Practitioner Brief

AI IN ASSET MANAGEMENT

Natural Language Processing

Francesco A. Fabozzi, PhD

Practitioner Brief written by Cathy Scott

CFA Institute Research Foundation

CFA Institute Research & Policy Center

Markets run on language. Every day, earnings calls, filings, press releases, research reports, and social media shape capital flows. For decades, natural language processing (NLP) has provided tools to turn text into usable signals. The rise of large language models (LLMs) marks a paradigm shift, extending the scope of NLP from narrow tasks to broad, general-purpose applications.

This chapter, part of *AI in Investment Management: Tools, Applications, and Frontiers*, explains how classical NLP techniques remain useful while LLMs unlock new frontiers in finance. Together, they allow practitioners to systematically process unstructured text, monitor risks, and generate insights at scale.

This brief distills the chapter's findings into practical guidance: who should read it, why it matters now, and how practitioners can use NLP and LLMs effectively.

Who Should Read This Chapter?

Portfolio managers and chief information officers (CIOs) will learn how to use NLP to monitor sentiment, ESG, and disclosures in real time. Risk officers, including chief risk officers (CROs) will learn how to use it to detect events and exposures earlier. For quants and data scientists, NLP can be applied to convert unstructured text into structured investment factors. Compliance leaders can use it to scale review of filings and monitor regulatory changes. Finally, C-suite executives will learn how to use NLP to set governance policies on AI adoption and deployment.

Why This Chapter Matters Now

For multiple reasons, NLP in finance has never been more important to understand:

- **Explosion of text data:** Filings, calls, and unstructured sources are multiplying faster than analysts can read.
- **Evolution of NLP:** From sentiment dictionaries to embeddings, NLP has long helped extract signals from language.
- **Generative AI shift:** LLMs expand NLP's reach by handling summarization, classification, and reasoning tasks with little training.
- **Competitive urgency:** Firms that integrate both classical NLP and LLMs will move faster on sentiment, compliance, and ESG signals.
- **Governance imperative:** Leaders must ensure models are reliable, explainable, and secure.

What Does This Chapter Deliver?

This chapter provides a clear, actionable guide to NLP in finance and its transformation under LLMs, covering the following:

- A history of NLP techniques, from dictionaries to transformers—showing continuity and what still works;
- An explanation of LLMs as a paradigm shift: pretrained, adaptable, and capable of general-purpose tasks;
- Applications with direct business value: compliance, ESG, risk monitoring, summarization, quant investing;
- Risks executives must weigh, including hallucinations, evaluation gaps, forward-bias in backtests, and data governance; and
- A roadmap for adoption that clarifies where classical NLP suffices and where LLMs offer unique advantages.

Practical Applications

- **Analyze sentiment and tone:** Use NLP to measure tone in earnings calls, press releases, and news. Apply FinBERT for finance-specific sentiment classification or OpenAI GPT application programming interfaces (APIs) for zero- or few-shot scoring when you lack labeled data. Run this analysis before and after earnings releases or to track changes in CEO and CFO language across quarters. Portfolio managers, quants, and strategists can use these signals to sharpen investment views.
- **Monitor ESG disclosures:** Extract and classify ESG statements from filings, sustainability reports, and media. Deploy models such as ESGBERT for classification, and use entity recognition to highlight mentions of climate, labor, or governance issues. Apply these tools during proxy season or when evaluating reputational and supply-chain risks. Chief investment officers (CIOs), ESG officers, and risk teams can strengthen stewardship and risk oversight.
- **Check compliance and regulatory filings:** Automate reviews of filings against current regulations. Combine retrieval-augmented generation (RAG) to fetch relevant rules with encoder models like BERT to compare new filings against prior versions. Apply this during 10-K/10-Q reviews or when monitoring shifts from regulators, such as the US Securities and Exchange Commission, the Financial Conduct Authority, or the European Securities and Markets Authority. Compliance leaders, chief risk officers (CROs), and legal teams can reduce manual workload while maintaining audit trails.
- **Scan streaming news for risk events:** Continuously monitor news and social media feeds for signals of market stress. Build pipelines with Hugging Face models that combine sentiment scoring with entity recognition to flag specific firms, sectors, or geographies. Use these systems for early warning of volatility spikes or to identify systemic exposures. Risk managers, traders, and CROs gain faster visibility into emerging threats.
- **Summarize financial documents:** Condense lengthy 10-Ks, 8-Ks, or earnings transcripts into actionable insights. Use ChatGPT-style models for fluent abstractive summaries or Sentence-BERT (SBERT) embeddings with extractive methods for structured bullet-point takeaways. Apply these tools during reporting season or for preparing quick internal memos. Analysts, portfolio managers, and CIOs free up time for higher-value judgment and decision-making.
- **Generate quantitative signals from text:** Turn unstructured documents into usable features for financial models. Use SBERT or FinBERT embeddings to convert text into vectors, then cluster them into interpretable "textual factors." Incorporate these factors into regression or forecasting models to improve return prediction or volatility analysis. Quants, data scientists, and CIOs can enhance model accuracy and explainability.

