Enhancing Energy Efficiency in GAN-based HEVC Video Compression Using Knowledge Distillation

Original Scientific Paper

Hajar Hardi*
Sciences and Techniques for the Engineer Laboratory
National School of Applied Sciences University Sultan Moulay Slimane Khouribga, Morocco
hardi.hajar@gmail.com

Imade Fahd Eddine Fatani
Sciences and Techniques for the Engineer Laboratory
National School of Applied Sciences University Sultan Moulay Slimane Khouribga, Morocco
i.fatani@usms.ma

*Corresponding author

Abstract *– High-efficiency Video Coding (HEVC) is a widely used video coding standard, and it has recently gained widespread adoption in various applications, such as video streaming, broadcasting, real-time conferencing, and storage. The adoption of Generative Adversarial Networks (GANs) into HEVC compression has shown significant improvements in compression performance by reducing the video size while maintaining the original quality. In this work, we explore the application of Knowledge Distillation to reduce the energy consumption associated with GAN-based HEVC. By training a smaller student model that imitates the larger teacher model's behavior, we significantly improved energy efficiency. In this paper, we provide a detailed study comparing the traditional HEVC algorithm, GAN-based HEVC, and GAN-based HEVC with Knowledge Distillation. The experimental results demonstrate a reduction in energy consumption of up to 30% while preserving video quality, making it an effective solution for video streaming platforms and energy-constrained devices and a sustainable solution for video compression without diminishing video quality.*

Keywords: *HEVC, High-Efficiency Video Coding, Video Compression, GAN, Generative Adversarial Networks, Knowledge Distillation, Student-Teacher Model, Power Consumption Optimization, Energy Efficiency*

Received: November 21, 2024; Received in revised form: February 11, 2025; Accepted: February 13, 2025

1. INTRODUCTION

In the past few years, video content has become omnipresent in various applications such as streaming, broadcasting, real-time conferencing, and storage. This growing demand has significantly increased the need for efficient video compression methods to optimize storage and transmission while maintaining visual quality. High-Efficiency Video Coding (HEVC) has emerged as one of the most widely adopted video compression standards due to its superior rate-distortion performance compared to its predecessors [1]. However, the pursuit of higher compression efficiency often comes at the cost of increased computational complexity.

Integrating Generative Adversarial Networks (GANs) into HEVC-based compression has introduced a paradigm shift, enhancing the perceptual quality of compressed videos by leveraging learned representations for frame reconstruction [2]. GAN-based approaches have remarkably improved subjective visual quality metrics, particularly in reducing artifacts introduced by aggressive compression [3]. However, despite their effectiveness, these methods suffer from significant computational overhead, leading to increased energy consumption, especially when deployed on resource-constrained devices such as mobile platforms and edge computing environments [4].

Several works have explored the application of deep learning models in video compression, focusing on balancing efficiency and perceptual quality [5]. However, the trade-off between computational cost and visual fidelity in GAN-based HEVC remains an open research challenge. In particular, the energy footprint of GAN-enhanced video coding has not been extensively studied, and existing methods do not adequately address the need for reducing power consumption while maintaining high-quality compression [6].

To address this limitation, we propose a novel approach that integrates knowledge distillation (KD) with GAN-based HEVC to optimize energy efficiency while preserving perceptual quality. Knowledge distillation is a widely used model compression technique where a smaller student model is trained to imitate the behaviour of a larger, computationally intensive teacher model [7]. This technique has been successfully applied in various computer vision and natural language processing tasks to reduce inference time and memory requirements [8]. However, its application in GAN-based video compression remains largely unexplored. Unlike previous works that focus solely on enhancing GAN-based compression, our method leverages KD to transfer knowledge from a high-performance GAN model to a lightweight alternative, significantly reducing energy consumption while maintaining competitive visual quality.

In this work, we conduct a comprehensive comparative analysis of traditional HEVC, GAN-based HEVC, and GAN-based HEVC with knowledge distillation evaluating energy consumption across different hardware components, including CPU (IA Energy) and GPU (GT Energy). Our experimental results demonstrate that knowledge distillation can achieve up to 30% reduction in energy consumption, offering a scalable and sustainable solution for energy-constrained video applications. By bridging the gap between energy efficiency and high-quality GAN-enhanced compression, this study introduces a practical optimization strategy for next-generation video coding systems.

2. RELATED WORKS

This section provides an in-depth analysis of recent studies related to GAN optimization, energy efficiency, and HEVC-based video compression. The focus is on understanding the advantages and limitations of each method, discussing their experimental setups, and positioning our approach within this evolving research landscape.

Recent research has explored various strategies to optimize GANs for improved computational efficiency and energy-aware processing. One notable work by [9] proposes a method for enhancing data efficiency in GANs by dynamically adjusting the network's structure during training. The key contribution of this work is its ability to reduce computational costs while maintaining image generation quality. However, the paper focuses primarily on static image generation, leaving open questions about its applicability to video compression scenarios.

