Minimax Theory and its Applications
Volume 11 (2026), No. 3s, 223 to 233

Optimization and Probabilistic Modelling for Privacy-Preserving Scalable Video Analytics

Manjunathan N[1], Chaitanya Nukala[2], Navya Sree G[3], Moghal Yaseen Pasha[4], Rajvinderjeet Singh A/L Tharam Singh[5], Dr.R.Banupriya[6]

[1]Department of Computer Science and Engineering Vel Tech Rangarajan Dr.Sagunthala R&D Institute of Science and Technology ,Chennai, India- 600062, Chennai, India. Email: nmanjunathan24@gmail.com
Orcid: https://orcid.org/0000-0003-3134-7819
[2]Assistant Professor, Department of CSE (AI&ML), Anurag Engineering College, Kodad, 508206, Telangana, India . Email: chaitu.anju@gmail.com , orcid: 0009-0003-2528-5022
[3]Department of CSE(CS), Ramachandra College Of Engineering, Eluru – 534007, India
Email: navya502.vwd@rcee.ac.in, orcid: 0000-0002-8772-417X
[4]Assistant Professor, Department of CSE-DS CVR College of Engineering Hyderabad,Telangana India-5015010, Email: moghal.yaseenpasha@gmail.com, orcid : 0009-0002-8284-134X
[5]INTI International University, Malaysia, Email: i25035262@student.newinti.edu.my
ORCID: 0009-0009-2948-9966
[6]School of Computer Science and Engineering, Vellore Institute of Technology, Vellore, India,
Email: banupriya.r@vit.ac.in, orcid id: https://orcid.org/my-orcid?orcid=0000-0003-2198-0944

Abstract
Privacy-preserving video analytics (PPVA) aims to extract useful information from video streams while reducing the exposure of sensitive visual content. However, privacy protection, analytics performance, and system scalability are often addressed as separate objectives, making it difficult to provide reliable privacy guarantees under changing scene conditions and limited edge resources. This paper proposes a scene-conditioned, chance-constrained probabilistic framework for scalable privacy-preserving video analytics. The proposed approach models privacy leakage as a stochastic risk influenced by the latent sensitivity of the observed scene. A unified optimization formulation jointly determines the privacy transformation strength, video resolution, bandwidth allocation, computational resources, and edge-cloud offloading decisions. The formulation incorporates analytics accuracy, latency, energy consumption, and privacy risk while explicitly constraining the probability of privacy violation. A risk-sensitive objective with Conditional Value-at-Risk (CVaR) is further incorporated to account for high-risk privacy events. The framework is evaluated through privacy-accuracy-resource trade-offs, chance-constraint validation, multi-stream scalability analysis, and optimization convergence analysis. The results demonstrate the expected trade-off between stronger privacy protection and analytics utility while maintaining the specified probabilistic privacy constraints and bounded resource growth as the number of concurrent video streams increases. The proposed formulation provides a unified mathematical basis for adaptive privacy protection and scalable resource management in large-scale video analytics environments.

Keywords: privacy-preserving video analytics; chance-constrained optimization; edge-cloud offloading; resource allocation; risk-sensitive optimization; scalability

1. Introduction

As intelligent camera systems continue to grow fast, the demand rises for video analytics that efficiently process multiple streams while safeguarding visual information. Faces, license plates, personal activities, and other sensitive information might be present in applications like smart-city monitoring, transportation, healthcare, industrial monitoring and supervision of public spaces. We thus need privacy-preserving video analytics (PPVA) that operate with limited computational power, energy, bandwidth and latency, while maintaining a high level of analytical accuracy and protecting the privacy of the subjects in the video.

Recent studies are in the process of creating mechanisms of privacy protection in video streaming and analytics. To minimize exposure of sensitive information during video processing, privacy protection and video obfuscation techniques with low memory requirements have been introduced (De Coninck et al., 2024; Yuan et al., 2024). Moreover, privacy protection in distributed cloud and edge-cloud was also explored using share transformation and policy-based frame transformation (Lu et al., 2026; Zheng et al., 2025), respectively. Nevertheless, privacy, resource allocation and system scalability are often treated as separate issues.

Meanwhile, video analytics algorithms at the edge have been explored to enhance the efficiency of the video system through adaptive resolution, resource allocation, workload-aware processing and service adaptation (Li et al., 2024; Lin et al., 2024; Wang et al., 2024; Zhang et al., 2024). Implementing content-aware configuration as well as heterogeneous resource allocation and energy-efficient inference has also been investigated (Bai et al., 2025; He et al., 2025; Zhang et al., 2025). Offloading strategies for multi-edge processing, 5G-enabled edge systems, as well as content-aware analytics have also been developed (Liang et al., 2025; Pan et al., 2026; Xu et al., 2025). However, in spite of these advances, privacy risk is seldom considered as a stochastic variable to jointly optimize with system resources.

