

Enhancing Risk Assessment in Transformers with Loss-at-Risk Functions

1st Jinghan Zhang
Portland State University
Portland, USA
jinghanz@pdx.edu

2nd Henry Xie
Westview High School
Portland, USA
henryjxie@gmail.com

3rd Xinhao Zhang
Portland State University
Portland, USA
xinhaoz@pdx.edu

4th Kunpeng Liu
Portland State University
Portland, USA
kunpeng@pdx.edu

Abstract—In the financial field, precise risk assessment tools are essential for decision-making. Recent studies have challenged the notion that traditional network loss functions like Mean Square Error (MSE) are adequate, especially under extreme risk conditions that can lead to significant losses during market upheavals. Transformers and Transformer-based models are now widely used in financial forecasting according to their outstanding performance in time-series-related predictions. However, these models typically lack sensitivity to extreme risks and often underestimate great financial losses. To address this problem, we introduce a novel loss function, the Loss-at-Risk, which incorporates Value at Risk (VaR) and Conditional Value at Risk (CVaR) into Transformer models. This integration allows Transformer models to recognize potential extreme losses and further improves their capability to handle high-stakes financial decisions. Moreover, we conduct a series of experiments with highly volatile financial datasets to demonstrate that our Loss-at-Risk function improves the Transformers' risk prediction and management capabilities without compromising their decision-making accuracy or efficiency. The results demonstrate that integrating risk-aware metrics during training enhances the Transformers' risk assessment capabilities while preserving their core strengths in decision-making and reasoning across diverse scenarios.

Index Terms—Loss Functions, Transformer, Risk Assessment

I. INTRODUCTION

In the financial sector, decision support systems depend critically on precise and reliable risk assessments, especially in evaluating the maximum potential loss from a decision [1]–[5]. Transformer models, particularly Large Language Models (LLMs) and extensive time series models based on Transformer architectures have emerged as key tools in complex risk management and decision analysis. These models typically utilize standard loss functions such as Mean Squared Error (MSE) for training and fine-tuning, which ensure high accuracy in predictions and judgments [6]–[9]. However, their loss functions in training and fine-tuning process are primarily designed to minimize average prediction errors, and they often overlook the potential for extreme risk events. This oversight can lead Transformers to make overly optimistic decisions during significant market upheavals or unusual situations, potentially resulting in substantial financial losses. Although extreme risk events are infrequent in the financial sector, they can cause severe economic losses and credit risks, such as financial crises and market collapse. Thus, it is essential to

improve Transformers' awareness of risk and their risk assessment capabilities to better influence their decision-making processes.

Previous studies have addressed the issue of generalization for Transformer models handling small datasets by introducing prior knowledge and improved loss function methods [10]–[12]. For instance, some research enhances model generalization by quantifying uncertainty or using task-specific loss function variants to boost performance in particular scenarios [13]–[17]. However, these methods often fail to address the sensitivity of Transformer models to high-risk decisions directly. While these techniques can improve prediction accuracy under normal conditions, they also tend to overlook risk assessments under extreme conditions. This oversight can lead Transformers to make insufficiently cautious decisions when faced with potential risks.

These approaches primarily focus on improving models' handling of common data rather than optimizing responses to rare or extreme events [18], [19]. This limitation is particularly critical in the financial sector, where even rare market anomalies can lead to major economic disruptions. Thus, traditional methods fall short in enabling Transformers to adequately assess and manage high-risk events that could result in severe consequences, affecting the safety and reliability of decisions at crucial moments.

Therefore, we consider that during the training process of Transformers, introducing risk-aware metrics would make the model more sensitive to potential extreme losses. For Transformers, the backpropagation is an optimal stage for such adjustments because the model has already acquired a general understanding of data patterns from its initial training on large datasets [20]–[22]. During this process, incorporating specialized loss functions can refine the model's responses to be particularly cautious about rare but high-impact risks crucial in high-stakes environments like finance. This method leverages the model's ability to generalize from broad data while honing in on the nuanced understanding needed for specific tasks, which preserves the model's overall performance and enhances its precision in critical risk assessments.

Our Target. To address these problems, we target three main aspects: (1) We aim to make our Transformers more sensitive to extreme losses during the backpropagation stage. By integrating metrics into the loss function, we focus on

minimizing losses in worst-case scenarios. This approach is crucial for high-stakes applications where errors can lead to significant consequences. (2) We need to ensure these changes do not compromise the Transformer’s core strengths. This includes its time-series-related prediction, attention mechanism and robustness. The model must still excel across standard tasks while managing the demands of high-risk environments. We manage the backpropagation process to boost the Transformer’s overall abilities. (3) We then rigorously test the updated loss function in financial settings. We evaluate the model performance on financial datasets marked by high volatility and risk, to compare how the model performs before and after introducing the new loss functions. Our goal is to confirm that our changes enhance the model’s ability to predict and manage risks without losing accuracy or efficiency in decision-making.

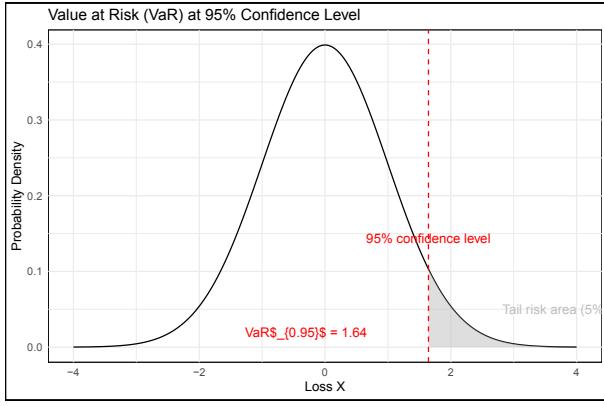


Fig. 1: A schematic diagram of VaR.