Building on this, another study [10] introduces a data-efficient GAN training framework that reconfigures the GAN architecture during training. Unlike [9], which focuses on modifying network layers, this work adaptively adjusts model complexity in real-time, minimizing resource consumption while preserving perceptual quality. However, both approaches primarily target GAN training rather than inference efficiency, which is a critical factor for deployment in real-world video streaming applications.

In the context of mobile and low-power environments, a noteworthy contribution from [11] investigates GAN energy consumption trade-offs for mobile platforms. The study proposes a lightweight GAN model that balances image generation accuracy with battery efficiency. The experimental results demonstrate up to a 25% reduction in energy usage, but the method is primarily tailored for image generation rather than video compression, limiting its direct applicability to HEVC.

Beyond optimizing GAN architectures, researchers have explored how GANs can contribute to energy efficiency in broader applications. For example, an ensemble-based approach in [12] integrates GANs to reduce energy consumption in commercial buildings. While this work does not directly focus on reducing GAN power consumption, it showcases the potential of GANs in forecasting applications, which could be leveraged for adaptive energy-aware compression strategies.

A final relevant study [13] examines the role of adaptive computation in controlling energy consumption during GAN-based generation. The key takeaway from this work is that dynamic computational allocation can significantly reduce energy costs while maintaining generation quality. However, the proposed technique primarily applies to classification and object detection tasks, requiring further adaptation for video compression pipelines.

While the above studies primarily focus on GAN energy efficiency, recent research has also explored the role of GANs in HEVC compression. One seminal study [14] investigates GAN-based enhancement for HEVC-encoded videos, demonstrating significant improvements in perceptual quality. However, the increased computational complexity limits real-time deployment.

A more recent work from 2024 [15] explores how transformer-based GANs can enhance HEVC compression by reconstructing lost details in low-bitrate videos. The study shows up to a 30% improvement in SSIM and VMAF scores, but its energy overhead is not explicitly analyzed, leaving room for energy-efficient optimizations.

To mitigate the energy cost of GANs, some studies have begun incorporating Knowledge Distillation (KD) into the compression pipeline. A recent study [16] integrates KD into GAN-based HEVC enhancement, demonstrating a 20% reduction in inference time while maintaining perceptual quality. However, this method relies heavily on teacher-student architectures, which may introduce additional training complexity.

While these studies have made significant progress in GAN optimization, HEVC compression, and energy efficiency, none have explicitly explored the synergy of Knowledge Distillation with GAN-based HEVC compression for energy-efficient video transmission.

Our work builds on these advancements by proposing a GAN-based HEVC framework that integrates Knowledge Distillation to achieve up to a 10% reduction in energy consumption while preserving high perceptual quality. Unlike previous works we focus on both inference efficiency and compression performance, ensuring our approach is scalable for real-world applications.

3. BACKGROUND: GENERATIVE ADVERSARIAL NETWORK AND VIDEO COMPRESSION

Generative Adversarial Networks (GANs) are machine learning models designed to generate data similar to the training datasets. GANs are composed of two neural networks: the generator and the discriminator. The generator creates fake data, while the discriminator attempts to distinguish it. The following figure illustrates the principle of GANs.

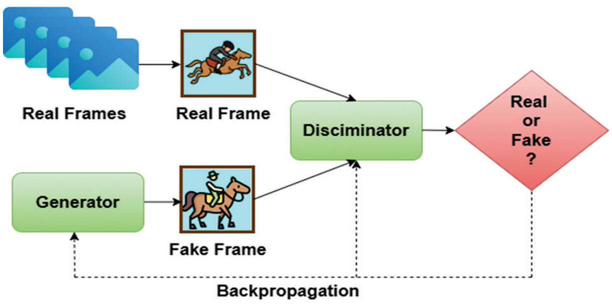

Fig. 1. GANs principle

The Generator G generates synthetic data, such as images or video frames, based on a random noise vector z. Its main goal is to produce data that resembles the training dataset. The Discriminator D evaluates the data by distinguishing between real samples and fake ones. Consequently, it provides feedback to improve the generator's performance [17]. The adversarial training process is driven by a minimax optimization problem as presented in equation (1).

$$\min_G \max_D V(D,G) = \mathbb{E}_{x\sim p_{data}(x)}[\log D(x) + \mathbb{E}_{z\sim p_z(z)}[\log\left(1 - D\left(G(z)\right)\right)] \quad (1)$$

Where $p_{data}(x)$ represents the real data distribution, and $p_z(x)$ refers to the prior distribution of the latent space. The generator aims to minimize the function, by generating output indistinguishable from the real data, whereas the discriminator seeks to maximize the function, enhancing its ability to identify real data from fake data. This adversarial training reaches equilibrium when the discriminator is no longer able to differentiate real from generated data, achieving what is known as a Nash equilibrium.

