

Explainable Neuro-Symbolic Intelligence for Constraint-Aware Financial Decision Making Under Uncertainty

Gopichand Talluri

Overland Park, Kansas, USA

Abstract — Under uncertain conditions, financial decision-making needs methods that combine predictive accuracy, risk management, regulatory requirements, and explainability. Existing financial models, which include both statistical forecasting and conventional machine learning techniques, may fail in terms of describing complex nonlinear market dynamics, quantifying uncertainty in financial decisions, and explaining the effects of financial constraints on decisions. This research presents an ENSC-FI approach to generating financial decisions using a framework consisting of neural learning, symbolic financial rules, uncertainty quantification, and constraint-aware optimization. ENSC-FI seeks to create clear financial decisions using market indicators, risk constraints, budget limitations, and explanation of decisions made. Performance of the proposed framework may be compared to existing traditional machine learning and optimization techniques in terms of return, Sharpe ratio, maximum drawdown, constraint satisfaction rate, and explanation quality. As for an evaluation target of ENSC-FI, the improvement of constraint satisfaction up to 95% and decrease in decision errors by 15% compared to one of the selected baselines will be considered. Those are the targets to be achieved but not actual experimental results yet.

Keywords— Neuro-symbolic artificial intelligence, explainable AI, financial decision-making, uncertainty quantification, constraint-aware optimization, risk management, financial intelligence.

I. INTRODUCTION

Financial decision-making has become increasingly complex due to market volatility, uncertainty, changing economic conditions, and the growing volume of financial data. Decisions involving investment allocation, portfolio management, lending, and risk assessment require models that can analyse financial patterns while satisfying predefined constraints. Artificial intelligence (AI) and machine learning techniques have demonstrated their ability to identify complex relationships in financial data and support data-driven predictions. However, conventional machine learning models often operate as black boxes, making it difficult to understand the reasoning behind their decisions. Furthermore, predictive accuracy alone does not guarantee that financial decisions satisfy budget limits, risk thresholds, regulatory requirements, and other operational constraints.

Explainable artificial intelligence (XAI) has emerged as an important approach for improving transparency in automated decision-making. It provides insights into the factors influencing model predictions and helps users understand the basis of financial recommendations. Nevertheless, many explanation techniques are applied after a prediction has been generated and may not ensure that the final decision complies with predefined financial rules. Similarly, symbolic reasoning systems can enforce explicit rules but may lack the flexibility

required to learn complex patterns from large and dynamic financial datasets

Neuro-symbolic AI provides a potential solution due to the integration of neural learning techniques with symbolic knowledge and logic reasoning. In such an approach, neural networks are able to discover non-linear relationships between features, while symbolic aspects can describe financial rules, risk constraints, and other decision constraints. Adding uncertainty estimation and constraint-aware optimization to the neuro-symbolic approach will help make reliable decisions in changing environments.

In order to solve these problems, this paper introduces the concept of Explainable Neuro-Symbolic Constraint-Aware Financial Intelligence (ENSC-FI). It is supposed that the framework will combine data-driven financial analysis, symbolic representation of rules, uncertainty estimation, decision optimization under constraints, and explanation techniques for the generated decisions. The goal of this work is to make decisions that are data-driven, transparent, risk-aware, and comply with certain financial constraints. This framework needs to be tested on a set of financial experiments.

1. Research Objectives

The primary research objectives for the proposed ENSC-FI framework include:

- Building a neuro-symbolic financial decision-making framework that combines neural learning and symbolic reasoning for the analysis of financial data.
- Making explainable decisions in finance through identification of the factors, rules, and reasoning that determine the recommendations provided by the model.
- Implementing constraint-based decision optimization taking into account certain preconditions related to financial decision-making like budget limitations and risk levels.
- Managing uncertainty in financial decision making through uncertainty assessment and testing the robustness of decisions made.
- Experiments with the proposed framework compared to selected baseline models using appropriate performance metrics such as decision accuracy, risk-adjusted return, maximum drawdown, constraint satisfaction rate, and explanation.

II. LITERATURE SURVEY

A provenance-aware access control framework was introduced by L. Sun et al. (2016) to include typed provenance information in the access control decisions. The data provenance is data regarding the origin and history and transformations of data. The framework also took into account the type of provenance to gain a better understanding of access control, and attempted to consider both the identity of the user and the requested resource, as well as the context and history of the data. This approach was applicable to secure information sharing, accountability and auditing in distributed systems. The research presented a basis to incorporate data lineage and provenance information within security mechanisms to make informed and traceable access control decisions in organizations.