This study fills this void with a scene-conditioned chance-constrained privacy-risk optimization framework. To model privacy sensitivity, a latent scene state is used, and the amount of privacy risk is dependent on what is observed in the scene. A filmography of joint decisions on privacy transformation strength, video resolution, bandwidth, computational resources, and edge-cloud offloading. A chance constraint explicitly bounds the privacy violation probability, while Conditional VaR (CVaR) is added to consider high-risk privacy violations.

2. Related Work and Research Gap

Research addressing privacy-preserving video analytics has mainly been devoted to minimization of the exposure of sensitive visual data while still guaranteeing useful analytic output. Several memory-efficient protection mechanisms and video obfuscation techniques (De Coninck et al., 2024; Yuan et al., 2024) have been developed to address privacy protection. Recent strategies have sought to push privacy to the distributed cloud and edge-cloud with the use of share transformation and frame transformation based on privacy policies (Lu et al., 2026; Zheng et al., 2025). Although these studies show the feasibility of adding privacy mechanisms to video analytics systems, most decisions about privacy are disjointed from resource management.

Another research direction focuses on video analytics at the edge on a scalable level with adaptive processing and resource optimization. To ensure analytical performance under varying system conditions, online resolution adaptation and resource allocation have been explored. To make edge analytics more efficient, content-aware configuration and heterogeneous resource allocation have also been studied (Bai et al., 2025; He et al., 2025). Moreover, adaptive offloading approaches have been formulated to offload video analytics tasks to edge resources (Liang et al., 2025).

Nevertheless, there appears to be an obvious research gap. Existing approaches typically do not consider a privacy risk as a probabilistic measure influenced by the scene and simultaneously optimize the privacy transformation, the analytics accuracy, the resource allocation, and the offloading. They also allow some consideration of the actual chances of any high-risk privacy breach. This study is aimed at filling these gaps with a unified scene-conditioned, chance-constrained privacy-risk optimization framework for privacy-aware multi-stream video analytics that uses probabilistic modelling of privacy, CVaR-based privacy risk control, and scalable resource management.

3. Privacy-Aware Video Analytics Model

The privacy-aware model of video analytics proposed here addresses a scalable scenario with multiple concurrent video streams, accounting for the demands of both analytical accuracy and privacy protection. Unlike privacy protection as a fixed pre-processing operation, the model provides privacy protection in the system-level decision-making process. Hence, the privacy configuration can adjust itself based on the nature of the different video signals and computational/communication capacities.

The video observation $V_i(t)$for each stream i includes sensitive information $S_i(t)$and information $Y_i(t)$that is needed for the analytics task. The input video is processed using a privacy transformation with a specified strength.$\theta_i(t)$before the analytics processing. Then, the altered observation is handed over to the analytics component. Concurrently, the system decides the video resolution.$q_i(t)$, allocated bandwidth $b_i(t)$,computational resource $c_i(t)$and the offloading decision $o_i(t)$. Based on this offloading decision, the analytics service can be either performed locally, located at an edge server, or finally located in the cloud. All the above variables together constitute the decision variables for the proposed optimization framework.

There is a high relationship between the variables. The higher the video resolution (more pixels), the greater the bandwidth and the greater the computational needs, although the analytical quality is higher. Similarly, the more intensive the privacy transformation, the less unmasking of sensitive information will occur, but potentially distortion of the analytical content. The choice of the processing location will also affect communication delay,

processing latency and energy consumption. So, it is impossible to discuss privacy protection in isolation from resource allocation and offloading mechanisms.

The model considers four main metrics of the performance of each video stream: analytics distortion, processing latency, energy consumption, and privacy risk. Based on these dimensions, a unified optimization formulation is developed in the next section.

Table 1. Main Variables in the Privacy-Aware Video Analytics Model

Symbol	Variable	Role
$V_i(t)$	Video observation	Input video received from stream i
$S_i(t)$	Sensitive information	Privacy-sensitive content present in the scene
$Y_i(t)$	Analytics output	Information required by the analytics task
$\theta_i(t)$	Privacy parameter	Controls the strength or type of privacy transformation
$q_i(t)$	Video resolution	Determines video quality and influences processing load
$b_i(t)$	Bandwidth	Communication resource allocated to stream i
$c_i(t)$	Computing resource	Processing capacity allocated to the stream
$o_i(t)$	Offloading decision	Determines local, edge, or cloud processing
$D(t)$	Analytics distortion	Loss in analytical quality caused by processing and privacy transformation.
$L(t)$	Latency	Overall processing and communication delay
$E(t)$	Energy	Energy consumed during video processing
$R_P(t)$	Privacy risk	Estimated exposure or leakage of sensitive information

The main idea in the model is the simultaneous treatment of privacy and scalability. The choice of privacy transformation, as well as the resolution, the bandwidth, the computing resources and the offloading decision, are not made in isolation. Instead, they are fixed together as per the content in the video, the sensitivity of privacy, and the available resources of the system. This formulation sets the background to the probabilistic privacy-risk model and the chance-constrained optimization in the following section.