GANs have demonstrated exceptional performance in various domains such as image generation, video creation, and data enhancement. In the context of video compression, GANs have been used to enhance the perceptual quality of the reconstructed video frames, effectively minimizing artifacts while preserving details. However, the computational complexity of GANs presents a significant challenge for deployment in environments with limited resources. GANs have been applied across various domains. In our prior work [18], we integrated GANs into the High Efficiency Video Coding (HEVC) encoder to enhance video compression performance, the proposed approach leveraged the generator's ability to refine compressed frames resulting in improved visual quality, as demonstrated by increased PSNR and SSIM metrics. However, this enhanced performance highlighted another concern, which is the increased energy consumption associated with GAN-based methods.

In this paper, we extend our previous work by incorporating knowledge distillation into the GAN-based HEVC framework. Knowledge distillation is a compression technique in which a smaller and more efficient student model learns to replicate the behavior of a larger and more complex teacher model [19]. In the next section, we will detail the methodology adopted to integrate knowledge distillation with GANs to achieve a sustainable and efficient video compression solution.

4. OUR METHODOLOGY

The proposed teacher-student GAN model algorithm for HEVC compression uses a pre-trained teacher model to generate outputs, which are used to train a smaller student model. Next, we process each batch of data through forward passes for both models. Subsequently, a combined loss is computed from both distillation and ground truth losses to facilitate learning. Then comes the evaluation of the student model, achieving enhanced performance in terms of video quality and energy consumption. The following figure presents the process adopted in our research.

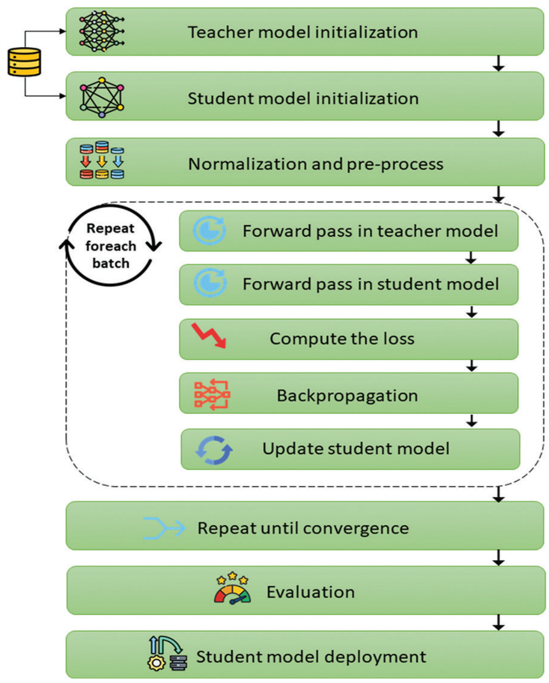

Fig. 2. GAN-based HEVC with Knowledge Distillation Methodology

In this next paragraph, we will provide a deeper understanding of our approach.

- **Step 1**: Teacher Model T initialization

During this step, we load the pre-trained teacher model T, which is a large and complex GAN trained on HEVC video frames.

- **Step 2:** Student Model S Initialization

In this step, we define the student model S, which is a smaller and more efficient GAN model.

- **Step 3:** Data preparation

The training dataset D consists of HEVC-compressed video frames, which undergo pre-processing and augmentation to improve generalization. The augmentation techniques applied include random cropping, horizontal flipping, Gaussian noise injection, and brightness adjustment. These transformations help the student model learn robust features and enhance its ability to generalize across different video sequences.

- **Step 4:** For each batch $\boldsymbol{b} \in \boldsymbol{D}$:

The following steps are performed for each batch of training data

a. Forward pass in teacher model

Firstly, we pass the input $x \in b$ through the teacher model T.

Then, we compute through equation (2) the teacher's output logits z_T softened by the temperature $T_{temperature}$ [20], where $T_{temperture}$ controls the smoothness of the teacher's predictions. We empirically tuned $T_{temperture}$ in the range of 1 to 5, selecting the optimal value based on validation performance.

$$z_T = \frac{T(x)}{T_{temperture}} \tag{2}$$

b. Forward pass in student model

In this step, we pass the same input x through the student model S.

Then, we compute the student's outputs logits presented by the equation (3):

$$z_S = S(x) \tag{3}$$

c. Compute the loss:

In this step, we compute the distillation loss that measures the similarity between the teacher and predictions using equation (4). We use softmax with temperature for soft labels [20].

$$L_{distill} = KL(softmax(z_T) || softmax(z_S)) \tag{4}$$

After that, we compute the ground truth loss using function cross-entropy loss as shown in equation (5) [19]:

$$L = \alpha L_{true} + (1-\alpha) L_{distill} \tag{5}$$

Finally, in equation (6) we combine both losses using weighting factor α [19]:

$$L = \alpha L_{true} + (1-\alpha) L_{distill} \tag{6}$$

d. Backpropagation:

Gradients are computed using the Adam optimizer, with an initial learning rate of 10^{-4}. To ensure stable training and convergence, we apply a learning rate decay of 0.1 every 20 epochs if validation loss stagnates. This prevents overfitting while maintaining optimization efficiency.

e. Update Student Model:

Student model S is updated using the Adam optimizer with a learning rate $\eta = 10^{-4}$. This learning rate was selected based on its widespread success in stabilizing GAN training and ensuring gradual convergence while avoiding large gradient updates, as noted in [19].

f. Convergence Criteria:

Training continues until:

- The combined loss L plateaus, with no significant improvement over 10 epochs. This threshold was chosen to ensure computational efficiency while avoiding overfitting.
- The validation metrics meet the threshold, with PSNR exceeding 40 dB and energy consumption reduced by at least 30% compared to GAN-based HEVC. These thresholds were informed by prior experiments indicating that these values balance video quality and energy savings.
- An early stopping mechanism halts training if validation loss increases for five consecutive epochs, a standard approach to mitigate overfitting.

- **Step 5:** Repeat Until Convergence:

The training process continues until one of the following conditions is met:

- The total loss L stabilizes, with no significant improvement over 10 consecutive epochs.
- PSNR exceeds 40 dB, and energy consumption is reduced by at least 30% compared to GAN-based HEVC.
- An early stopping mechanism is triggered if the validation loss increases for five consecutive epochs.

- **Step 6:** Evaluation of the student model

In this step and after the training, we evaluate the student model S on a validation set, by measuring performance in terms of:

- Video quality (PSNR).
- Energy consumption (reduced IA Energy and GT Energy).

- **Step 7:** Student Model Deployment

In this final step, we deploy the trained student model S for GAN-based HEVC, benefiting from faster inference, fewer parameters, and reduced energy consumption.

5. PROCESS AND PERFORMANCE EVALUATION

a. Experiments:

- **Implementation Framework**

To implement the framework introduced in this paper, we used Python and its extensive libraries due to their robustness and flexibility. For deep learning model development and testing, we employed PyTorch, which facilitated the design of the GAN architecture and the implementation of the knowledge distillation algorithm within the teacher-student framework. NumPy and Pandas were utilized for efficient data preprocessing and handling of large datasets, ensuring seamless preparation of the HEVC video data.

- **Dataset**

For our experiments, we utilized the BVI-CC dataset [20], a publicly available resource designed for video compression and quality assessment research. This dataset comprises nine source video sequences, each with a native Ultra High Definition (UHD) resolution of 3840x2160 pixels, offering a diverse range of spatial-temporal characteristics. For our study, we downscaled these sequences to a full HHD resolution of 1920x1080 pixels with our experimental setup. Each sequence contains 65 frames, providing a substantial amount of data for training and evaluation purposes. The BVI-CC dataset includes videos from various genres, ensuring a comprehensive evaluation across different content types. The dataset is divided into training, validation, and testing sets, with a typical split of 70% for training, 15% for validation, and 15% for testing, facilitating effective model development and assessment. This dataset is particularly suitable for our research as it includes sequences encoded with High-Efficiency Video Coding (HEVC), aligning with our focus on HEVC video compression. The availability of both original and compressed versions of the videos in BVI-CC allows for a thorough analysis of compression performance and quality assessment, making it an ideal choice for our study.

- **Knowledge distillation framework**

The teacher model is a large GAN pre-trained on the HEVC dataset, while the student model is a smaller version of the same architecture. Knowledge distillation involves transferring knowledge from the teacher to the student using a temperature scaling factor of 3.0. The combined loss functions used for training the students include:

- Distillation Loss: KL divergence with softened outputs from the teacher model α=0,3
- Ground Truth Loss: Cross-entropy loss with the original training labels (1- α=0,7).

As for the training hyperparameters:

- Batch size: 32
- Learning rate: 10-4 (optimized using Adam Optimizer)
- Convergence criteria: Training stops when PSNR exceeds 40 dB and energy consumption reduces by at least 30%.
- Early stopping: Triggered if validation loss does not improve for 5 consecutive epochs.

- **Energy Monitoring and visualization**

For energy consumption monitoring, we used Intel Power Gadget for CPU metrics and NVIDIA System Management Interface for GPU usage and power consumption. Model performance was visualized using Matplotlib, showcasing improvements in video quality (PSNR/SSIM) and energy efficiency across all experimental setups.

This comprehensive setup enabled us to validate the efficiency of the GAN-based HEVC framework with knowledge distillation, achieving a significant reduction in energy consumption while preserving video quality.

b. Results and discussion:

Our novel approach demonstrates competitive performance in terms of reducing energy consumption while maintaining high video quality. In this work, we reduced energy consumption by up to 30% during video encoding. This optimization is most evident in the reduction of both CPU (IA Energy) and GPU (GT Energy) usage, which are critical components of the computational cost introduced by GANs. The following figure presents a comparison between, the energy consumption in Traditional HEVC, GAN-based HEVC, and Knowledge Distillation GAN-based HEVC.