Our Method. According to the three targets, we design **Loss-at-Risk** function specifically to enhance the robustness of Transformers in high-stakes financial scenarios. Our method integrates both Value at Risk (VaR) [23]–[25] and Conditional Value at Risk (CVaR) [26]–[28] into the traditional MSE loss framework. (1) We enhance sensitivity to extreme losses. The Loss-at-Risk calculates the traditional MSE for average scenarios and adds a weighted component based on VaR and CVaR metrics. VaR provides a threshold beyond which a given percentage of losses will fall, introducing a focus on worst-case scenarios. CVaR goes further by averaging the losses that exceed the VaR threshold, ensuring the model accounts for the severity of extreme outcomes. By focusing on these metrics during backpropagation, the model becomes more sensitive to major potential losses, while ignoring common, less critical errors. (2) We also maintain core capabilities at the Transformer’s loss function. While integrating these risk metrics, we carefully calibrate their influence to maintain the Transformer’s essential skills in reasoning, decision-making, and robustness across general tasks. The weighting of the VaR and CVaR components within the Loss-at-Risk function is adjusted to ensure that risk management enhancements do not compromise the model’s effectiveness in less volatile environments. This balanced approach allows the model to

remain versatile and effective, operating reliably in typical and high-risk financial contexts. (3) Furthermore, to validate the effectiveness of our Loss-at-Risk function, we implement a rigorous testing regime using diverse financial datasets. By comparing the model’s performance on these datasets before and after applying the Loss-at-Risk function, we assess how well the model predicts and manages financial risks. The goal is to demonstrate that our approach not only preserves but enhances the model’s decision-making accuracy and efficiency under stress, and further proving that the model can handle the complexities of real-world financial applications effectively.

In summary, our contribution includes:

- 1) We integrate risk assessment with MSE loss for Transformer models’ training process to enhance their ability in extreme loss scenarios. This Loss-at-Risk equips Transformers the risk-awareness to handle high-stakes financial risks effectively.
- 2) Our method maintains the Transformer model’s core decision-making and time-series-based predicting capabilities while enhancing risk sensitivity. This ensures effective performance in both typical and high-risk financial situations.
- 3) We conduct rigorous testing on volatile financial datasets to validate our model function. The results confirm improved risk management without sacrificing decision-making accuracy or efficiency.

II. RELATED WORK

A. Transformers in Time Series Data

Transformers have revolutionized the analysis of time series data due to their ability to model complex dependencies across time steps. Their architecture, particularly effective in capturing long-range interactions in data sequences, has been extensively applied in various domains, including financial markets [29]–[34]. However, Transformers require substantial data to train effectively, which poses a challenge in scenarios where time series data is limited or highly volatile, such as in financial forecasting during unusual market conditions. The primary limitation of Transformers in these contexts is their potential overfitting to small datasets and their sensitivity to the noise commonly found in financial time series. Traditional training methods often fail to equip Transformers with the robustness needed to handle the unpredictability and sparse data points characteristic of financial time series [35]–[40].

B. Loss Function for Transformers

In research on Transformers, the predominant loss functions typically include classic forms such as cross-entropy [18], [41], [42] and MSE [43]–[46]. These are often augmented with task-specific variants to tackle challenges like data imbalance [47]–[49] or to facilitate multi-task learning [50], [51]. Our work marks a significant departure by incorporating risk assessment into the loss function during the backpropagation process, specifically enhancing the model’s ability to manage risks in scenarios where the stakes are exceptionally high. Integrating risk assessment with traditional loss functions can

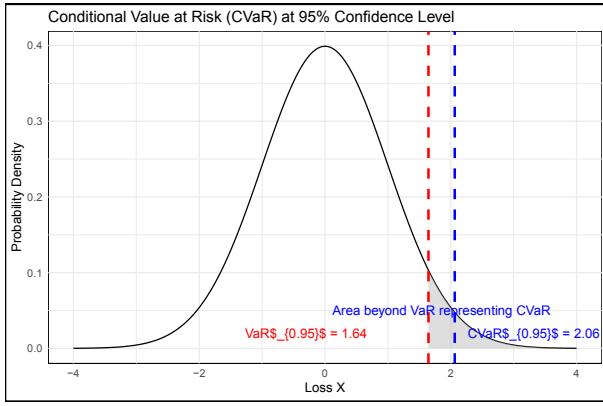


Fig. 2: A schematic diagram of CVaR.

quantify and minimize potential extreme losses. Moreover, with this integrated loss, we introduce a novel dimension to Transformer training that prioritizes risk sensitivity—an essential factor for applications in fields like finance, where accurately predicting and mitigating risks can have profound implications.

C. Value at Risk and Conditional Value at Risk

VaR is a statistical measure used to assess the level of financial risk within a firm or investment portfolio over a specific time frame, as we show in Fig. 2. It is defined as the maximum loss expected (under normal market conditions) over a target horizon within a given confidence interval. Mathematically, VaR is represented as:

$$\text{VaR}_\alpha(X) = \inf \{x \in \mathbb{R} : P(X > x) \leq 1 - \alpha\}, \quad (1)$$

where X represents the loss random variable and α is the confidence level. VaR is widely used due to its ability to provide a clear metric that quantifies potential losses at a specific confidence level.

CVaR, also known as Expected Shortfall (ES), measures the expected loss assuming that a loss is greater than the VaR. It is considered a more coherent and risk-sensitive measure than VaR because it accounts for the tail risk beyond the VaR threshold. CVaR is defined as the average of losses that exceed the VaR value at a certain confidence level:

$$\text{CVaR}_\alpha(X) = E[X | X > \text{VaR}_\alpha(X)]. \quad (2)$$

Alternatively, CVaR can be calculated through the integration of the tail distribution of losses:

$$\text{CVaR}_\alpha(X) = \frac{1}{1 - \alpha} \int_\alpha^1 \text{VaR}_u(X) du, \quad (3)$$

where u spans from the confidence level α to 1.

III. METHODOLOGY

A. Methodology Overview

In this section, we break down our approach, focusing on how we designed and tested a new loss function for LLMs, combining MSE with VaR. We keep things practical by

detailing the key math concepts and their proofs, pinpointing exactly how this helps in training LLMs. We also introduce CVaR, explaining how it builds on VaR to cover the average of more extreme losses. This part also tackles the specific challenges we faced in optimizing the model and how we solved them.