Al-Rakhami, M. S. and Al-Mashari, M. (2022) introduced a framework called ProChain that is a new approach to provenance-aware traceability for an Internet of Things (IoT)-based supply chain system. Supply chains have several stakeholders, devices, and information system involved, that continuously create and share data. It is important to preserve these data and their history for transparent, authentic and accountable research. ProChain dealt with the problem of tracking data and activities across the supply chain by introducing "provenance" information to the traceability chain. A framework like this can facilitate the stakeholders to ensure product histories, determine the origins of inconsistencies, and strengthen the consistency and reliability of product information in the supply chain. The study underlined the need for provenance management in IoT where data comes from

sources that are heterogeneous and is subjected to several processing steps.

In this work, the construction and application of knowledge graphs for fault handling of power grid have been studied by N. Xiao et al. (2023). The management and control of power grid faults involve the analysis of complex information, such as the characteristics of the equipment, fault records, operational procedures and the interaction relations of the components. A knowledge graph is a graph that represents such information by entities and relationships, and is suitable for organizing, connecting and recalling domain knowledge. The suggested solution was to organize power grid fault-related information to facilitate fault analysis and handling process. A knowledge graph can be used to model relationships between equipment, faults, and relevant operational knowledge, which can help employees determine potential causes and make informed decisions about actions to take. The results of this research confirmed that knowledge graph technology can be applied to specialized industrial fields and served as a reference on how to build a structured and interconnected knowledge representation.

C. Yan, et al. (2023) investigated a solution to combine multi-source heterogeneous data in building an enterprise knowledge graph. Organizations generally have information spread out in a variety of databases, documents, business applications and other systems, which means that data formats, structures, meanings and levels of detail can differ. The combination of these sources involves data extraction, entity recognition, relationship identification and integration into a single representation. This study overcame these challenges by proposing an enterprise knowledge graph construction method that aims to link scattered organizational knowledge. This kind of representation can lead to better information retrieval, knowledge sharing, and finding relations that could not be discovered in isolated information sources. This work is especially significant in the context of enterprise information management as it provides an example of the need for data integration and semantic relationships when building a coherent knowledge base for an organization.

In this work, a new method for automatically summarizing the survey class, namely KeyphraseDS, was presented by S. Yang et al., (2017). When creating a survey or literature review, one must be able to identify important concepts, organise the relevant information and present the relationships between the topics of the various documents. Keyphrases identify the main ideas of a document and give useful clues for selecting and organizing information that is relevant to the document. The

study examined the possibility of using keyphrase information to aid in the automatic production of surveys to simplify some of the manual effort required to review and synthesize large document collections. It can be applied in automated text analysis, document summarization and knowledge discovery. In a larger information extraction system, this keyphrase identification can also serve to structure the information for the specific domain and to present meaningfully important representations of the important concepts.

Han et al. (2022) aimed at formulating a relation-oriented model which is based on the concept of embedding global contextual information to achieve joint entity and overlapping relation extraction from text. Named entity recognition (NER) is the task of recognizing important entities like persons, organizations, places and technical terms, and relation extraction (RE) is the task of deriving semantic relationships between these entities. These tasks are more difficult if there are multiple relationships in a sentence, and/or multiple relationships involving the same entity. The independent representation of such an overlapping structure may be difficult using the traditional extraction methods. To tackle this challenge, the proposed model focused on the relationships embedded in the text and added additional context in the sentences to enhance the joint extraction process. The method is beneficial for building knowledge graph, if the entity and relation are accurately extracted, the structured information in the text can be accurately represented as relation graph.

P.-T. Lai et al. (2023) proposed a new model named BioREx, which aims to enhance biomedical relation extraction using heterogeneous datasets. P.-T. Lai et al. (2023) presented a new approach called BioREx to leverage heterogeneous datasets to improve biomedical relation extraction. This study introduced BioREx which aims to enhance biomedical relation extraction based on heterogeneous datasets. There are complex relationships between genes, proteins, diseases, chemicals, drugs, etc., and other biological entities in biomedical literature. The automatic extraction of these relationships can help in biomedical knowledge discovery, literature mining and creation of biomedical knowledge graphs. Biomedical data sets can still vary in their data distributions, terms, relations and annotation schemes. Such variations may restrict the utility of the models trained on a single data set. In this regard, BioREx tried to tackle this challenge by leveraging information obtained from multiple data types for biomedical relation extraction. The study is also applicable to specialized information extraction, as it demonstrates the benefits of integrating complementary datasets to enable systems to learn

more complete relationships and the extraction of structured knowledge from specialized scientific texts.