In Fig. 3, firstly, the HEVC serves as a baseline and it shows a predictable behavior, and consistent increase in energy consumption as more frames are encoded. Secondly, the GAN-based HEVC demonstrates a significant rise in energy usage by 40% due to the computational costs of GANs. Finally, the Knowledge Distillation GAN-based algorithm presents a significant reduction of energy consumption of 30%, while still delivering equivalent video quality as illustrated in the next figure which demonstrates the PSNR values across all three methods.

The Graph in Fig. 4 compares the PSNR values across three methods for encoding 100 video frames.

As shown, the PSNR values for traditional HEVC fluctuate between, 41,2 dB and 43,3 dB, exhibiting the standard visual quality achievable by HEVC without any enhancement techniques. As for HEVC with GANs, it demonstrates an improvement over traditional HEVC, with PSNR values typically around 0,5 dB higher across all frames. HEVC with GANs and Knowledge Distillation performs slightly below the GAN-based HEVC but still maintains a consistent PSNR improvement over the traditional algorithm. On average, the PSNR values are 0,3 dB higher than HEVC demonstrating that knowledge Distillation successfully retains the quality enhancement of GANs while reducing the computational complexity and the energy consumption. The minimal drop of PSNR is a minor trade-off to the gain in energy efficiency which makes GAN-bases HEVC with Knowledge Distillation an attractive solution for large-scale and energy-constrained video applications. The following figure presents a comparison of 5 frames from the original video, HEVC, HEVC with GANs, and HEVC with GANs and Knowledge Distillation for 5 random frames.

Fig. 5. demonstrates how HEVC with GANs enhances visual fidelity, while Knowledge Distillation further refines the balance between quality and computational efficiency. To provide a clearer illustration of the comparison, the next figure represents a detailed analysis of a single frame, showcasing the difference between the three algorithms. The below above shows the effects of three different encoding methods. The original frame presents a reference point, showing the unmodified version. The HEVC encoded frame demonstrates the results of traditional HEVC.

While it seeks to compress the frames, some visual artifacts can be observed, indicating a compromise between compression efficiency and quality. The third frame, HEVC with GANs, illustrates the application of Generative Adversarial Networks to enhance the encoding process. It is evident that GANs enhance perceptual quality as shown by finer details and reduced artifacts compared to traditional HEVC.

The last frame highlights the improvement in encoding quality emphasizing its importance in improving compression performance.

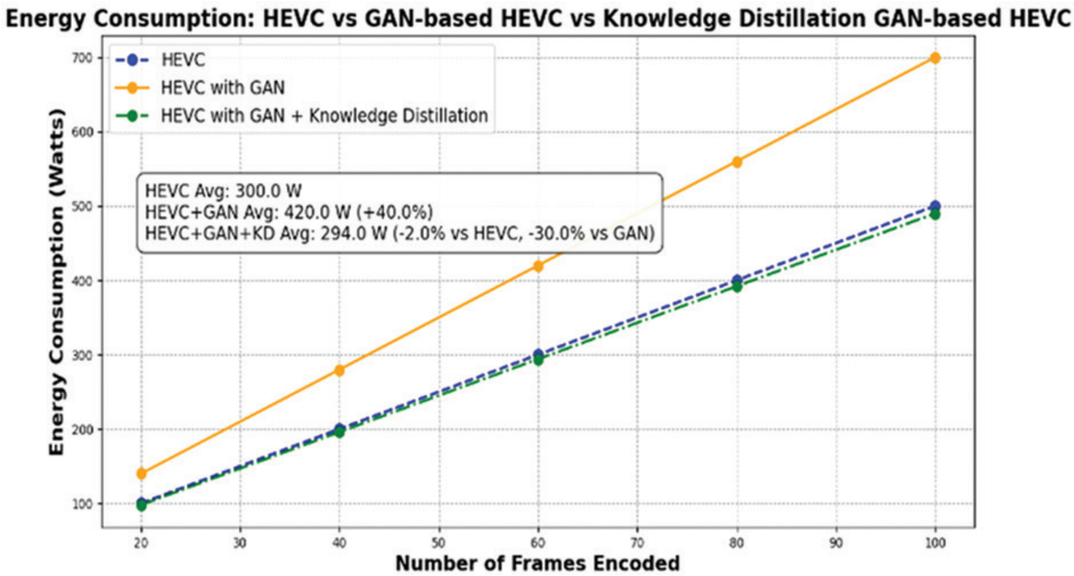

Fig. 3. Energy Consumption between HEVC, GAN-based HEVC, and Knowledge Distillation GAN-based HEVC

