

FinNLP-2025

**The 10th Workshop on Financial Technology and
Natural Language Processing**

Proceedings of the Workshop

November 9, 2025

Suzhou, China

Preface

Welcome to FinNLP, a forum dedicated to fostering international collaboration and knowledge sharing in applying NLP to the dynamic domain of FinTech. We are especially delighted that the 10th edition of FinNLP coincides with the launch of the ACL Special Interest Group on Economic and Financial Natural Language Processing (SIG-FinTech). With SIG-FinTech, we aim to establish a sustainable presence for this research area within the ACL community. The mission of SIG-FinTech goes beyond advancing NLP research in finance and economics. We also seek to bring together experts from diverse fields to propose important research directions and exchange insights. Each year, we invite speakers from various disciplines to share their perspectives. This year, in addition to distinguished industry experts—Dr. Shi-Xiong (Austin) Zhang and Dr. Sambit Sahu from Capital One—we are honored to host Prof. Saeed Abdullah (Penn State), an expert in Human–Computer Interaction (HCI), and Prof. Andrea Rocci (Università della Svizzera italiana), a linguist, who will share their perspectives on financial information and applications from different disciplinary angles. Another highlight of FinNLP-2025 is the FinEval initiative, where we discuss methods for evaluating generated reports. A central challenge remains how to assess text quality in terms of its usefulness for user decision-making—a task that is still difficult to standardize. We hope that, starting with the 10th FinNLP, this conversation will pave the way toward more robust evaluation frameworks in the era of generative AI.

We are indebted to all the program committee members who dedicated significant time and expertise to provide insightful feedback on submissions and guide the selection process for FinNLP-2025. This year, we received 31 submissions to the main track and accepted 17. The accepted papers cover a broad spectrum of topics, including financial language modeling, evaluation frameworks, document auditing, sentiment analysis, fraud detection, and relation extraction. They also highlight advances in retrieval-augmented generation (RAG), multimodal analysis of charts and tables, and responsible AI practices in finance and law. We look forward to the stimulating discussions at this year’s FinNLP and hope you enjoy the workshop.

Chung-Chi Chen, Genta Indra Winata, Stephen Rawls, Anirban Das, Hsin-Hsi Chen,
Hiroya Takamura
FinNLP-2025 Organizers

Organizations

General Chairs

- Chung-Chi Chen - AIST, Japan
- Genta Indra Winata - Capital One, US
- Stephen Rawls - Capital One, US
- Anirban Das - Capital One, US
- Hsin-Hsi Chen - National Taiwan University, Taiwan
- Hiroya Takamura - AIST, Japan

Organizing Committee (Advisor)

- Sophia Ananiadou - University of Manchester, UK
- Hen-Hsen Huang - Academia Sinica, Taiwan
- Shi-Xiong Zhang - Capital One, US
- Supriyo Chakraborty - Capital One, US
- Sambit Sahu - Capital One, US

Program Committee (Reviewer)

- Mrinal Ahlawat - Google
- Paulo Alves - Católica Porto Business School
- Yung-Chun Chang - Taipei Medical University
- Emmanuele Chersoni - Aix-Marseille University
- Nelson Correa - Andinum, Inc.
- Giulia D'Agostino - Università della Svizzera Italiana
- Akshar Prabhu Desai - Indian Institute of Technology
- Pablo Duboue - Textualization Software Ltd.
- Himaanshu Gauba - Google
- Yu-Shiang Huang - National Taiwan University
- Kiyoshi Izumi - The University of Tokyo
- Karthik Kushala
- Chit-Kwan Lin - Revelata, Inc.
- Mohammad Luqman - Google
- Chenyang Lyu - MBZUAI
- Hilal Pataci - University of Texas at San Antonio
- Vijai Anand Ramar
- Hiroki Sakaji - Hokkaido University

- Takehiro Takayanagi - The University of Tokyo
- Ke Tian - Rakuten Group, Inc
- Kanji Uchino - FUJITSU RESEARCH OF AMERICA, INC.
- Chuan-Ju Wang - Academia Sinica
- Shih-Hung Wu - Chaoyang University of Technology
- Ajay Yadav - Google LLC
- Sin-Han Yang - Appier
- Annie T.T. Ying - IBM

Shared Task Organizers

- Takehiro Takayanagi - The University of Tokyo, Japan
- Tomas Goldsack - University of Sheffield, UK
- Kiyoshi Izumi - The University of Tokyo, Japan
- Chenghua Lin - University of Manchester, UK
- Hiroya Takamura - AIST, Japan
- Chung-Chi Chen - AIST, Japan

Table of Contents

LLM as a Guide: an Approach for Unsupervised Economic Relation Discovery in Administrative Documents

Thomas Sebbag, Solen Quiniou and Emmanuel Morin----- 1

Zero-Shot Extraction of Stock Relationship Graphs with LLMs

Hao Zhou, Luis Felipe Costa Sperb and Tiejun Ma----- 12

Enhancing Financial RAG with Agentic AI and Multi-HyDE: A Novel Approach to Knowledge Retrieval and Hallucination Reduction

Akshay Govind Srinivasan, Ryan George, Jayden Koshy Joe, Hrushikesh Kant, Harshith M R, Sachin Sundar, Sudharshan Suresh, Rahul Vimalkanth, and Vijayavallabh----- 19