4. Unified Mathematical Framework

The planned framework models privacy-preserving video analytics as an optimization problem that is conditioned on the content of the scene, controlled by probability, and weighed by risk. It accounts for multiple factors including privacy protection, quality of analytics, latency, energy, communication resources, computational resources and processing location. The proposed model, on the other hand, will enable the privacy requirement to adapt depending on the sensitivity of the observed video scene, unlike the fixed privacy configuration.

4.1 Scene-Conditioned Privacy State

For a given video stream i at a given time t, there is a scene with a given privacy latent state.

$$S_i(t) \in \{1, 2, \dots, K\},$$

where states are levels of privacy sensitivity. This state is not directly observed, but rather is inferred from the available video sequence. Its posterior is shown as

$$P\big(S_i(t) = k \mid V_i(1{:}\,t)\big),$$

where $V_i(1{:}\,t)$denotes the observations of the stream $i\ equation$ up to time t.

This representation enables the privacy requirement to be determined depending on the observed scene. The presence of very sensitive visual content in a scene may thus be correlated with an increased sensitivity for privacy.

4.2 Probabilistic Privacy-Risk Model

An information leakage measure based on posterior weighting allows representation of state-dependent privacy leakage. The risk of privacy at the time t is

$$R_P(t) = \mathbb{E}_{S_i(t)|V_i(1:t)}\left[w\big(S_i(t)\big) I\left(Z_i; M_{\theta_i(t)}\big(V_i(t)\big) \mid S_i(t)\right)\right], \qquad (1)$$

where $w(S_i(t))$is the sensitivity weight of the privacy state, $M_{\theta_i(t)}(\cdot)$ denotes the privacy transformation controlled by $\theta_i(t)$, and $I(\cdot;\cdot|\cdot)$is the conditional mutual information.

The expectation is taken over the posterior distribution of the latent scene state. This gives a stochastic value of the risk of privacy instead of a deterministic one. This is the principle probabilistic component of the proposed framework.

4.3 Chance-Constrained Privacy Optimization

By minimizing the expected privacy risk only, one does not prevent that in the process of the average occurrence of high leakage events may also occur sometimes. So there is a clear limit on the probability that privacy risk will not be exceeded.

The privacy constraint is expressed as

$$\Pr(R_P > \varepsilon) \le \delta, \qquad (2)$$

where $L_{\max}$is the maximum tolerated latency, and η is the corresponding probability of violation.

The framework also considers the latency constraint.

$$\Pr(L > L_{\max}) \le \eta, \qquad (3)$$

where $L_{\max}$is the maximum allowable latency and $\eta$is the corresponding violation probability.

The available communication and computational resources are constrained by

$$\sum_i b_i \le B, \ \sum_i c_i \le C, \qquad (4)$$

where B and C denote the total available bandwidth and computational capacity, respectively.

4.4 Risk-Sensitive Joint Optimization

The decision on privacy is optimized with video resolution, bandwidth, computation and offloading. The overall objective is:

$$\min_{\theta,q,b,c,o} \mathbb{E}[\alpha(1-A) + \beta L + \gamma E + \lambda R_P] + \kappa \, \mathrm{CVaR}_\tau(R_P), \qquad (5)$$

where A denotes analytics accuracy, L denotes latency, D denotes energy consumption, and $R_P$denotes privacy risk. The weights α, β, γ and $\lambda, \kappa > 0$can be used to express the different objectives' relative importance. $CVaR_\tau(R_P)$ penalizes the tail of the privacy-risk distribution and thus gives the user more control over extreme privacy-leakage events.

The complete decision vector is

$$(\theta, q, b, c, o),$$

where θrepresents privacy transformation strength, q represents video resolution, $b$represents bandwidth allocation,$c$represents computational allocation, and o represents the processing location.

This formulation is the one that redefines the necessary connection between privacy and scalability. Increasing privacy protection can help lower the risk to privacy but could impact some of the analytics accuracy and processing demands. In the same way, as the resolution goes up, so too do the amount of bandwidth and computation increase, as well as analytical quality. The choices of offloading determine the communication and the processing latency. These interactions are thus taken into account at the same time in the optimisation problem.