In this study, a neural biomedical named entity recognition and normalization (NERN) tool, namely BERN2 was introduced by M. Sung et al., (2022). There are many specialised words related to diseases, genes, chemicals, species and other biological concepts used in biomedical texts. These terms can be recognized in textual data using named entity recognition, and standardized identifiers for these terms can be mapped from textual data to biomedical databases or knowledge resources using entity normalization. This normalisation is significant since the same biomedical concept can have several abbreviations, synonyms or different names. BERN2 focuses on the development of tools to aid with more complex biomedical entity recognition and normalization, with the aim of assisting biomedical researchers in processing biomedical literature and linking extracted terms to structured biomedical knowledge resources. It has relevance in the context of knowledge graph construction because it can be used to recognize and normalize biomedical entities, which is vital to integrating information from various documents and preventing inconsistencies resulting from the different terms used to describe the same concept.

X. Wang et al. (2023) explored the problem of extracting the domain-specific Chinese named entities from aviation safety reports. The aviation safety documents are filled with aviation jargon associated with aircraft parts, operating methods, incidents, personnel and safety conditions. These terms might not be recognized well by general-purpose NER systems since they focus on a wider range of language tasks than aviation specific ones. The study aimed to solve the problems of entity recognition in aviation safety reports from China, and proved the necessity of customizing the information extraction technique to the specific domain of a certain safety report. Correct entity extraction can be used for aviation safety knowledge organization, aviation safety incident classification and the identification of recurring risk factors. The work is of interest for domain-specific knowledge graph development because the accurate identification of specialized entities is the basis to extract the relationships and set up structured representations of technical information.

III. PROPOSED MODEL

The proposed Explainable Neuro-Symbolic Constraint-Aware Financial Intelligence (ENSC-FI) framework integrates neural learning, symbolic reasoning, uncertainty assessment, and constraint-aware optimization to support transparent financial

decision-making. The methodology consists of financial data acquisition and preprocessing, neural feature learning, symbolic knowledge representation, uncertainty estimation, constrained decision optimization, and explainability-based evaluation. Each stage contributes to generating decisions that account for financial patterns, predefined requirements, and uncertainty in market conditions.

1. Financial Data Acquisition and Preprocessing

The first stage collects historical financial data relevant to the selected decision-making task, such as portfolio allocation, investment selection, or financial risk assessment. Depending on the application, the input variables may include asset prices, historical returns, trading volume, volatility indicators, and relevant economic variables. The dataset is organized chronologically to preserve temporal relationships and avoid information leakage between training and testing sets.

Data preprocessing includes handling missing values, identifying duplicate records, checking anomalous observations, and normalizing numerical features where appropriate. Financial returns and volatility indicators can be derived from historical prices to represent market behavior more effectively. The processed dataset is divided into training, validation, and testing subsets, with the testing subset reserved for evaluating generalization to unseen observations.

Let X represent the pre-processed financial feature matrix and y represent the target variable associated with the selected financial task:

$$X = \{x_1, x_2, \dots, x_n\}, \quad y = \{y_1, y_2, \dots, y_n\}$$

For a financial asset with price P_t at time t , the simple return is calculated as follows:

$$r_t = \frac{P_t - P_{t-1}}{P_{t-1}}$$

2. Neural Learning and Financial Pattern Extraction

The neural learning component identifies nonlinear relationships and complex patterns in the preprocessed financial data. Depending on the dataset and prediction task, a multilayer perceptron, recurrent neural network, or transformer-based model can be selected. The neural model learns a mapping between the input features and the target financial outcome, which may represent expected return, risk level, or a decision-relevant prediction.

The model parameters are learned using a suitable loss function and an optimization algorithm. Training performance is monitored using validation data to reduce overfitting and support generalization. The resulting predictions are passed to the subsequent symbolic reasoning and optimization stages rather than being treated as final financial decisions.

The neural prediction function is expressed as:

$$\hat{y} = f_{\theta}(X)$$

where f_{θ} represents the neural model, θ denotes its trainable parameters, and $\hat{y}$ denotes the predicted financial outcome.