StockGenChaR: A Study on the Evaluation of Large Vision-Language Models on Stock Chart Captioning

Le Qiu and Emmanuele Chersoni----- 33

FinEval-KR: A Financial Domain Evaluation Framework for Large Language Models' Knowledge and Reasoning

Shaoyu Dou, Yutian Shen, Mofan Chen, Zixuan Wang, Jiajie Xu, Qi Guo, Kailai Shao, Chao Chen, Haixiang Hu, Haibo Shi, Min Min and Liwen Zhang----- 47

Towards Efficient FinBERT via Quantization and Coreset for Financial Sentiment Analysis

Avinash Kumar Sharma, Aisha Hamad Hassan and Tushar Shinde----- 70

LAVA: Logic-Aware Validation and Augmentation Framework for Large-Scale Financial Document Auditing

Ruoqi Shu, Xuhui Wang, Isaac Wang, Yanming Mai and Bo Wan----- 75

FinCoT: Grounding Chain-of-Thought in Expert Financial Reasoning

Natapong Nitarach, Warit Sirichotedumrong, Panop Pitchayarthorn, Pittawat Taveekitworachai, Potsawee Manakul and Kunat Pipatanakul----- 93

Assessing RAG System Capabilities on Financial Documents

Oscar Lithgow-Serrano, David Kletz, Vani Kanjirangat, David Adametz, Marzio Lunghi, Claudio Bonesana, Matilde Tristany-Farinha, Yuntao Li, Detlef Replinger, Marco Pierbattista, Stefania Stan and Oleg Szehr-----124

Do Companies Reveal Their Own Fraud? - A Novel Data Set for Fraud Detection Based on 10-K Reports

Moustafa Amin and Matthias Aßenmacher----- 148

Synthesizing Behaviorally-Grounded Reasoning Chains: A Data-Generation Framework for Personal Finance LLMs

Akhil Theerthala----- 167

Detecting Evasive Answers in Financial Q&A: A Psychological Discourse Taxonomy and Lightweight Baselines	
Khaled Al Nuaimi, Gautier Marti, Alexis Marchal and Andreas Henschel-----	191
Enhancing Post Earnings Announcement Drift Measurement with Large Language Models	
Samuel Hadlock, Jesse Roberts and Joohun Lee-----	197
Natural Language Inference as a Judge: Detecting Factuality and Causality Issues in Language Model Self-Reasoning for Financial Analysis	
Yilin Wu, Han Yuan, Li Zhang and Zheng Ma-----	210
SEC-QA: A Systematic Evaluation Corpus for Financial QA	
Viet Dac Lai, Michael Krumdick, Charles Lovering, Varshini Reddy, Craig Schmidt and Chris Tanner-----	221
An Automatically Improving Method for Generating Descriptions of Financial Data Quality Grading with LLMs	
Yang Zhao, Yohei Ikawa and Bishwaranjan Bhattacharjee-----	237
Earnings2Insights: Analyst Report Generation for Investment Guidance	
Takehiro Takayanagi, Tomas Goldsack, Kiyoshi Izumi, Chenghua Lin, Hiroya Takamura and Chung-Chi Chen-----	246
Beyond Summaries: Multi-Agent Generation of Investment Reports with Text, Tables, and Charts	
Weijie Yang, Junbo Peng and Weijie Yang-----	252
From Earnings Calls to Investment Reports: Evaluating Role-based Multi-Agent LLM Systems	
Ranjan Satapathy, Raphael Liew, Joyjit Chattorj, Erik Cambria and Rick Goh-----	258
FinDebate: Multi-Agent Collaborative Intelligence for Financial Analysis	
Tianshi Cai, Guanxu Li, Nijia Han, Ce Huang, Zimu Wang, Changyu Zeng, Yuqi Wang, Jingshi Zhou, Haiyang Zhang, Qi Chen, Yushan Pan, Shuihua Wang and Wei Wang-----	268
Structured Adversarial Synthesis: A Multi-Agent Framework for Generating Persuasive Financial Analysis from Earnings Call Transcripts	
Saisab Sadhu, Biswajit Patra and Tannay Basu-----	283
Meta Prompting for Analyst Report Generation: Turning Earnings Calls into Investment Guidance	
Pulkit Chatwal, Mann Bajpai, Priyanshu Deswal, Harish Pratap Singh and Santosh Kumar Mishra-----	292
LangKG at the FinNLP 2025 - Earnings2Insights: Task-Adaptive LLMs To Generate Human-Persuasive Investment Reports	
Shivika Prasanna and Hui Su-----	297
Jetsons at the FinNLP-2025 - Earnings2Insights: Persuasive Investment Report Generation Using Single And Multi-Agent Frameworks	

Parag Dakle, Sai Krishna Rallabandi, Nikhi Kohli, Khyati Morparia, Ojas Raundale and Preethi Raghavan-----	305
Agentic LLMs for Analyst-Style Financial Insights: An LLM Pipeline for Persuasive Financial Analysis	
Gaurangi Sinha, Rajarajeswari Palacharla and Manoj Balaji Jagadeeshan -----	321
Multi-Agent Collaboration for Investment Guidance: Earnings2Insights Report Generation	
Mingrui Tan, Yang Liu, Kun Gao, Fei Gao and Yuting Song-----	327