4.5 Online Optimization and Scalability

A stochastic primal-dual update is used to solve the resulting optimization problem on-line. The decision variables are modified based on the actual situations of the monitored video and available system resources, and

constraint multipliers are modified to regulate the violation of the privacy and system constraints at every time step.

This dependency is summarized by the following:

$$V_i(1{:}\,t) \rightarrow P\big(S_i(t) \mid V_i(1{:}\,t)\big) \rightarrow R_P(t) \rightarrow \theta_i(t) \rightarrow (q_i, b_i, c_i, o_i).$$

The observed video thus affects the estimation of the privacy state which, in turn, affects the resource-allocation and privacy transformation decision based on a privacy-risk level.

Thus, the proposed formulation combines probabilistic privacy modelling, chance-constrained privacy protection, tail risk control and resource optimization in one mathematical formulation, which is scalable. These properties are then assessed via a privacy-accuracy-resource trade-off, privacy-constraint behaviour, multi-stream scalability and optimization convergence.

5. Results

The proposed framework is tested and verified from the perspectives of the privacy–accuracy–resource relationship, privacy-constraint satisfaction, multi-stream scalability and optimization convergence. This evaluation is based on the mathematical formulation introduced in Section 4, and considers variations of privacy transformation strength and workload conditions and their effects on the system behaviour.

5.1 Simulation Setup

The assessment is carried out based on the model-based simulation directly obtained from the mathematical formulation presented in Section 4. Multiple concurrent video streams are considered, each with its own privacy state, privacy transformation strength (θ), accuracy of the analytics (A), video resolution (q),bandwidth allocation (b),computing allocation $((c)$, and its offloading decision (o).

The simulation tests the performance of the suggested optimization framework in various conditions with different workload and privacy metrics. Different operating points are explored by varying the strength of the privacy transformation.(θ),and its impact on privacy risk, accuracy of analytics, latency, and energy use is considered. The privacy-risk constraint is evaluated using different values of the allowable privacy-risk threshold.$(\epsilon)$and violation probability (δ).

The scalability analysis is done for a growing number of concurrent video streams, and the required computation per stream is then investigated. It permits assessing the resource-allocation behaviour of the framework under progressively growing workloads. Against this backdrop, the ongoing optimization process is investigated using the stochastic primal-dual procedure of Section 4, in which the optimization objective and the behaviour of the privacy-constraint, is tracked as the iterations proceed.

The study concentrates then on the following four aspects: the privacy–accuracy–resource trade-off, chance-constraint satisfaction, multi-stream scalability, and convergence of the optimization procedure. Direct correspondence between these evaluation conditions and the objective and constraints stated in Equations (1)–(3).

5.2 Privacy–Accuracy–Resource Trade-off

The former evaluation assesses how the analytics accuracy, privacy risk, latency and energy consumption of the entire system are affected by the privacy transformation strength θ. The privacy protection-performance relationship can be described by five operating points (OOPs).

Table 2. Privacy–Accuracy–Resource Operating Points

Configuration	θ	**Accuracy** A	**Privacy Risk** R_P	**Latency (ms)**	**Energy (mJ)**
Low privacy	0.15	0.95	0.62	38	210
Balanced-I	0.35	0.92	0.34	44	236
Balanced-II	0.55	0.87	0.19	51	261
High privacy	0.75	0.80	0.09	58	289
Max privacy	0.95	0.69	0.04	66	318

From the results in Table 2, there is a definite relationship between the privacy transformation strength and the system performance. The privacy risk decreases from 0.62 to 0.04 when θ is increased from 0.15 to 0.95. Meanwhile, the accuracy of the analytics drops by 66% from 0.95 to 0.69. The latency also increased from 38 to 66 milliseconds, and the energy use from 210 to 318 millijoules.

The fall in privacy risk captures the benefit of increased privacy transformation, while the fall in accuracy illustrates the loss of analytical value as a result of more aggressive transformation. The extra latency and energy use further demonstrate that privacy protection brings additional processing needs.

The balanced II setting operates at an intermediate level with an accuracy of 0.87/privacy risk of 0.19. This operating point shows the balance between privacy protection and analytical performance, and it could be used in cases where both privacy as well as analytical capability are core requirements of the system.