For a regression-based prediction task, the mean squared error (MSE) can be used as the training loss:

$$L_{\text{MSE}} = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2$$

3. Symbolic Knowledge Representation and Rule-Based Reasoning

The symbolic reasoning component represents financial policies, risk limits, and operational requirements as explicit logical rules. Examples include limiting the allocation to an individual asset, ensuring that the total investment does not exceed the available budget, and preventing decisions that violate a predefined risk threshold. The symbolic layer receives neural predictions and evaluates candidate decisions against the encoded rules. If a decision violates a hard constraint, the system rejects or modifies that decision before final selection. This separation between neural prediction and symbolic verification is consistent with recent neuro-symbolic financial research, which uses a neural component to interpret inputs and a symbolic verifier to enforce formal business rules. In ENSC-FI, the precise rules depend on the selected financial application and must be defined before evaluation.

Let d denote a candidate financial decision and $C_j(d)$ denote the j -th constraint. A feasible decision must satisfy all hard constraints:

$$C_j(d) \leq 0, \quad j = 1, 2, \dots, m$$

The feasibility indicator is defined as:

$$F(d) = \begin{cases} 1, & \text{if all constraints are satisfied,} \\ 0, & \text{otherwise.} \end{cases}$$

4. Uncertainty Estimation

Financial markets are affected by unpredictable price movements, economic events, and changing volatility. Therefore, a model's predictions should not be treated as perfectly certain. The uncertainty estimation stage assesses the reliability of the neural predictions and communicates this information to the decision optimizer. Depending on the selected implementation, uncertainty can be estimated using an ensemble of neural models, Monte Carlo dropout, or probabilistic prediction intervals.

For example, an ensemble generates several predictions for the same input and measures their variation. Greater variation indicates that the model's output is less stable, which can motivate a more conservative decision. The uncertainty estimate is used as an additional input to the optimization process rather than as a replacement for explicit financial risk controls.

For K predictive models, the ensemble mean and predictive variance can be represented as:

$$\bar{y} = \frac{1}{K} \sum_{k=1}^K \hat{y}^{(k)}$$

$$U(X) = \frac{1}{K} \sum_{k=1}^K (\hat{y}^{(k)} - \bar{y})^2$$

Here, $\hat{y}^{(k)}$ is the prediction from the k -th model, and $U(X)$ represents the estimated predictive disagreement.

5. Constraint-Aware Financial Decision Optimization

The optimization stage combines predicted financial outcomes, risk measures, uncertainty estimates, and symbolic constraints to select a feasible decision. For portfolio allocation, the decision variables may represent the proportion of available capital assigned to each asset. The objective can balance expected return against risk and uncertainty while respecting investment limits.

A risk-aware objective may penalize portfolio variance and predictive uncertainty. The penalty weights determine the relative importance assigned to each component and should be selected using validation data or a predefined financial policy. Hard constraints are enforced separately so that a favourable objective value cannot compensate for violating mandatory requirements.

For portfolio weights w , expected returns μ , covariance matrix Σ , uncertainty penalty $U(w)$, and penalty coefficients λ and γ , the optimization problem can be expressed as:

$$\max_w \quad \mu^T w - \lambda w^T \Sigma w - \gamma U(w)$$

subject to:

$$\sum_{i=1}^N w_i = 1, \quad w_i \geq 0$$

$$w_i \leq w_i^{\max}, \quad i = 1, 2, \dots, N$$

Here, N is the number of assets, w_i is the allocation to asset i , and $w_i^{\max}$ is its maximum permitted allocation. The covariance term penalizes portfolio risk, while $U(w)$ penalizes uncertainty associated with the selected allocation. Other financial tasks require different objective functions and constraints.

6. Explainability and Decision Justification

The explainability stage provides a human-readable justification for the selected decision. Instead of presenting only the final recommendation, ENSC-FI is designed to identify the influential input features, applicable symbolic rules, uncertainty level, and constraints that affected the optimization outcome. For a portfolio decision, an explanation might indicate that an asset received a lower allocation because of elevated predicted risk or a maximum-allocation rule.

Explanations should be evaluated rather than assumed to be correct simply because they are understandable. The evaluation can examine whether the stated factors are consistent with the model's behavior, whether explanations remain stable under small changes in input, and whether the identified constraints accurately describe the decision. Where appropriate, counterfactual analysis can show how a change in a financial input or constraint would affect the selected action.

A feature-attribution explanation for input x can be represented abstractly as:

$$E(x) = \{e_1, e_2, \dots, e_p\}$$

where e_i denotes the contribution attributed to feature i , and p is the number of input features. The specific explanation method, such as SHAP or an intrinsically interpretable reasoning trace, must be selected and documented in the implementation.

7. Performance Evaluation

The proposed framework should be evaluated against suitable baseline models, such as a conventional neural predictor, a neural model combined with a separate optimizer, and a symbolic constraint-verification baseline. All methods should use the same chronological data splits, financial assumptions, transaction-cost settings, and test periods to ensure a fair comparison. Performance should be measured on unseen data and, where feasible, under different market conditions.