B. Loss Function Definition

a) *Loss-at-Risk with VaR*: We define our primary loss function $L(y, y_{\text{true}})$ using MSE, given by:

$$\text{MSE}(y, y_{\text{true}}) = (y - y_{\text{true}})^2, \quad (4)$$

where y is the model's output and y_{true} is the actual value. MSE is favored in regression tasks due to its emphasis on penalizing larger errors more significantly, which helps in stabilizing the training process.

For risk-sensitive applications, we incorporate the concept of VaR, defined at a confidence level α . VaR measures the maximum loss not exceeded with a probability of α , thus focusing on worst-case scenarios:

$$\text{VaR}_\alpha(L) = \inf \{\xi \in \mathbb{R} : P(L(y, y_{\text{true}}) \leq \xi) \geq \alpha\}. \quad (5)$$

To combine these measures into a single coherent loss function suitable for LLMs, we introduce the Loss-at-Risk:

$$L_{\text{VaR-MSE}}(y, y_{\text{true}}) = L_{\text{MSE}}(y, y_{\text{true}}) + \lambda \cdot \text{VaR}_\alpha(L_{\text{MSE}}), \quad (6)$$

where:

- $L_{\text{VaR-MSE}}(y, y_{\text{true}})$ is the combined MSE and VaR loss function.
- λ is a parameter that balances the impact of the MSE and VaR components.
- $\text{VaR}_\alpha(L_{\text{MSE}})$ measures the maximum loss at the confidence level α , based on the MSE calculation.

This hybrid function is designed to optimize standard performance and reduce the risk of extreme losses, allowing for adjustments based on the specific risk management needs of the application.

b) *Loss-at-Risk with CVaR*: Building on our initial Loss-at-Risk with VaR framework, we further enhance it by incorporating CVaR, defined at a confidence level α . CVaR is advantageous as it calculates the mean loss beyond the VaR threshold, offering a more comprehensive measure of tail risks:

$$\text{CVaR}_\alpha(L) = \frac{1}{1 - \alpha} \int_\alpha^1 \text{VaR}_u(L) du, \quad (7)$$

which captures the average of the worst $100(1 - \alpha)\%$ outcomes, making it a stringent measure for scenarios with potentially catastrophic financial implications.

To tailor our loss function for heightened risk sensitivity in LLMs, we propose the CVaR-enhanced Loss-at-Risk function:

$$L_{\text{CVaR-MSE}}(y, y_{\text{true}}) = L_{\text{MSE}}(y, y_{\text{true}}) + \lambda \cdot \text{CVaR}_\alpha(L_{\text{MSE}}), \quad (8)$$

where λ finely tunes the balance between MSE and CVaR, facilitating a targeted approach to manage extreme risk exposures effectively.

This adaptation of the loss function allows LLMs to remain robust in their predictive capabilities while enhancing their precision in assessing and mitigating extreme risks. CVaR's focus on the average of the most severe losses provides a more realistic and useful risk assessment, particularly suitable for financial sectors and other high-stake applications where overlooking such risks could lead to significant repercussions.

C. Mathematical Rigor and Suitability

a) Continuity and Differentiability: In Loss-at-Risk, the MSE part is a quadratic function, and all polynomial functions are continuous and differentiable across their entire domain. This property is crucial for the application of gradient-based optimization methods. The derivative of MSE is given by:

$$\frac{d}{dy} L_{\text{MSE}}(y, y_{\text{true}}) = 2(y - y_{\text{true}}), \quad (9)$$

demonstrating smoothness and suitability for such methods.

The continuity of the VaR component depends on the continuity of the loss distribution. If the loss distribution is continuous, the Cumulative Distribution Function (CDF) is also continuous, and hence VaR is continuous.

Differentiability of VaR can be complex and is contingent on the continuity and non-zero property of the Probability Density Function (PDF) at the quantile. Assuming a smooth distribution, VaR can be approximated as differentiable.

For CVaR, the continuity and differentiability are similarly conditioned on the underlying distribution of losses. Since CVaR is essentially an average of VaR over a range of confidence levels above α , its continuity depends directly on the continuity of VaR across these levels. This typically means that if VaR is continuous for a given distribution of losses, CVaR will also be continuous.

The differentiability of CVaR, however, can be more nuanced due to its dependence on the integral of quantiles which themselves may have points of non-differentiability depending on the loss distribution's characteristics. Nevertheless, assuming that the loss distribution is smooth and well-behaved, CVaR can often be treated as differentiable for practical purposes in optimizing models' risk awareness. This allows the use of gradient-based methods in training, optimizing both the average and the extreme risks associated with predictions.

b) Convexity: The convexity of the MSE is validated by its second derivative:

$$\frac{d^2}{dy^2} L_{\text{MSE}}(y, y_{\text{true}}) = 2, \quad (10)$$

which is positive, confirming that MSE is convex.

The convexity of VaR depends on the characteristics of the loss distribution. Under regular conditions (e.g., unimodal, not heavily skewed distributions), VaR is typically convex. Thus in Loss-at-Risk, the function's continuity and differentiability are typically preserved, with convexity contingent on the convexity of the VaR component.

The convexity of CVaR, much like VaR, is contingent upon the characteristics of the underlying loss distribution. CVaR, defined as the expected value of the losses exceeding the VaR

threshold, also tends to be convex under conditions where the loss distribution is smooth and regular. This integration over a range of VaR values tends to smooth out irregularities, potentially enhancing the convexity of CVaR compared to VaR alone. Therefore, when the loss distribution is unimodal and not heavily skewed, the CVaR component generally exhibits convexity. Consequently, when combined with MSE in the Loss-at-Risk function, the overall function maintains desirable convex properties, assuming the parameter λ is chosen appropriately. In this way, we confirm that the Loss-at-Risk function is mathematically robust and suitable for use in gradient-based optimization scenarios.

c) Robustness and Sensitivity: (1) Impact on accuracy. The Mean Squared Error, $L_{\text{MSE}}(y, y_{\text{true}})$, ensures that the model is penalized for any deviation from the true values, which inherently promotes high accuracy in predictions. The key property here is that MSE treats all errors equally, squaring the magnitude of deviations, which emphasizes larger errors more heavily, thus preventing large deviations in model predictions. (2) Enhanced sensitivity to extreme risks. The addition of the VaR or CVaR component introduces a focus on the tail of the loss distribution. Specifically, VaR focuses on the maximum potential loss that can occur with a certain probability, while CVaR provides an average of the worst losses beyond the VaR threshold. This aspect of the Loss-at-Risk function causes the model to not only minimize average prediction errors but also to be particularly cautious about the potential for significant outliers or extreme values in predictions.