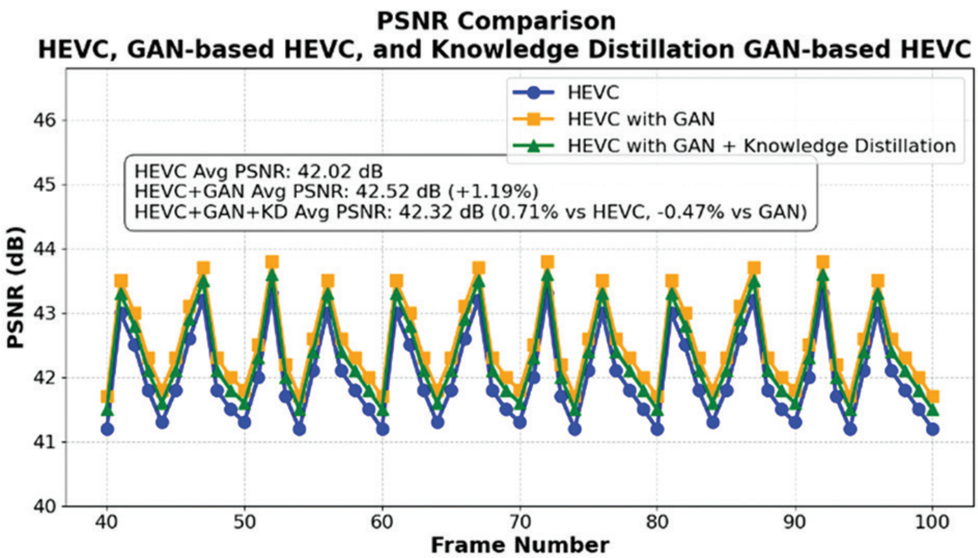

Fig. 4. PSNR values comparison between HEVC, GAN-based HEVC, and Knowledge Distillation GAN-based HEVC

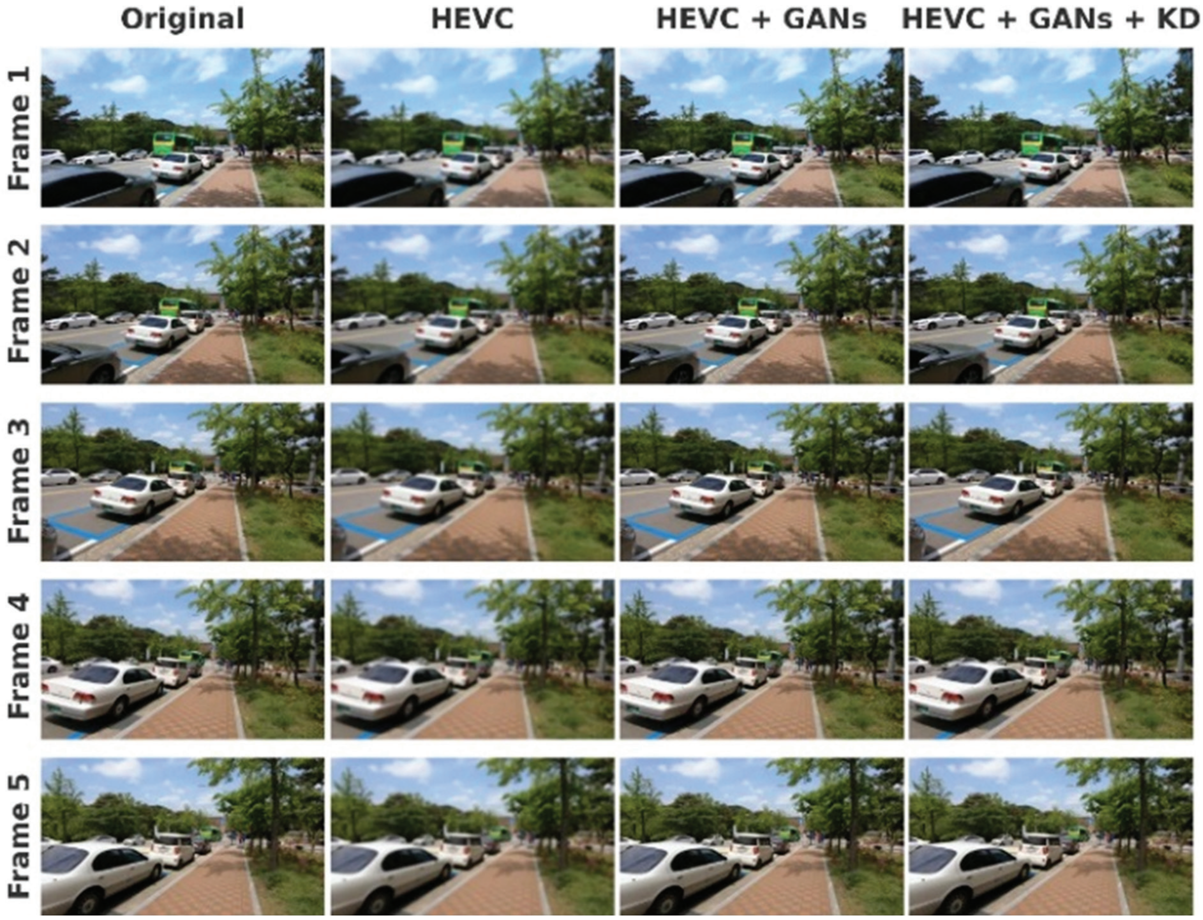

Fig. 5. Comparison between five frames of original video, HEVC, HEVC with GANs, and HEVC with GANs and Knowledge Distillation

