neural Information Processing Systems (NeurIPS)* (Vol. 33, pp. 11570–11581).
- Lu, J. & Yan, W. Q. (2020). Deep learning methods for human behavior recognition. In *International Conference on Image and Vision Computing New Zealand*.
- Luo, R., Zhao, Z., Yang, M., Dong, J., Qiu, M., Lu, P., ... Wei, Z. (2023). Valley: Video assistant with large language model enhanced ability. *arXiv preprint arXiv:2306.07207*.
- Ma, B. & Yan, W. Q. (2024). JudPriNet: Video transition detection based on semantic relationship and Monte Carlo sampling. *IEEE Intelligent and*

Converged Networks.

- Ma, J., Liu, W., Miller, P. & Yan, W. (2009). Event Composition with Imperfect Information for Bus Surveillance. In *2009 Sixth IEEE International Conference on Advanced Video and Signal Based Surveillance* (p. 382-387). doi: 10.1109/AVSS.2009.25
- Maaz, M., Rasheed, H., Khan, S. & Khan, F. S. (2023). Video-ChatGPT: Towards detailed video understanding via large vision and language models. *arXiv preprint arXiv:2306.05424*.
- McKinzie, B., Gan, Z., Fauconnier, J.-P., Dodge, S., Zhang, B., Dufter, P., ... others (2024). Mm1: Methods, analysis & insights from multimodal llm pre-training. *arXiv preprint arXiv:2403.09611*.
- Miller, E. (2023). *Attention is off by one*. July. Retrieved from <https://www.evanmiller.org/attention-is-off-by-one.html>
- OpenAI. (2023a). *Chatgpt*. <https://openai.com/blog/chatgpt/>.
- OpenAI. (2023b). GPT-4v(ision) system card.
- OpenAI, Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., ... Zoph, B. (2024). *GPT-4 Technical Report*. Retrieved from <https://arxiv.org/abs/2303.08774>
- Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., ... others (2022). Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35, 27730–27744.
- Pang, G., Shen, C., van den Hengel, A., Bai, J. & Chang, D. (2020, August). Deep learning for video anomaly detection: A review. *IEEE Trans. Circuits Syst. Video Technol.*, 31(8), 3145–3168.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., ... others (2019). PyTorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems (NeurIPS)* (pp. 8026–8037).
- Piękos, P., Csordás, R. & Schmidhuber, J. (2025). *Mixture of sparse attention: Content-based learnable sparse attention via expert-choice routing*. Retrieved from <https://arxiv.org/abs/2505.00315>
- Qin, Z., Sun, W., Li, D., Shen, X., Sun, W. & Zhong, Y. (2024a). Lightning attention-2: A free lunch for handling unlimited sequence lengths in large language models. *arXiv preprint arXiv:2401.04658*.

- Qin, Z., Sun, W., Li, D., Shen, X., Sun, W. & Zhong, Y. (2024b). Various lengths, constant speed: Efficient language modeling with lightning attention. *arXiv preprint arXiv:2405.17381*.
- Qiu, Z., Wang, Z., Zheng, B., Huang, Z., Wen, K., Yang, S., ... Lin, J. (2025). *Gated attention for large language models: Non-linearity, sparsity, and attention-sink-free*. Retrieved from <https://arxiv.org/abs/2505.06708>
- Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., ... others (2021). Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning* (pp. 8748–8763).
- Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., ... Liu, P. J. (2020). Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(140), 1–67.
- Rahman, W., Hasan, M. K., Lee, S., Zadeh, A., Mao, C., Morency, L.-P. & Hoque, E. (2020). Integrating multimodal information in large pretrained transformers. In *Proceedings of the Conference. Association for Computational Linguistics* (Vol. 2020, p. 2359).
- Ramapuram, J., Danieli, F., Dhekane, E., Weers, F., Busbridge, D., Ablin, P., ... others (2024). Theory, analysis, and best practices for sigmoid self-attention. *arXiv preprint arXiv:2409.04431*.
- Redmon, J., Divvala, S., Girshick, R. & Farhadi, A. (2016, June). You Only Look Once: Unified, Real-Time Object Detection. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 779–788). Las Vegas, NV, USA: IEEE Computer Society.
- Sai Hareesh, S. K. L. V., Kumari, P. L. & Sreenivasu, S. V. N. (2021). Deep learning for intelligent video surveillance: A survey. *Multimed. Tools Appl.*, 80(20), 30233–30275.
- Shazeer, N. (2020). GLU variants improve transformer. *arXiv preprint arXiv:2002.05202*.
- Soomro, K., Zamir, A. R. & Shah, M. (2012). *UCF101: A dataset of 101 human actions classes from videos in the wild* (Tech. Rep.). CRCV-TR-12-01, University of Central Florida.
- Srivastava, R. K., Greff, K. & Schmidhuber, J. (2015). Highway networks. *arXiv preprint arXiv:1505.00387*.
- Sun, C., Myers, A., Vondrick, C., Murphy, K. & Schmid, C. (2019). VideoBERT:

- A joint model for video and language representation learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (pp. 7464–7473).
- Sun, M., Chen, X., Kolter, J. Z. & Liu, Z. (2024). Massive activations in large language models. *arXiv preprint arXiv:2402.17762*.
- Sun, Q., Fang, Y., Wu, L., Wang, X. & Cao, Y. (2023). EVA-CLIP: Improved training techniques for clip at scale. *arXiv preprint arXiv:2303.15389*.
- Sun, Y., Dong, L., Huang, S., Ma, S., Xia, Y., Xue, J., ... Wei, F. (2023). *Retentive network: A successor to transformer for large language models*. Retrieved from <https://arxiv.org/abs/2307.08621>
- Team, G., Anil, R., Borgeaud, S., Wu, Y., Alayrac, J.-B., Yu, J., ... others (2023). Gemini: A family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Tong, S., Liu, Z., Zhai, Y., Ma, Y., LeCun, Y. & Xie, S. (2024). Eyes wide shut? exploring the visual shortcomings of multimodal llms. *arXiv preprint arXiv:2401.06209*.
- Touvron, H., Cord, M., Douze, M., Massa, F., Sablayrolles, A. & Jegou, H. (2021, 18–24 Jul). Training data-efficient image transformers & distillation through attention. In M. Meila et al. (Eds.), *Proceedings of the 38th International Conference on Machine Learning* (Vol. 139, pp. 10347–10357). PMLR. Retrieved from <https://proceedings.mlr.press/v139/touvron21a.html>
- Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., ... others (2023a). Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.-A., Lacroix, T., ... others (2023b). Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... Polosukhin, I. (2017). Attention is all you need. In *Advances in Neural Information Processing Systems (NeurIPS)* (pp. 5998–6008).
- Wake, N., Kanehira, A., Sasabuchi, K., Takamatsu, J. & Ikeuchi, K. (2023). GPT-4V (ision) for robotics: Multimodal task planning from human demonstration. *arXiv preprint arXiv:2311.12015*.
- Wang, J., Kankanhalli, M. S., Yan, W. & Jain, R. (2003). Experiential sampling for video surveillance. In *First ACM SIGMM International Workshop on*

- Video Surveillance* (pp. 77–86). New York, New York, USA: ACM Press.
- Wang, J., Yan, W., Kankanhalli, M., Jain, R. & Reinders, M. (2004, 01). Adaptive Monitoring for Video Surveillance. In *International Conference on Information, Communications and Signal Processing, 2003 and the Fourth Pacific Rim Conference on Multimedia*. (p. 1139 - 1143 vol.2). doi: 10.1109/ICICS.2003.1292638
- Wang, W., Lv, Q., Yu, W., Hong, W., Qi, J., Wang, Y., ... others (2023). CogVLM: Visual expert for pretrained language models. *arXiv preprint arXiv:2311.03079*.
- Wei, J., Bosma, M., Zhao, V. Y., Guu, K., Yu, A. W., Lester, B., ... Le, Q. V. (2021). Finetuned language models are zero-shot learners. *arXiv preprint arXiv:2109.01652*.
- Wu, C., Lei, J., Zheng, Q., Zhao, W., Lin, W., Zhang, X., ... others (2023). Can gpt-4v (ision) serve medical applications? case studies on gpt-4v for multimodal medical diagnosis. *arXiv preprint arXiv:2310.09909*.
- Wu, W., Luo, H., Fang, B., Wang, J. & Ouyang, W. (2023). Cap4video: What can auxiliary captions do for text-video retrieval? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 10704–10713).
- Wu, W., Yao, H., Zhang, M., Song, Y., Ouyang, W. & Wang, J. (2023). GPT4Vis: What Can GPT-4 Do for Zero-shot Visual Recognition? *arXiv preprint arXiv:2311.15732*.
- Wu, W., Sun, Z., Song, Y., Wang, J. & Ouyang, W. (2023). Transferring vision-language models for visual recognition: A classifier perspective. *International Journal of Computer Vision*, 1–18.
- Xiao, G., Tian, Y., Chen, B., Han, S. & Lewis, M. (2023). Efficient streaming language models with attention sinks. *arXiv preprint arXiv:2309.17453*.
- Yan, W. Q. (2019a). *Introduction to Intelligent Surveillance: Surveillance Data Capture, Transmission, and Analytics*. Germany: Springer.
- Yan, W. Q. (2019b). Visual Event Computing. In *Texts in Computer Science* (pp. 155–165). Cham: Springer International Publishing.
- Yan, W. Q. (2023). *Computational Methods for Deep Learning: Theory, Algorithms, and Implementations*. Germany: Springer Nature.
- Yan, W. Q. & Liu, F. (2016). Event analogy based privacy preservation in visual surveillance. In *Image and Video Technology – PSIVT 2015 Workshops* (pp.