D. Optimization in Transformers

Then we explore the mathematical underpinnings of the MSE and VaR in the context of Transformers. Our aim is to analyze the convexity of VaR concerning MSE and outline the optimization conditions for this model.

We start by establishing the risk measure function $V(\xi)$ to understand VaR's role in Transformers:

$$V(\xi) = E \left[\xi + \frac{1}{1-\alpha} (L(X) - \xi)_+ \right], \quad (11)$$

where $(z)_+$ represents the non-negative part, expressed as $\max(z, 0)$.

a) Derivative of VaR: The derivative $V'(\xi)$ is crucial for understanding VaR's behavior in the optimization process:

$$V'(\xi) = \frac{d}{d\xi} E \left[\xi + \frac{1}{1-\alpha} (L(X) - \xi)_+ \right], \quad (12)$$

The function $(L(X) - \xi)_+$ is piecewise linear, with its derivative breaking down:

- When $L(X) > \xi$, $(L(X) - \xi)_+ = L(X) - \xi$ and its derivative with respect to ξ is -1.
- When $L(X) \leq \xi$, $(L(X) - \xi)_+ = 0$, derivative is 0.

Thus, we derive:

$$V'(\xi) = 1 - \frac{1}{1-\alpha} P(L(X) > \xi). \quad (13)$$

b) *Optimization Condition:* The condition for optimization $V'(\xi) = 0$ is essential for setting the correct threshold in VaR:

$$1 - \frac{1}{1 - \alpha} P(L(X) > \xi) = 0. \quad (14)$$

Reformulating gives:

$$P(L(X) > \xi) = 1 - \alpha. \quad (15)$$

This indicates that we need to find a specific ξ , where the probability of the model's loss exceeding ξ matches $1 - \alpha$. This threshold is crucial in LLM optimization as it delineates the boundary beyond which the model's loss is unlikely to exceed under the chosen confidence level α .

This rigorous approach ensures that the Loss-at-Risk function, when calibrated with the appropriate ξ value, effectively balances between minimizing typical prediction errors and preventing large unexpected losses, hence optimizing the LLM's performance in risk-sensitive applications.

Algorithm 1 Loss-at-Risk

```

1: Input: Predictions  $y_{\text{pred}}$ , True values  $y_{\text{true}}$ , Risk level  $\alpha$ , Weight  $\lambda$ 
2: Output: Loss-at-Risk  $L_{\text{VaR-MSE}}(y, y_{\text{true}})$ 
3: Initialize MSE loss  $L_{\text{MSE}} \leftarrow \emptyset$ 
4: for each sample  $i_{\text{pred}}, i_{\text{true}}$  in  $y_{\text{pred}}$  and  $y_{\text{true}}$  do
5:   Compute MSE( $i_{\text{pred}}, i_{\text{true}}$ ) according to Eq.4
6:    $L_{\text{MSE}} \leftarrow \text{MSE}(i_{\text{pred}}, i_{\text{true}})$ 
7: end for
8: Calculate the  $\alpha$ -quantile of  $L_{\text{MSE}}$ :  

 $L_{\text{VaR}} \leftarrow \text{quantile}(L_{\text{np}}, \alpha)$ 
9: Compute the mean of  $L_{\text{MSE}}$ :  $\bar{L}_{\text{MSE}} \leftarrow \text{mean}(L_{\text{MSE}})$ 
10: Compute the total loss with VaR weighting:  

 $L_{\text{VaR-MSE}}(y, y_{\text{true}}) \leftarrow \bar{L}_{\text{MSE}} + \lambda \cdot L_{\text{VaR}}$ 
11: return  $L_{\text{VaR-MSE}}(y, y_{\text{true}})$ 

```

IV. EXPERIMENTS

In the experimental section of our study, we conduct a series of tests to evaluate the performance of the Loss-at-Risk function in various scenarios. We aim to answer three key questions through the following experiments:

- 1) How does the Loss-at-Risk function compare in baseline performance to traditional MSE loss on standard tasks?
- 2) How effective is the Loss-at-Risk function in handling extreme values?
- 3) What is the impact of key parameters such as the confidence level α and weight λ in the Loss-at-Risk function on the performance of the models?

In answering these questions, we design the following experiments: 1) Baseline Performance Comparison. We use the standard financial dataset FNSPID [52] from Nasdaq, train models using both MSE loss and the Loss-at-Risk function and compare their accuracy. 2) Performance Under Extreme Conditions. We perform the models on financial transaction data during extreme market conditions, fine-tune the models using

both MSE and Loss-at-Risk, and compare their performance on accuracy and sensitivity to high-risk scenarios. 3) Parameter Sensitivity Analysis. We vary α and λ while keeping other conditions fixed during fine-tuning, assess their performance under different parameter settings, and illustrate the impact of parameter changes on model performance through graphs.

A. Experimental Settings

a) *Datasets:* For the first experiment, we use the FN-SPID dataset in our experiments, which combines stock prices and financial news across a broad range of S&P 500 companies. This dataset features over 29.7 million stock price records and 15.7 million news entries from 1999 to 2023. Specifically, we focus on the stocks with tickers KO, AMD, TSM and WMT. These stocks cover various sectors, from beverages to big tech, providing a diverse view of market reactions to news. We choose this dataset to explore how news impacts stock behavior and to leveraging its rich time-series data to analyze the direct effects of financial news on stock movements.

TABLE I: Analysis of AMD's overall and extreme performance (top and bottom 5% data). The best results are highlighted in **bold**, and the runner-up results are highlighted in underline.