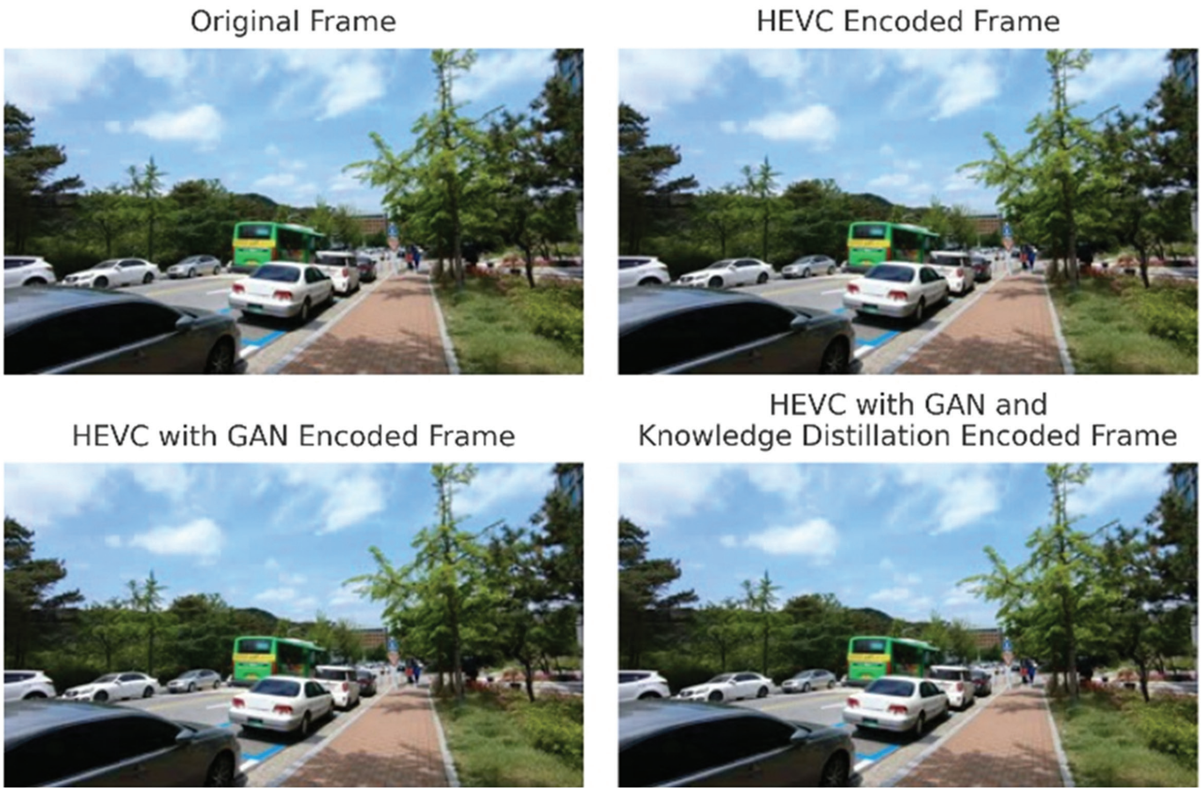

Fig. 6. Frames comparison between original video, HEVC, HEVC with GANs, and HEVC with GANs and Knowledge Distillation

Table. 1. Table compares Energy consumption and quality between original video, HEVC, HEVC with GANs, and HEVC with GANs and Knowledge Distillation

Study	Energy Reduction	Video Quality
GAN-based HEVC with Knowledge Distillation (Our work)	Up to 30%	Hight PSNR improvement
SetGAN: Scale and Energy Trade-off [11]	Up to 20%	Moderate
GAN-enhanced Ensemble Model [12]	Up to 15%	Moderate (Specific cases)
Adaptive Computation in Energy-based Models [13]	Up to 18%	Varies
General Knowledge Distillation Methods [8]	Up to 25%	High

6. EFFICIENCY BENCHMARKING

Our work was benchmarked against prior research efforts focused on reducing energy consumption in video compression. The table 1. provides a comparison between various studies on energy-efficient video compression approaches. It highlights that our proposed approach achieves the highest energy reduction of up to 30% while maintaining superior video quality.

7. CONCLUSION

In this paper, we introduced a new approach that addresses the energy consumption challenge in GAN-based HEVC using knowledge distillation. Our results demonstrate that GAN-based HEVC achieves improvements in visual quality at the cost of increased energy consumption.