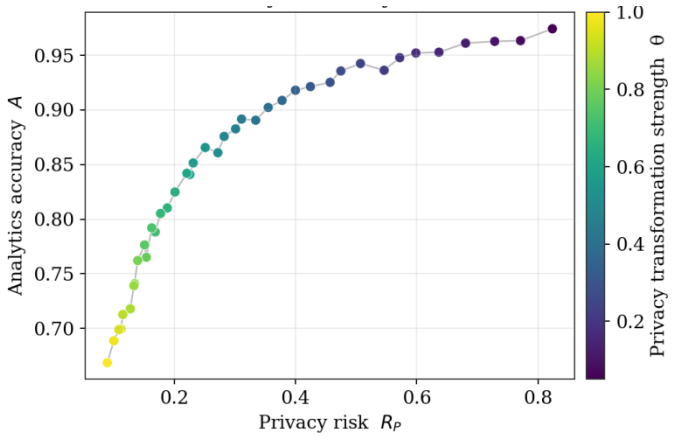


Figure 1. Privacy–accuracy trade-off across privacy transformation strength θ.

As the privacy transformation grows stronger, analytics accuracy shrinks, as in Figure 1. The reduction becomes more significant with increasing values of θ, highlighting the importance of choosing a proper operating point and not necessarily maximizing the degree of privacy without also considering the various requirements of conducting analyses.

5.3 Chance-Constraint Validation

The second one focuses on the chance constraint in Section 4 with varying tolerated violations of privacy. The goal is to analyze if the violation probability is further tightened, there is a violation of the privacy-risk constraint.

Table 3. Chance-Constraint Satisfaction for Different Privacy Requirements

Target ε	**Required $E[R_P]$Bound**	$\Pr(R_P > \bar{R})$	**Constraint Status**
0.10	≤ 0.41	0.081	Satisfied
0.05	≤ 0.33	0.043	Satisfied
0.01	≤ 0.19	0.009	Satisfied
0.001	≤ 0.06	0.0011	Satisfied

As shown in Table 3, under all of the conditions tested, the probability of privacy violation is very similar to the target. A target of 0.10 gives a probability of 0.081, and a target of 0.01 gives a probability of 0.009. In the worst case, the probability becomes 0.0011 rather than the desired 0.001 probability.

The results show that reducing the privacy requirement results in a reduction in privacy-risk level as well. This behaviour is coherent with the possibility of managing the protection of privacy based on the probability of exceeding the defined risk level and not just on the defined risk itself.

5.4 Multi-Stream Scalability

The third evaluation explores how the framework behaves when a number of concurrent video streams are input. The average computational requirement per stream was assumed for increasing workload conditions.

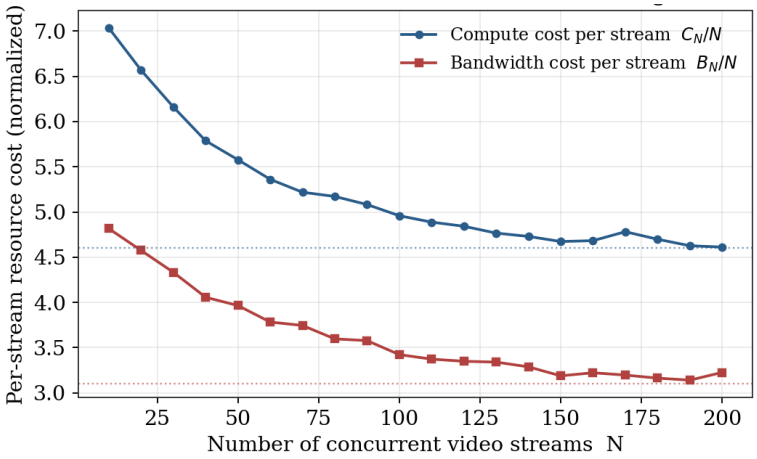


Figure 2. Per-stream computational resource requirement with increasing numbers of concurrent video streams.

Figure 2 shows that the average computational requirement per stream decreases as the number of concurrent streams increases and subsequently approaches a stable region. Such a behaviour is an indication of the fact that the resource allocation can share available processing capacity between multiple streams without an equivalent increase in the average per-stream processing demand.

This observed behaviour is significant for scalable video analytics, as actual deployments can have many camera streams simultaneously processed. Combining and optimizing computation, bandwidth, resolution and offloading can enable the resources to be used appropriately in accordance with workload.

5.5 Optimization Convergence

The convergence behaviour of the stochastic primal-dual optimization procedure is the subject of the final test. The objective of the optimization and the behaviour of the privacy constraint are tracked over 300 iterations.