For portfolio applications, the evaluation can include annualized return, Sharpe ratio, maximum drawdown, and constraint-satisfaction rate. Prediction error can also be reported if the model forecasts returns or prices. Explainability and uncertainty quality should be evaluated using suitable metrics, such as explanation fidelity and expected calibration error, when applicable.

The Sharpe ratio is calculated as:

$$S = \frac{R_p - R_f}{\sigma_p}$$

where R_p is the portfolio return, R_f is the risk-free rate over the same period, and σ_p is the standard deviation of portfolio returns, with all quantities expressed on a consistent time basis. The maximum drawdown is defined as:

$$MDD = \max_t \left(\frac{H_t - V_t}{H_t} \right)$$

where V_t is the portfolio value at time t , and H_t is the highest portfolio value observed up to that time.

The constraint-satisfaction rate is calculated as:

$$CSR = \frac{N_{\text{feasible}}}{N_{\text{total}}} \times 100\%$$

where N_{feasible} is the number of evaluated decisions satisfying all hard constraints and N_{total} is the total number of evaluated decisions.

The proposed ENSC-FI architecture consists of eight interconnected modules. Financial data are first pre-processed and supplied to a neural learning module to predict financial outcomes. A symbolic reasoning module checks financial rules and constraints, while uncertainty estimation evaluates the reliability of predictions.




Fig. 1. Proposed Model Architecture of Explainable Neuro-Symbolic Constraint-Aware Financial Intelligence (ENSC-FI).

A constraint-aware optimizer then selects a feasible financial decision. The explainability module generates human-readable justifications, and the evaluation module measures predictive performance, risk, constraint satisfaction, and explanation quality. The final output includes the recommended action, influential factors, applied rules, and associated uncertainty




Fig. 2. Architectural Data Flow Diagram of the Proposed ENSC-FI Framework.

- **Input:** Historical financial data, market indicators, and economic variables enter the system.
- **Preprocessing:** Missing values and anomalies are handled, and relevant features are prepared.
- **Prediction:** The neural module predicts financial outcomes; the uncertainty module estimates prediction variability.
- **Rule verification:** The symbolic knowledge base supplies financial and regulatory constraints to the optimizer.
- **Decision optimization:** The optimizer selects a feasible decision using predicted outcomes, risk, and uncertainty.
- **Explanation and evaluation:** The system generates a decision rationale and evaluates financial performance and constraint satisfaction.
- **Output:** The framework produces an interpretable financial recommendation, including its key factors, applicable rules, and uncertainty.

Key Implementation

The proposed Explainable Neuro-Symbolic Constraint-Aware Financial Intelligence (ENSC-FI) framework can be implemented using the following key components.

- **Financial Data Processing:** Collect historical prices, returns, market indicators, and relevant economic data. Clean missing values, handle outliers, normalize features, and generate technical indicators.
- **Neural Learning Module:** Implement an ML model, such as an LSTM or Transformer, to learn nonlinear financial patterns and predict expected returns or risk indicators.
- **Symbolic Reasoning Engine:** Encode financial rules, budget limits, risk thresholds, and asset-allocation constraints to verify whether candidate decisions are valid.
- **Uncertainty Estimation:** Use ensemble predictions, Monte Carlo dropout, or probabilistic models to estimate prediction uncertainty and support risk-aware decisions.
- **Constraint-Aware Optimization:** Combine predicted returns, portfolio risk, and uncertainty penalties to identify feasible decisions that satisfy the predefined constraints.
- **Explainability Module:** Use SHAP feature attribution, rule traces, and constraint-impact analysis to explain why a particular financial decision was selected.
- **Performance Evaluation:** Compare ENSC-FI with baseline models using prediction error, Sharpe ratio, maximum drawdown, constraint-satisfaction rate, and explanation quality.

Key Highlights

- **Hybrid Intelligence:** Combines neural learning with symbolic reasoning.

- **Explainable Decisions:** Identifies important financial factors and applicable rules.
- **Constraint Compliance:** Incorporates budget, risk, and asset-allocation limits.
- **Uncertainty Awareness:** Considers prediction reliability during decision-making.
- **Risk-Aware Optimization:** Balances expected returns against risk and uncertainty.
- **Measurable Evaluation:** Supports comparison with existing models using quantitative metrics.
- **Modular Architecture:** Allows individual components to be updated or replaced as requirements evolve.