Model	MSE	MAE	R2	Max AE	Min AE
L_{MSE}	20.1314	3.2918	0.9790	8.5783	<u>5.2252</u>
$L_{\text{VaR-MSE}}$	<u>19.0661</u>	<u>3.2640</u>	<u>0.9801</u>	7.9585	5.7073
$L_{\text{CVaR-MSE}}$	18.8994	3.1816	0.9803	<u>8.0151</u>	4.6462

b) *Baseline Performance Comparison:* We conduct an analysis on the AMD stock, specifically focusing on the highest and lowest 5% of data points, corresponding to an α value of 5%, to simulate the performance under extreme conditions. The metrics we evaluated include MSE-VaR loss and MSE-CVaR loss. The outcomes of this analysis are presented in Table I to show how these models perform under stress scenarios.

The Max Absolute Error (Max AE) and Min Absolute Error (Min AE) are metrics designed to capture the absolute discrepancies between predicted values $y_{\text{pred}}(t)$ and true values $y_{\text{true}}(t)$ at specific points in time. Specifically, Max AE is the absolute error computed at the time t when y_{true} reaches its maximum value, while Min AE is calculated at the time t when y_{true} reaches its minimum.

$$t_{\text{max}} = \arg \max_t y_{\text{true}}(t), \quad (16)$$

$$\text{Max AE} = |y_{\text{true}}(t_{\text{max}}) - y_{\text{pred}}(t_{\text{max}})|.$$

$$t_{\text{min}} = \arg \min_t y_{\text{true}}(t), \quad (17)$$

$$\text{Min AE} = |y_{\text{true}}(t_{\text{min}}) - y_{\text{pred}}(t_{\text{min}})|.$$

B. Standard Prediction

In this experiment, our goal is to compare the performance of Transformers using the Loss-at-Risk function with those using traditional MSE loss regarding prediction accuracy. This

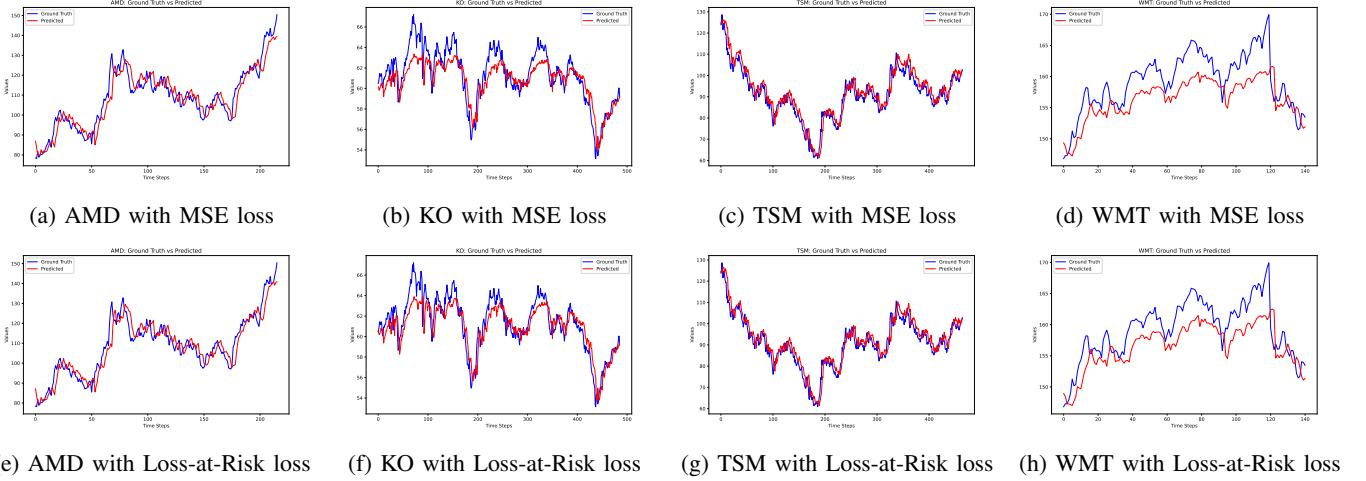


Fig. 3: Comparison of MSE and Loss-at-Risk performance across selected stocks on Transformer.

is to evaluate whether the new loss function impacts the model’s overall performance. Here we show that introducing sensitivity to extreme risks does not compromise the model’s performance on routine tasks. As shown in Fig. 3, Transformers with Loss-at-Risk exhibit similar or slightly better performance than those using MSE loss. These results indicate that the Loss-at-Risk function enhances the model’s ability to perceive extreme risks without degrading its overall predictive capabilities. These findings support applying this new loss function in risk-sensitive financial decision-making models.

C. Performance on Extreme Values

Then we test and evaluate the performance of Transformer with Loss-at-Risk under extreme value conditions and in predicting potential risk points. Both Loss-at-Risk and traditional MSE loss functions are evaluated on same Transformer and we observe their predictive behaviors during high-volatility scenarios.

As shown in Fig. 4 and Fig. 5, we observe a notable improvement in model performance with Loss-at-Risk than MSE. Specifically, the results demonstrate that the maximum absolute error (Max AE) and minimum absolute error (Min AE) are lower for models utilizing the Loss-at-Risk function, particularly with the CVaR-MSE configuration. The Max AE for the VaR-MSE model showed a significant decrease compared to the standard MSE model, indicating a better handling of worst-case scenarios. Similarly, the CVaR-MSE model demonstrated the lowest Min AE, which suggests superior performance in managing the most favorable loss outcomes under extreme conditions.