- 357–368). Springer International Publishing.
- Yang, A., Yang, B., Zhang, B., Hui, B., Zheng, B., Yu, B., ... others (2024). Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*.
- Yang, S., Kautz, J. & Hatamizadeh, A. (2024). Gated delta networks: Improving mamba2 with delta rule. *arXiv preprint arXiv:2412.06464*.
- Yang, Z., Dai, Z., Yang, Y., Carbonell, J., Salakhutdinov, R. R. & Le, Q. V. (2019). XLNet: Generalized autoregressive pretraining for language understanding. *Advances in Neural Information Processing Systems*, 32.
- Yang, Z., Li, L., Lin, K., Wang, J., Lin, C.-C., Liu, Z. & Wang, L. (2023). The dawn of lmms: Preliminary explorations with gpt-4v (ision). *arXiv preprint arXiv:2309.17421*, 9(1), 1.
- Yao, L., Huang, R., Hou, L., Lu, G., Niu, M., Xu, H., ... Xu, C. (2021). FILIP: Fine-grained interactive language-image pre-training. In *International Conference on Learning Representations (ICLR)*.
- Young, A., Chen, B., Li, C., Huang, C., Zhang, G., Zhang, G., ... others (2024). Yi: Open foundation models by 01. ai. *arXiv preprint arXiv:2403.04652*.
- Yuan, J., Gao, H., Dai, D., Luo, J., Zhao, L., Zhang, Z., ... others (2025). Native sparse attention: Hardware-aligned and natively trainable sparse attention. *arXiv preprint arXiv:2502.11089*.
- Zhai, X., Mustafa, B., Kolesnikov, A. & Beyer, L. (2023). Sigmoid loss for language image pre-training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (pp. 11975–11986).
- Zhang, H., Li, X. & Bing, L. (2023). Video-Llama: An instruction-tuned audio-visual language model for video understanding. *arXiv preprint arXiv:2306.02858*.
- Zhang, P., Zeng, G., Wang, T. & Lu, W. (2024). TinyLlama: An open-source small language model. *arXiv preprint arXiv:2401.02385*.
- Zhang, S., Roller, S., Goyal, N., Artetxe, M., Chen, M., Chen, S., ... others (2022). OPT: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*.
- Zheng, A. & Yan, W. Q. (2024). Attention-based multimodal fusion model for breast cancer diagnostics. In *ICONIP*.
- Zheng, A. & Yan, W. Q. (2025). Attention-Pool: 9-ball game video analytics system with object attention and temporal context gated attention. *Computers*.

-
- Zhou, L., Yan, W., Shu, Y. & Yu, J. (2018). CVSS: A cloud-based visual surveillance system. *Int. J. Digit. Crime Forensics*, 10(1), 79–91.
- Zhu, D., Chen, J., Shen, X., Li, X. & Elhoseiny, M. (2023). MiniGPT-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*.
- Zuhri, Z. M. K., Fuadi, E. H. & Aji, A. F. (2025). *Softpick: No attention sink, no massive activations with rectified softmax*. Retrieved from <https://arxiv.org/abs/2504.20966>