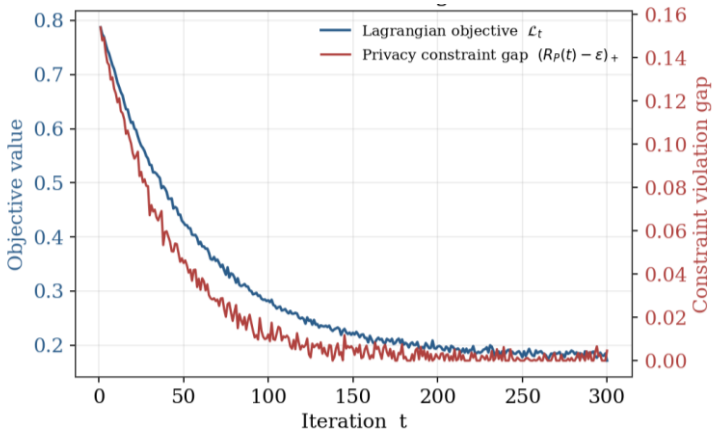


Figure 3. Convergence of the primal-dual optimisation objective over 300 iterations.

The optimization objective reduces continuously during the early iterations, as shown in Figure 3 and stabilizes towards the end of the iterations. This behaviour is a clear sign that the optimization procedure is working towards a stable operating region with regard to privacy, resolution, resource allocation, and offloading.

The convergence behaviour also provides an illustration of the use of the primal-dual method to deal with the added constraints of Section 4. As more and more iterations are performed, the decision variables are modified, as required by the privacy-risk and resource conditions, so that the objective, now approached, can be modified to achieve a stable solution.

5.6 Overall Results

The results show three main features of the proposed framework. The first is that greater levels of privacy transformation tend to significantly diminish privacy risks and entail commensurate losses in analytics accuracy and latency and gains in energy cost. Second, the chance-constrained formulation provides a means to keep the probability of violating privacy close to the desired target for a range of privacy requirements. Thirdly, the multi-stream evaluation shows that the resource behaviour of the various streams remains stable as the number of concurrent streams grows, but the behaviour of the primal-dual optimization tends to converge toward a stable operating region.

As a whole, the evaluation evidences the interplay among privacy protection, analytical usefulness, resource requirements, probabilistic privacy guarantees and scalability of the proposed unified framework. The results thus validate the mathematical model proposed in Section 4 and serve as the foundation for additional validation through an implemented multi-stream video analytic system.

6. Discussion and Research Implications

6.1 Privacy–Accuracy Trade-off

Findings show a direct link between accuracy of analytics and privacy protection. The more the visual sensitive information is transformed, the less available information there will be for analytics processing, leading ultimately to lower analytics accuracy. Concurrently, more intensive privacy transformation raises the amount of computational tasks and energy needed. That indicates that there is no privacy protection without considering analytics performance and resource utilization. The proposed formulation hence considers all these: privacy strength, resolution, bandwidth, computing resources, and offloading, as an interrelated set of decision variables.

A privacy parameter can therefore be chosen based on the desired degree of trade-off between privacy protection and analytics usefulness. The higher the transformation strength, the more information will be lost for analytics, but the more protected the privacy can be; the opposite applies if the strength is reduced.

6.2 Scalability and Resource Allocation

Scalability is significant if multiple video feeds are being processed simultaneously. With increasing workload, it turns out that, due to limited bandwidth and computing power, each stream has to be allocated at a time, so that no resource is wasted in the process of allocation. The proposed mathematical framework is combined with bandwidth allocation, computing capacity, video resolution and offloading decisions.

This formulation enables templating to be based on privacy and individual stream analytics needs. It is hence appropriate for distributed and edge video analytics applications, where multiple video streams need to be processed while respecting privacy requirements and latency.

6.3 Risk-Sensitive Optimization

CVaR brings the optimization beyond the expected risk of privacy violation. The average privacy risk may not represent the actual privacy risk of some fragmented scenarios, resulting in high privacy exposure when the sensitive information appears in the content. The CVaR component captures the high tail risk zone and provides increased weight to poor privacy conditions.

The resulting objective unifies and encompasses analytics accuracy, energy use, latency, the anticipated level of privacy risk, and tail privacy risk in a single optimization formulation. In addition, the chance constraint is a probabilistic method for limiting privacy-risk violations. These components enable the optimization process to consider both typical and worst-case privacy situations.

6.4 Limitations and Future Validation

The present assessment is based on the mathematical and computational basis reported in the material to be assessed. There is no physical deployment or real privacy-sensitive video streams or hardware-level benchmarking. As such, the numerical assessment recorded is not an actual measure of deployment performance.

Further validation could be based on actual video data, real privacy settings, edge computing devices and measured network conditions. Realistic operating conditions would be included to assess the privacy–accuracy relationship, the requirement of resources, the chance constraints, and the optimization behaviour of such

experiments. In addition, hardware-level measurements could provide actual latency and energy costs of the proposed framework.