D. Ablation Study

In the ablation study, we examine the impact of varying α values on the performance of Loss-at-Risk functions that incorporate VaR and CVaR. The α , confidence level for quantifying risk, is crucial in determining how the model perceives and manages potential extreme losses. We systematically adjust α

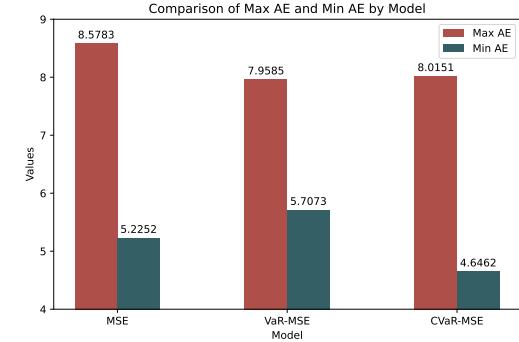


Fig. 4: Model performance under extreme values (top and bottom 5%). Lower values indicate better outcomes.

values from 0.5 to 1.0 to observe changes in MAE for both VaR and CVaR settings. Then we further explore the model’s sensitivity to different proportions of data considered under extreme risk conditions.

As shown in Fig. 6a and Fig. 6b, the results demonstrate a clear trend where increasing α values lead to a decrease in MAE, with an enhanced prediction accuracy with higher confidence levels. Specifically, the CVaR measure shows a more pronounced reduction in MAE than VaR, indicating its superior capability in managing tail risks.

For VaR, the MAE is higher at lower α levels, which implies that a lower proportion of data considered at risk leads to reduced prediction accuracy. Conversely, the accuracy improves as the α value increases, indicating a more conservative risk assessment. This trend underscores the importance of carefully selecting the α parameter to achieve an optimal balance between risk sensitivity and prediction accuracy in financial modeling, particularly in environments where accurate risk assessment is crucial. As CVaR focuses on averaging losses beyond the VaR threshold, it is more effective in scenarios demanding stringent risk management. Thus, it is a valuable tool for enhancing model reliability under extreme conditions.

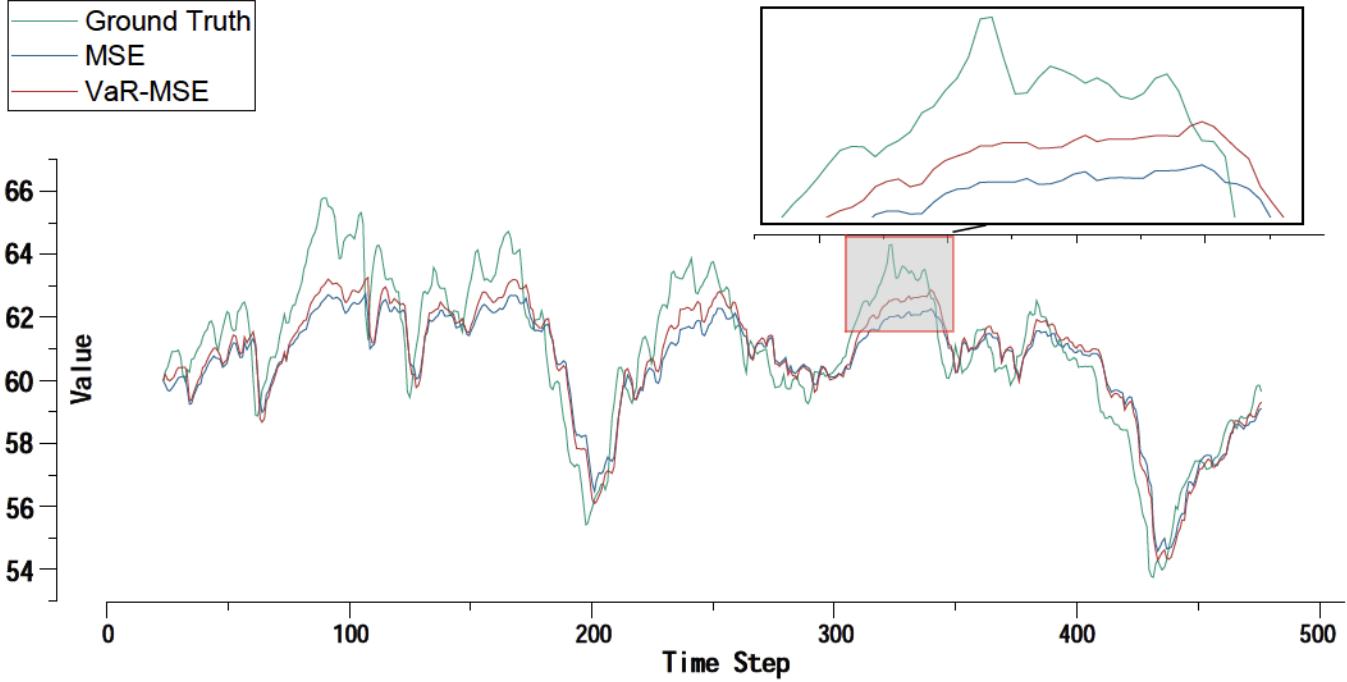


Fig. 5: Model performance under extreme values (top and bottom 5%). Lower values indicate better outcomes.

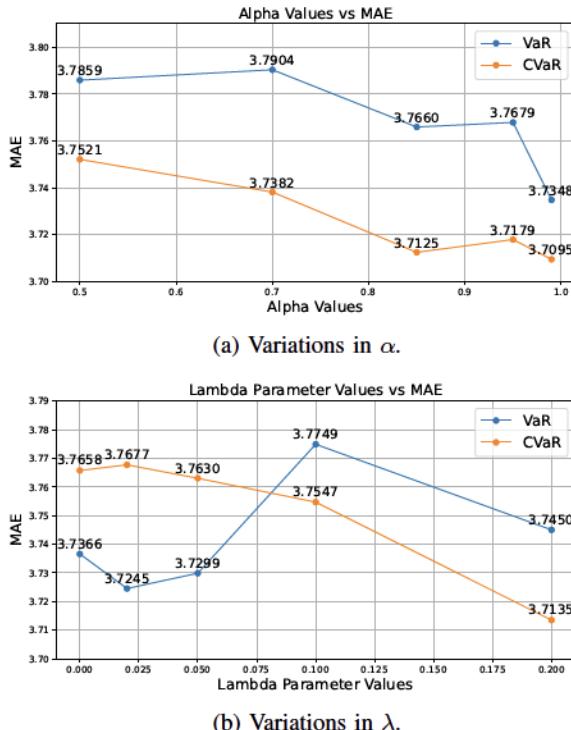


Fig. 6: $L_{\text{VaR-MSE}}$ and $L_{\text{CVaR-MSE}}$ performance on the AMD dataset for as λ and α variations.