IV. RESULTS AND DISCUSSIONS

The proposed Explainable Neuro-Symbolic Constraint-Aware Financial Intelligence (ENSC-FI) framework is evaluated through illustrative comparisons with Random Forest, XGBoost, Long Short-Term Memory (LSTM), and Multilayer Perceptron (MLP) models. The evaluation considers prediction error, risk-adjusted performance, maximum drawdown, constraint satisfaction, uncertainty calibration, explanation fidelity, and robustness under increasing market uncertainty. The comparison values are synthetic and manually specified for demonstration; therefore, they illustrate the intended evaluation methodology rather than verified experimental performance.

1. Comparative Prediction Error Analysis

Prediction error is an important metric for evaluating the reliability of financial forecasting models. Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE) are used to compare the models. In the illustrative results, ENSC-FI records an MAE of 0.013 and an RMSE of 0.021, compared with 0.021 and 0.032, respectively, for Random Forest.




Fig. 3. Comparative Prediction Error of Financial Models.

These values suggest how the proposed framework could be evaluated for predictive accuracy. However, because the values are synthetic, they do not establish that ENSC-FI outperforms the baseline models in real financial data

2. Risk-Adjusted Return Analysis

The Sharpe ratio measures portfolio excess return relative to return volatility and is commonly used to evaluate risk-adjusted performance. The illustrative values assign ENSC-FI a Sharpe ratio of 1.42, compared with 1.18 for LSTM, 1.05 for XGBoost, 0.96 for MLP, and 0.82 for Random Forest. This comparison demonstrates how risk-adjusted performance can be reported when evaluating the proposed framework. Actual financial performance would require portfolio simulation using historical observations, consistent transaction-cost assumptions, and out-of-sample testing.




Fig. 4. Risk-Adjusted Return Comparison Using the Sharpe Ratio.

3. Maximum Drawdown Analysis

Maximum drawdown measures the largest decline from a previous portfolio peak to subsequent trough. It provides an indication of downside exposure during adverse market conditions. In the illustrative comparison, ENSC-FI has a maximum drawdown of 14.6%, whereas the corresponding values are 18.2% for LSTM, 19.8% for XGBoost, 21.0% for MLP, and 22.5% for Random Forest. These values demonstrate how drawdown can be used to compare risk profiles. The apparent advantage of ENSC-FI is hypothetical and must be tested using a consistent back testing methodology before drawing conclusions about real-world risk reduction.




Fig. 5. Maximum Drawdown Comparison Across Financial Models.

4. Constraint Satisfaction Analysis

Constraint satisfaction is particularly important in ENSC-FI because the framework is designed to integrate neural predictions with symbolic financial rules and constraint-aware optimization. The illustrative results assign the proposed model a constraint-satisfaction rate of 97%, compared with 88% for XGBoost, 86% for LSTM, 84% for Random Forest, and 82% for MLP. This metric represents the percentage of evaluated decisions that satisfy all predefined hard constraints. In an actual implementation, hard constraints should be enforced and independently verified rather than relying solely on an average satisfaction score. The synthetic results illustrate the intended comparison but do not prove that the proposed framework achieves a 97% satisfaction rate.




Fig. 6. Constraint Satisfaction Rate Comparison.

5. Uncertainty Calibration Analysis

Financial decisions are affected by market volatility, incomplete information, and unpredictable events. Therefore, evaluating the reliability of predictive uncertainty is important when integrating uncertainty estimates into the decision process. The illustrative calibration-error values are 0.041 for ENSC-FI, 0.068 for LSTM, 0.075 for XGBoost, 0.083 for

MLP, and 0.092 for Random Forest. Lower calibration error is generally preferable when the metric is appropriately defined and calculated on a comparable scale. In practical experiments, the selected calibration metric and its calculation procedure must be specified, and uncertainty estimates should be evaluated on held-out data.




Fig. 7. Uncertainty Calibration Error Comparison.

6. Explainability Analysis

Explainability is a central component of ENSC-FI because financial recommendations should be accompanied by understandable reasons. The illustrative explanation-fidelity score is 0.89 for ENSC-FI, compared with 0.61 for XGBoost, 0.58 for LSTM, 0.55 for Random Forest, and 0.52 for MLP. This comparison illustrates a possible evaluation of how faithfully an explanation represents the behaviour of the corresponding model. However, explanation fidelity depends on the chosen definition, reference model, and evaluation procedure. Since the values are manually specified, they cannot be interpreted as evidence that ENSC-FI produces more faithful explanations in practice.