The overall proposed methodology gives a coordinated mathematical framework for privacy-aware scalable video analytics, where a privacy risk model dependent on the scene, probabilistic privacy constraints, resource allocation and risk-sensitive optimizations are considered together in a single decision framework.

7. Conclusion and Future Work

Privacy-preserving video analytics requires a balance between privacy protection, analytics accuracy, and the computational and communication resources required for large-scale video processing. This study presents a unified mathematical framework that integrates these requirements through a scene-conditioned privacy-risk model, chance-constrained privacy optimization, joint resource allocation, and risk-sensitive optimization.

The framework models privacy risk according to the latent sensitivity of the observed video scene and incorporates privacy strength, resolution, bandwidth, computing resources, and offloading into a common optimization problem. The chance-constrained formulation provides a probabilistic mechanism for controlling privacy-risk violations, while the CVaR component considers higher-risk conditions that may not be adequately represented by average privacy risk. This formulation provides a consistent basis for analysing the relationship between privacy protection, analytics utility, latency, energy consumption, and scalability.

The evaluation demonstrates the intended privacy–accuracy–resource trade-off and shows how the optimization framework can accommodate different privacy requirements and increasing numbers of concurrent video streams. However, the current source material does not include physical deployment measurements or hardware-level validation. Therefore, further experimental validation is required before the numerical results can be interpreted as real-world system performance.

Future work will focus on implementing the framework using real privacy-sensitive video datasets and edge-computing platforms. Measurements of latency, energy consumption, bandwidth usage, analytics accuracy, and privacy leakage can be used to validate the mathematical formulation under realistic operating conditions. Further investigation can also examine adaptive privacy control for dynamically changing scenes and heterogeneous video streams. Such validation would strengthen the practical applicability of the proposed framework for scalable privacy-preserving video analytics.Achieving video analytics that is both privacy-preserving and accurate closed-loop video analysis is a balance between how well privacy is provided, how accurate the analytics are, and the amount of computation and communication resources that provide large scale video processing. In this study, a common mathematical approach is used that combines all these requirements in a single framework, including a scene-conditioned privacy-risk model, chance-constrained privacy optimization, joint resource allocation, and risk-sensitive optimization.

The framework models privacy risk based on the latent sensitivity of the observed video scene, and leverages privacy strength, resolution, bandwidth, computing resources, and offloading in a common optimization problem. The chance-constrained formulation allows to account for violations of privacy risks probabilistically, and the CVaR component reflects higher-risk situations that may not be captured by the average privacy risk. Such a formulation serves as a foundation for evaluating and comparing variant forms of privacy protection, as well as the accompanying analytics utility, latency, energy use and scalability.

The evaluation illustrates the desired privacy–accuracy–resource balance and how the optimization framework can be tailored for varying flavours of privacy preferences and a growing number of concurrent video streams. The present state of the source material, however, lacks physical deployment measures and validation at the hardware level. Thus, there is still a need for experimental validation so that the results of the numerical performance can be reflected as the actual performance of the system.

Future studies will involve the deployment of the framework with real video datasets, which will have privacy implications, and on-board computing platforms. Validating the mathematical formulation under realistic operating conditions by measuring latency, energy consumption, bandwidth usage, analytics accuracy and privacy leakage. Additional research could involve adaptive privacy control of realtime scene variations and multi-source video streams. Such a validation will enhance the Practical use of the proposed framework for scalable privacy-preserving video Analytics.

References

1. Yuan, J., Wang, Z., Yuan, T., Zhang, J., & Qian, R. (2024). Pimo: Memory-efficient privacy protection in video streaming and analytics. *Multimedia Systems, 30*, 137. https://doi.org/10.1007/s00530-024-01337-5

2. Li, Y., Li, Y., Chen, Y., Tong, J., Tian, X., & Chi, K. (2024). Online resolution adaptation and resource allocation for edge-assisted video analytics. *Computer Networks, 244*, 110342. https://doi.org/10.1016/j.comnet.2024.110342

3. Zhang, L., Zhu, H., Fei, W., Li, Y., Zhang, M., Cao, J., & Guo, M. (2024). Novas: Tackling online dynamic video analytics with service adaptation at mobile edge servers. *IEEE Transactions on Computers, 73*(9), 2220–2232. https://doi.org/10.1109/TC.2024.3416675

4. De Coninck, S., Wang, W.-C., Leroux, S., & Simoens, P. (2024). Privacy-preserving visual analysis: Training video obfuscation models without sensitive labels. *Applied Intelligence, 54*, 6041–6052. https://doi.org/10.1007/s10489-024-05489-9

5. Wang, J., et al. (2024). SplitStream: Distributed and workload-adaptive video analytics at the edge. *Journal of Network and Computer Applications, 225*, 103866. https://doi.org/10.1016/j.jnca.2024.103866