V. CONCLUSION AND FUTURE WORK

In conclusion, our paper proposes a new loss function, Loss-at-Risk, integrating risk assessment into Transformer models' fine-tuning process for financial forecasting. In Loss-at-Risk,

we incorporate both VaR and CVaR with vanilla MSE loss to improve the models' capabilities to accurately assess and manage extreme financial risks. Our experiments with highly volatile financial datasets clearly demonstrate the enhanced ability of these Transformer models to predict and manage extreme losses. In this way, we enable safer and more reliable financial tools for decision-making processes.

For future work, further research could explore the impact of the Loss-at-Risk on large language models in various real-world scenarios to deepen the understanding of its practical applications and effectiveness in risk-sensitive environments. This could lead to a broader adoption of risk-aware methodologies across different sectors, and enhance the large language models' reliability and utility in complex decision-making scenarios.

ACKNOWLEDGMENT

The authors Jinghan Zhang, Xinhao Zhang and Kunpeng Liu are supported by NSF 2348485 and NSF 2426339.

REFERENCES

- [1] S. Wu, O. Irsøy, S. Lu, V. Dabrowski, M. Dredze, S. Gehrmann, P. Kambadur, D. Rosenberg, and G. Mann, "Bloombergpt: A large language model for finance," *arXiv preprint arXiv:2303.17564*, 2023.
- [2] Y. Li, S. Wang, H. Ding, and H. Chen, "Large language models in finance: A survey," in *Proceedings of the fourth ACM international conference on AI in finance*, pp. 374–382, 2023.
- [3] H. Yang, X.-Y. Liu, and C. D. Wang, "Fingpt: Open-source financial large language models," *arXiv preprint arXiv:2306.06031*, 2023.
- [4] E. Di Lorenzo, G. Piscopo, and M. Sibillo, "Addressing the economic and demographic complexity via a neural network approach: risk measures for reverse mortgages," *Computational Management Science*, vol. 21, no. 1, p. 11, 2024.

[5] J.-W. Chang, N. Yen, and J. C. Hung, "Design of a nlp-empowered finance fraud awareness model: the anti-fraud chatbot for fraud detection and fraud classification as an instance," *Journal of Ambient Intelligence and Humanized Computing*, vol. 13, no. 10, pp. 4663–4679, 2022.

[6] D. Chicco, M. J. Warrens, and G. Jurman, "The coefficient of determination r-squared is more informative than smape, mae, mape, mse and rmse in regression analysis evaluation," *Peerj computer science*, vol. 7, p. e623, 2021.

[7] A. Jadon, A. Patil, and S. Jadon, "A comprehensive survey of regression based loss functions for time series forecasting," *arXiv preprint arXiv:2211.02989*, 2022.

[8] Y. Xing, C. Yan, and C. C. Xie, "Predicting nvidia's next-day stock price: A comparative analysis of lstm, mlp, arima, and arima-garch models," *arXiv preprint arXiv:2405.08284*, 2024.

[9] C. Yan, Y. Weng, J. Wang, Y. Zhao, Y. Zou, Z. Li, and U. Baltimore, "Enhancing credit card fraud detection through adaptive model optimization,"

[10] Q. Wang, Y. Ma, K. Zhao, and Y. Tian, "A comprehensive survey of loss functions in machine learning," *Annals of Data Science*, pp. 1–26, 2020.

[11] Y. Tian, D. Su, S. Lauria, and X. Liu, "Recent advances on loss functions in deep learning for computer vision," *Neurocomputing*, vol. 497, pp. 129–158, 2022.

[12] F. Mo, K. Mao, Z. Zhao, H. Qian, H. Chen, Y. Cheng, X. Li, Y. Zhu, Z. Dou, and J.-Y. Nie, "A survey of conversational search," 2024.

[13] S. Fu, X. Wang, J. Tang, S. Lan, and Y. Tian, "Generalized robust loss functions for machine learning," *Neural Networks*, vol. 171, pp. 200–214, 2024.

[14] F. Mo, K. Mao, Y. Zhu, Y. Wu, K. Huang, and J.-Y. Nie, "Convqqr: Generative query reformulation for conversational search," 2024.

[15] X. Zhang, J. Zhang, B. Rekabdar, Y. Zhou, P. Wang, and K. Liu, "Dynamic and adaptive feature generation with llm," 2024.

[16] Y. Kang, Y. Song, and S. Huang, "Tie memories to e-souvenirs: Personalized souvenirs with augmented reality for interactive learning in the museum," *Preprints*, October 2024.

[17] X. Zhang, Z. Wang, L. Jiang, W. Gao, P. Wang, and K. Liu, "Tfwt: Tabular feature weighting with transformer," 2024.

[18] Z. Zhang and M. Sabuncu, "Generalized cross entropy loss for training deep neural networks with noisy labels," *Advances in neural information processing systems*, vol. 31, 2018.

[19] J. Zhang, X. Wang, W. Ren, L. Jiang, D. Wang, and K. Liu, "Ratt: A thought structure for coherent and correct llm reasoning," 2024.

[20] J. Kaplan, S. McCandlish, T. Henighan, T. B. Brown, B. Chess, R. Child, S. Gray, A. Radford, J. Wu, and D. Amodei, "Scaling laws for neural language models," *arXiv preprint arXiv:2001.08361*, 2020.

[21] W. Cai, J. Chen, and M. Li, "Exploring the encoding layer and loss function in end-to-end speaker and language recognition system," *arXiv preprint arXiv:1804.05160*, 2018.

[22] J. Howard and S. Ruder, "Universal language model fine-tuning for text classification," *arXiv preprint arXiv:1801.06146*, 2018.

[23] S. Benninga and Z. Wiener, "Value-at-risk (var)," *matrix*, vol. 32, p. s33, 1998.