Fig. 8. Explanation Fidelity Comparison Across Financial Models.

7. Robustness Under Market Uncertainty

Robustness analysis examines how model performance changes as market uncertainty increases. The illustrative evaluation considers five uncertainty levels, ranging from 0.10 to 0.50. ENSC-FI is assigned normalized performance values decreasing from 0.96 to 0.85, while the baseline models show larger decreases over the same range. This scenario demonstrates how a robustness graph can help assess whether uncertainty estimation and constraint-aware optimization support stable decisions. In real experiments, uncertainty levels should be derived from defined market scenarios or controlled perturbations, and performance should be measured consistently across all models.




Fig. 9. Model Robustness Under Increasing Market Uncertainty.

V. CONCLUSION

This paper presents Explainable Neuro-Symbolic Constraint-Aware Financial Intelligence (ENSC-FI), a proposed framework for transparent and risk-aware financial decision-making under uncertainty. The framework integrates neural learning, symbolic reasoning, uncertainty estimation, and constraint-aware optimization to address the limitations of conventional data-driven financial models. By combining learned financial patterns with explicit financial rules, ENSC-FI aims to generate decisions that consider predictive outcomes, risk exposure, budget limits, and predefined operational constraints. The explainability component is designed to provide understandable justifications for the selected decisions by identifying influential financial factors and applicable constraints. The illustrative evaluation considers prediction error, Sharpe ratio, maximum drawdown, constraint-satisfaction rate, uncertainty calibration, explanation fidelity, and robustness under varying market uncertainty. Although the synthetic comparison values demonstrate how the proposed

framework can be evaluated against Random Forest, XGBoost, LSTM, and MLP baselines, they do not establish actual performance improvements. Experimental validation using historical financial datasets, out-of-sample testing, and consistent benchmarking is therefore necessary to determine the framework's effectiveness.

Future work will focus on implementing ENSC-FI, integrating real-world financial data, improving uncertainty calibration, and evaluating its performance across different market conditions. Further investigation will also examine computational efficiency, scalability, and the reliability of generated explanations. Overall, ENSC-FI provides a methodological foundation for developing financial decision-support systems that aim to combine predictive intelligence, explicit constraint enforcement, and interpretable reasoning in uncertain financial environments.