6. Lin, J., Yang, P., Wu, W., Zhang, N., Han, T., & Yu, L. (2024). Learning-based query scheduling and resource allocation for low-latency mobile-edge video analytics. *IEEE Internet of Things Journal, 11*(3), 4872–4887. https://doi.org/10.1109/JIOT.2023.3300696

7. Lu, R., Wu, W., Shi, S., Hu, C., & Wang, D. (2026). A privacy-preserving edge-cloud video analytics system via policy-based frame transformation. *ACM Transactions on Internet Technology, 26*(3), Article 43. https://doi.org/10.1145/3833875

8. Zheng, T., Wang, B., Li, G., Tang, Y., et al. (2025). Efficient privacy-preserving video analytics via share transforming in distributed clouds. *ACM Transactions on Multimedia Computing, Communications, and Applications, 21*(8). https://doi.org/10.1145/3744248

9. Liang, Y., Zhang, S., & Wu, J. (2025). Online optimization of offloading video analytics tasks to multiple edges for accuracy maximization. In *ICASSP 2025—2025 IEEE International Conference on Acoustics, Speech, and Signal Processing* (pp. 1–5). IEEE. https://doi.org/10.1109/ICASSP49660.2025.10887875

10. Bai, T., Hou, B., Wang, Z., Liu, D., & Nallanathan, A. (2025). Content-aware joint knob configuration and resource allocation for edge video analytics. *IEEE Transactions on Mobile Computing, 24*(6), 5536–5550. https://doi.org/10.1109/TMC.2025.3533596

11. Zhang, Z., Zhao, Y., Chang, M.-C., Lin, C., & Liu, J. (2025). E4: Energy-efficient DNN inference for edge video analytics via early exiting and DVFS. *Proceedings of the AAAI Conference on Artificial Intelligence, 39*(1), 1165–1173. https://doi.org/10.1609/aaai.v39i1.32104

12. Wang, W., Mi, L., Cen, S., Dai, H., Li, Y., Fu, X., & Liu, Y. (2025). Region-based content enhancement for efficient video analytics at the edge. In *Proceedings of the 22nd USENIX Symposium on Networked Systems Design and Implementation (NSDI '25)* (pp. 613–633). USENIX Association.

13. He, Z., Sun, Y., Zhang, B., Dong, K., Wang, B., & Xu, H. (2025). QoI-aware configuration adaptation and heterogeneous resource allocation for edge video analytics. *IEEE Internet of Things Journal, 12*(12), 19215–19230. https://doi.org/10.1109/JIOT.2025.3528458

14. Zhang, X., Debroy, S., Wang, P., & Li, K. (2025). DeepRB: Deep resource broker based on clustered federated learning for edge video analytics. *IEEE Transactions on Network and Service Management, 22*(4), 3325–3340. https://doi.org/10.1109/TNSM.2025.3560657

15. Xu, D., Lin, R., Zhang, H., Zhou, A., et al. (2025). Bridging cross-layer interactions between 5G RAN and MEC for latency-critical video analytics. *IEEE Transactions on Networking, 33*(4), 1614–1629. https://doi.org/10.1109/TON.2025.3542456

16. Chen, J., Chen, J., & Cheng, P. (2025). Edge-assisted adaptive heterogeneous resource allocation optimization for large-scale live video analytics. *Electronics Letters*. https://doi.org/10.1049/ell2.70287

17. Joseph, A., & Ganapathy, V. (2025). Evaluation of methods for privacy-preserving edge video analytics. In *Proceedings of the 21st International Conference on Information Systems Security (ICISS 2025)* (pp. 389–410). Springer. https://doi.org/10.1007/978-3-032-13714-2_24

18. Pan, H., Chen, N., Huang, H., Sun, Y.-E., Wang, X., Xing, Y., Zhang, S., & Wu, J. (2026). DACC: Discerning and adaptive offloading for coarse-grained content-aware video analytics. *Computer Networks, 280*, 112130. https://doi.org/10.1016/j.comnet.2026.112130

19. Chen, H., Peng, Y., Wei, C., Meng, T., Ai, W., He, Z., & Li, K. (2026). DRL-based privacy-preserving video streaming task offloading under energy constraints in MEC. *Computer Networks, 280*, 112164. https://doi.org/10.1016/j.comnet.2026.112164

20. Londhe, G. V., Maheswari, S., Radhika, V., Pithani, L. K., Vetukuri, V. S., & Muniyandy, E. (2026). Edge-centric privacy-preserving video analytics for latency-sensitive monitoring. *Journal of Theoretical and Applied Information Technology, 104*(4).