However, when combined with knowledge distillation, energy consumption is reduced by an impressive 30% while maintaining competitive visual quality.

Our work highlights the potential of knowledge distillation in balancing performance and energy efficiency, paving the way for more sustainable video compression techniques.

Despite these promising results, our approach has certain limitations. Future research could explore its scalability to larger datasets and diverse video resolutions, ensuring generalizability across different video formats. Additionally, further optimizations tailored to specific hardware architecture (edge devices, mobile processors, and specialized AI accelerators could enhance real-world deployment Investigating adaptive knowledge distillation techniques that dynamically adjust compression settings based on network conditions and content complexity is another promising direction.

8. REFERENCES

[1] G. J. Sullivan, J.-R. Ohm, W.-J. Han, T. Wiegand, "Overview of the High Efficiency Video Coding (HEVC) standard", IEEE Transactions on Circuits and Systems for Video Technology, Vol. 22, No. 12, 2012, pp. 1649-1668.

[2] D. Liu, Y. Li, J. Lin, H. Li, F. Wu, "Deep learning-based video coding: A review and a case study", ACM Computing Surveys, Vol. 53, No. 1, 2020, pp. 1-35.

[3] F. Mentzer, G. D. Toderici, M. Tschannen, E. Augustsson, "High-fidelity generative image compression", Advances in Neural Information Processing Systems, Vol. 33, 2020, pp. 11913-11924.

[4] R. Erol, "Case Studies for Energy Efficient Machine Learning Inference Acceleration", University of Arkansas at Little Rock, AR, USA, PhD Thesis, 2024.

[5] Y. Dong, J. Tao, Y. Zhang, W. Lin, J. Ai, "Deep learning in aircraft design, dynamics, and control: Review and prospects", IEEE Transactions on Aerospace and Electronic Systems, Vol. 57, No. 4, 2021, pp. 2346-2368.

[6] Z. Wang, T. Luo, R. S. M. Goh, J. T. Zhou, "EDCompress: Energy-aware model compression for dataflows", IEEE Transactions on Neural Networks and Learning Systems, Vol. 35, No. 1, 2022, pp. 208-220.

[7] G. Hinton, "Distilling the Knowledge in a Neural Network", arXiv:1503.02531, 2015.

[8] J. Gou, B. Yu, S. J. Maybank, D. Tao, "Knowledge distillation: A survey", International Journal of Computer Vision, Vol. 129, No. 6, 2021, pp. 1789-1819.

[9] D. Saxena, J. Cao, J. Xu, T. Kulshrestha, "RG-GAN: Dynamic Regenerative Pruning for Data-Efficient Generative Adversarial Networks", Proceedings of the AAAI Conference on Artificial Intelligence, Vancouver, BC, Canada, 20-27 February 2024, pp. 4704-4712.

[10] D. Saxena, J. Cao, J. Xu, T. Kulshrestha, "Re-GAN: Data-efficient GANs training via architectural reconfiguration", Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Vancouver, BC, Canada, 17-24 June 2023, pp. 16230-16240.

[11] N. K. Jayakodi, J. R. Doppa, P. P. Pande, "SETGAN: Scale and Energy Trade-off GANs for Image Applications on Mobile Platforms", arXiv:2103.12896, 2021.

[12] D. Wu, K. Hur, Z. Xiao, "A GAN-enhanced ensemble model for energy consumption forecasting in large commercial buildings", IEEE Access, Vol. 9, 2021, pp. 158820-158830.

[13] Y. Du, I. Mordatch, "Implicit generation and modeling with energy-based models", Advances in Neural Information Processing Systems, Vol. 32, 2019.

[14] G. Zhong, J. Wang, J. Hu, F. Liang, "A GAN-based video intra coding", Electronics, Vol. 10, No. 2, 2021, p. 132.

[15] K. Wu, Z. Li, Y. Yang, Q. Liu, "Deep video compression based on long-range temporal context learning", Computer Vision and Image Understanding, Vol. 248, 2024, p. 104127.

[16] L. Helminger, R. Azevedo, A. Djelouah, M. Gross, C. Schroers, "Microdosing: Knowledge distillation for GAN-based compression", arXiv:2201.02624, 2022.

[17] I. Goodfellow et al. "Generative adversarial nets", Advances in Neural Information Processing Systems, Vol. 27, 2014.

[18] H. Hardi, I. F. E. Fatani, "Optimizing Video Compression Quality Using AI-Boosted HEVC", Proceedings of the 6th International Conference on Intelligent Computing in Data Sciences, 2024, pp. 1-4.

[19] I. Goodfellow et al. "Generative adversarial networks", Communications of the ACM, Vol. 63, No. 11, 2020, pp. 139-144.

[20] A. Katsenou, F. Zhang, M. Afonso, G. Dimitrov, D. R. Bull, "BVI-CC: A dataset for research on video compression and quality assessment", Frontiers in Signal Processing, Vol. 2, 2022, p. 874200.