[24] M. Better, F. Glover, G. Kochenberger, and H. Wang, "Simulation optimization: applications in risk management," *International Journal of Information Technology & Decision Making*, vol. 7, no. 04, pp. 571–587, 2008.

[25] R. Ismal, "Volatility of the returns and expected losses of islamic bank financing," *International Journal of Islamic and Middle Eastern Finance and Management*, vol. 3, no. 3, pp. 267–279, 2010.

[26] A. Tamar, Y. Glassner, and S. Mannor, "Optimizing the cvar via sampling," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 29, 2015.

[27] Y. Chow, A. Tamar, S. Mannor, and M. Pavone, "Risk-sensitive and robust decision-making: a cvar optimization approach," *Advances in neural information processing systems*, vol. 28, 2015.

[28] G. J. Alexander and A. M. Baptista, "A comparison of var and cvar constraints on portfolio selection with the mean-variance model," *Management science*, vol. 50, no. 9, pp. 1261–1273, 2004.

[29] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.

[30] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, et al., "Language models are few-shot learners," *Advances in neural information processing systems*, vol. 33, pp. 1877–1901, 2020.

[31] J. Zhang, X. Wang, Y. Jin, C. Chen, X. Zhang, and K. Liu, "Prototypical reward network for data-efficient rlhf," 2024.

[32] C. Yan, Y. Xing, S. Liu, E. Gao, and J. Wang, "Machine learning models for cardiovascular disease prediction: A comparative study," *bioRxiv*, pp. 2024–05, 2024.

[33] C. Yan, "Predict lightning location and movement with atmospherical electrical field instrument," in *2019 IEEE 10th Annual Information Technology, Electronics and Mobile Communication Conference (IEMCON)*, pp. 0535–0537, IEEE, 2019.

[34] X. Zhang, J. Zhang, F. Mo, Y. Chen, and K. Liu, "Tifg: Text-informed feature generation with large language models," *arXiv preprint arXiv:2406.11177*, 2024.

[35] L. Chen and G. Guo, "Nearest neighbor classification of categorical data by attributes weighting," *Expert Systems with Applications*, vol. 42, no. 6, pp. 3142–3149, 2015.

[36] Y. Chen and Y. Hao, "A feature weighted support vector machine and k-nearest neighbor algorithm for stock market indices prediction," *Expert Systems with Applications*, vol. 80, pp. 340–355, 2017.

[37] H.-C. Dan, B. Lu, and M. Li, "Evaluation of asphalt pavement texture using multiview stereo reconstruction based on deep learning," *Construction and Building Materials*, vol. 412, p. 134837, 2024.

[38] S. Chowdhury, N. Helian, and R. Cordeiro de Amorim, "Feature weighting in dbSCAN using reverse nearest neighbours," *Pattern Recognition*, vol. 137, p. 109314, 2023.

[39] X. Wang, Y. Wang, and L. Wang, "Improving fuzzy c-means clustering based on feature-weight learning," *Pattern recognition letters*, vol. 25, no. 10, pp. 1123–1132, 2004.

[40] D. S. Yeung and X. Wang, "Improving performance of similarity-based clustering by feature weight learning," *IEEE transactions on pattern analysis and machine intelligence*, vol. 24, no. 4, pp. 556–561, 2002.

[41] M. Wang and S. Liu, "Machine learning-based research on the adaptability of adolescents to online education," *arXiv preprint arXiv:2408.16849*, 2024.

[42] L. Feng, S. Shu, Z. Lin, F. Lv, L. Li, and B. An, "Can cross entropy loss be robust to label noise?," in *Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence*, pp. 2206–2212, 2021.

[43] Z. Wang and A. C. Bovik, "Mean squared error: Love it or leave it? a new look at signal fidelity measures," *IEEE signal processing magazine*, vol. 26, no. 1, pp. 98–117, 2009.

[44] T. Kim, J. Oh, N. Kim, S. Cho, and S.-Y. Yun, "Comparing kullback-leibler divergence and mean squared error loss in knowledge distillation," *arXiv preprint arXiv:2105.08919*, 2021.

[45] Y. Song, P. Arora, R. Singh, S. T. Varadharajan, M. Haynes, and T. Starner, "Going blank comfortably: Positioning monocular head-worn displays when they are inactive," in *Proceedings of the 2023 ACM International Symposium on Wearable Computers*, pp. 114–118, 2023.

[46] Z. Wang, H. Yan, C. Wei, J. Wang, S. Bo, and M. Xiao, "Research on autonomous driving decision-making strategies based deep reinforcement learning," *arXiv preprint arXiv:2408.03084*, 2024.

[47] A. Mustafa, A. Mikhailiuk, D. A. Iliescu, V. Babbar, and R. K. Mantlik, "Training a task-specific image reconstruction loss," in *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pp. 2319–2328, 2022.

[48] Y. Kang, Z. Zhang, M. Zhao, X. Yang, and X. Yang, "Tie memories to e-souvenirs: Hybrid tangible ar souvenirs in the museum," in *Adjunct Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*, pp. 1–3, 2022.

[49] L. Yang, H. Li, F. Meng, Q. Wu, and K. N. Ngan, "Task-specific loss for robust instance segmentation with noisy class labels," *IEEE Transactions on Circuits and Systems for Video Technology*, 2021.

[50] Z. Wang, S. Chang, J. Zhou, M. Wang, and T. S. Huang, "Learning a task-specific deep architecture for clustering," in *Proceedings of the 2016 SIAM International Conference on Data Mining*, pp. 369–377, SIAM, 2016.

[51] H. Gao, H. Wang, Z. Feng, M. Fu, C. Ma, H. Pan, B. Xu, and N. Li, "A novel texture extraction method for the sedimentary structures' classification of petroleum imaging logging," in *Pattern Recognition: 7th Chinese Conference, CCPR 2016, Chengdu, China, November 5–7, 2016, Proceedings, Part II 7*, pp. 161–172, Springer, 2016.

[52] Z. Dong, X. Fan, and Z. Peng, "Fnspid: A comprehensive financial news dataset in time series," *arXiv preprint arXiv:2402.06698*, 2024.