REFERENCES

1. L. Sun, J. Park, D. Nguyen, and R. Sandhu, "A provenance-aware access control framework with typed provenance," *IEEE Trans. Dependable Secure Comput.*, vol. 13, no. 4, pp. 411–423, Jul. 2016, doi: 10.1109/TDSC.2015.2410793.
2. M. S. Al-Rakhami and M. Al-Mashari, "ProChain: Provenance-aware traceability framework for IoT-based supply chain systems," *IEEE Access*, vol. 10, pp. 3631–3642, 2022, doi: 10.1109/ACCESS.2021.3135371.
3. P. N. Tinguely, J. Lee, and V. F. He, "Designing human resource management systems in the age of AI," *J. Org. Design*, vol. 12, no. 4, pp. 263–269, Sep. 2023, doi: 10.1007/s41469-023-00153-x.
4. N. Xiao, B. Peng, X. Li, J. Wu, J. Lou, and Y. Si, "Research on the construction and implementation of power grid fault handling knowledge graphs," *Energy Rep.*, vol. 9, pp. 182–189, Apr. 2023, doi: 10.1016/j.egy.2023.02.073.
5. C. Yan, X. Fang, X. Huang, C. Guo, and J. Wu, "A solution and practice for combining multi-source heterogeneous data to construct enterprise knowledge graph," *Frontiers Big Data*, vol. 6, Sep. 2023, Art. no. 1278153, doi: 10.3389/fdata.2023.1278153.
6. S. Yang, W. Lu, D. Yang, X. Li, C. Wu, and B. Wei, "KeyphraseDS: Automatic generation of survey by exploiting keyphrase information," *Neurocomputing*, vol. 224, pp. 58–70, Feb. 2017, doi: 10.1016/j.neucom.2016.10.052.
7. H. Han, J. Wang, and X. Wang, "A relation-oriented model with global context information for joint extraction of overlapping relations and entities," *Frontiers Neuroinformatics*, vol. 16, Jul. 2022, Art. no. 914705, doi: 10.3389/fnbot.2022.914705.
8. P.-T. Lai, C.-H. Wei, L. Luo, Q. Chen, and Z. Lu, "BioREx: Improving biomedical relation extraction by leveraging heterogeneous datasets," *J. Biomed. Informat.*, vol. 146, Oct. 2023, Art. no. 104487, doi: 10.1016/j.jbi.2023.104487.
9. M. Sung, M. Jeong, Y. Choi, D. Kim, J. Lee, and J. Kang, "BERN2: An advanced neural biomedical named entity recognition and normalization tool," *Bioinformatics*, vol. 38, no. 20, pp. 4837–4839, Oct. 2022, doi: 10.1093/bioinformatics/btac598.
10. X. Wang, Z. Gan, Y. Xu, B. Liu, and T. Zheng, "Extracting domain-specific Chinese named entities for aviation safety reports: A case study," *Appl. Sci.*, vol. 13, no. 19, p. 11003, Oct. 2023, doi: 10.3390/app131911003.
11. Y. Fang, H. Wang, L. Zhao, F. Yu, and C. Wang, "Dynamic knowledge graph based fake-review detection," *Appl. Intell.*, vol. 50, no. 12, pp. 4281–4295, Dec. 2020, doi: 10.1007/s10489-020-01761-w.
12. X. Suo, B. Guo, Y. Shen, W. Wang, Y. Chen, and Z. Zhang, "Embodying the number of an entity's relations for knowledge representation learning," *Int. J. Softw. Eng. Knowl. Eng.*, vol. 31, no. 10, pp. 1495–1515, Oct. 2021, doi: 10.1142/s0218194021500509.
13. Y. Dai, L. Shou, M. Gong, X. Xia, Z. Kang, Z. Xu, and D. Jiang, "Graph fusion network for text classification," *Knowl.-Based Syst.*, vol. 236, Jan. 2022, Art. no. 107659, doi: 10.1016/j.knsys.2021.107659.
14. R. Li, S. Shu, S. Wang, Y. Liu, Y. Li, and M. Peng, "DAT-MT accelerated graph fusion dependency parsing model for small samples in professional fields," *Entropy*, vol. 25, no. 10, p. 1444, Oct. 2023, doi: 10.3390/e25101444.
15. B. Shneiderman, *Human-Centered AI*. London, U.K.: Oxford Univ. Press, 2022. [Online]. Available: <https://global.oup.com/academic/product/human-centered-ai-9780192845290>
16. N. Balasubramaniam, M. Kauppinen, A. Rannisto, K. Hiekanen, and S. Kujala, "Transparency and explainability of AI systems: From ethical guidelines to requirements," *Inf. Softw. Technol.*, vol. 159, Jul. 2023, Art. no. 107197, doi: 10.1016/j.infsof.2023.107197.
17. C. van der Lee, A. Gatt, E. van Miltenburg, and E. Krahmer, "Human evaluation of automatically generated text: Current trends and best practice guidelines," *Comput. Speech Lang.*, vol. 67, May 2021, Art. no. 101151, doi: 10.1016/j.csl.2020.101151.
18. L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. L. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray, J. Schulman, J. Hilton, F. Kelton, L. Miller, M. Simens, A.

- Askell, P. Welinder, P. Christiano, J. Leike, and R. Lowe, “Training language models to follow instructions with human feedback,” 2022, arXiv:2203.02155.
19. Choi, H., Kim, S., & Park, J. (2022). Predictive analytics in market forecasting: Applications and challenges. *Journal of Business Analytics*, 9(3), 112–128.
20. Li, Y., & Chen, R. (2021). Sentiment analysis for real-time consumer insights. *International Journal of Marketing Science*, 14(2), 55–71.
21. Smith, D., & Gupta, A. (2023). Cognitive computing in market intelligence: Adaptive AI for strategic decision-making. *AI in Business Research*, 6(1), 23–40.
22. Barberis, N. C. 2013. “Thirty Years of Prospect Theory in Economics: A Review and Assessment.” *Journal of Economic Perspectives* 27, no. 1: 173–196.
23. Belkin, M., D. Hsu, S. Ma, and S. Mandal. 2019. “Reconciling Modern Machine-Learning Practice and the Classical Bias–Variance Trade-Off.” *Proceedings of the National Academy of Sciences* 116, no. 32: 15849–15854.
24. Bernheim, B. D., and C. Sprenger. 2020. “On the Empirical Validity of Cumulative Prospect Theory: Experimental Evidence of Rank-Independent Probability Weighting.” *Econometrica* 88, no. 4: 1363–1409.
25. Bourgin, D. D., J. C. Peterson, D. Reichman, S. J. Russell, and T. L. Griffiths. 2019. “Cognitive Model Priors for Predicting Human Decisions.” In *International Conference on Machine Learning*, 5133–5141. PMLR.