Case in Action: Sentiment Analysis for Financial Text

One of the most practical NLP use cases in finance is sentiment classification: identifying whether text is positive, negative, or neutral. The chapter illustrates this process with a hands-on example using both open-source models and a commercial API:

Dataset

- Financial Phrasebank (Malo et al., 2014): 4,800 sentences from financial news, each labeled as positive, negative, or neutral.
- Widely used in academia and practice to benchmark sentiment models.

Approach 1: Open-Source Pipeline

SBERT generates embeddings—dense vector representations of each sentence. A Ridge Classifier is then trained on these embeddings to predict sentiment labels.

- *Benefit:* Full control over the pipeline, explainability, and no data leaves the firm's environment.
- *Limitation:* Requires data science expertise, infrastructure, and labeled data.

Approach 2: API-Based Pipeline

OpenAI GPT-3.5 (or similar) is prompted with instructions: *"Classify this text as positive, neutral, or negative."* The model returns a sentiment score without any fine-tuning.

- *Benefit:* Rapid deployment, minimal setup, no training data required.
- *Limitation:* Data must be sent to a third-party API; costs can accumulate at scale.

Takeaway for Practitioners

Open-source models (such as SBERT or FinBERT) are ideal when control, customization, or compliance is paramount. API-based LLMs are attractive for speed, prototyping, or when labeled data are scarce. Many firms adopt a hybrid approach, using APIs for quick wins while developing internal models for sensitive or large-scale workflows.

Practitioner Toolkit

Role	Technique	Approaches	Practitioner Benefit
CIOs	Strategic integration of AI signals into investment processes	API-based LLMs for prototyping; open-source models (FinBERT, SBERT) for internal deployment	Faster insight generation; defensible communication to boards/regulators
Portfolio managers	Day-to-day market monitoring and efficiency gains	Document summarization platforms; tone analysis dashboards; embedding-based alerts	Reduced research workload; quicker reaction to disclosures and events
Risk leaders (CROs, risk managers)	Real-time event detection and systemic risk oversight	RAG pipelines connected to news/filings; monitoring dashboards with sentiment triggers	Early warning of shocks; stronger compliance posture
Data science teams	Model development and scaling	Hugging Face libraries; fine-tuning pipelines; quantization/distillation methods for efficiency	Robust models; lower infrastructure costs; integration of text with numeric signals
Compliance leaders	Automation of repetitive review tasks	Encoder models for document comparison; hybrid API/on-premise setups for secure audits	Scalable monitoring; reduced manual workload; defensible audit trails

The Bottom Line

NLP has been central to financial text analysis for decades. What is new is the transformative power of LLMs, extending NLP from narrow tools to flexible, general-purpose systems. Practitioners can use this combined toolkit today to perform the following:

- **Monitor** sentiment and tone in markets.
- **Surface** ESG and compliance issues at scale.
- **Detect** risk events faster.
- **Summarize** financial disclosures efficiently.
- **Generate** structured factors from text for quantitative investing.

The winning firms will recognize the continuity of NLP and the disruption of LLMs, embedding both into daily workflows with strong governance and clear strategic alignment.

Metrics That Matter

To move NLP from proof-of-concept to production, practitioners must quantify how models perform. The following metrics capture model accuracy, efficiency, and governance—core dimensions for evaluating NLP and LLM systems in finance.

- **F1 score:** Balance of precision and recall; critical for classification tasks (e.g., ESG disclosure detection).
- **ROUGE/BLEU scores:** Quality of text summarization; compare machine output with human reference.
- **Latency (milliseconds/seconds):** Processing speed; essential for real-time trading and risk monitoring.
- **Coverage:** Percentage of relevant documents/sources analyzed; a measure of breadth.
- **Human-in-the-loop checks:** Validation rate; percentage of model outputs confirmed/edited by humans.

Glossary

Embedding: Vector representation of text; foundation for clustering and factors.

FinBERT/SBERT: Variants of BERT fine-tuned for finance or sentence embeddings.

Natural language processing (NLP): AI field focused on analyzing human language.

Quantization/distillation: Techniques to shrink model size and speed up inference.

Retrieval-augmented generation (RAG): Combining search with generation for up-to-date answers.

Summarization: Condensing long text into shorter, meaningful summaries.

Textual factor: A theme or feature derived from text, usable in financial models.

Related Content

Cao, Larry, ed. 2023. *Handbook of Artificial Intelligence and Big Data Applications in Investments*. CFA Institute Research Foundation. https://rpc.cfainstitute.org/research/foundation/2023/ai-and-big-data-in-investments-handbook.

Pisaneschi, Brian. 2025. "RAG for Finance: Automating Document Analysis with LLMs." CFA Institute Research and Policy Center. https://rpc.cfainstitute.org/research/the-automation-ahead-content-series/retrieval-augmented-generation.

Spyrou, Alex, and Brian Pisaneschi. 2024. "Practical Guide for LLMs in the Financial Industry." CFA Institute Research and Policy Center. https://rpc.cfainstitute.org/research/the-automation-ahead-content-series/practical-guide-for-llms-in-the-financial-industry.

Wilson, Cheryll-Ann. 2025. "Explainable AI in Finance: Addressing the Needs of Diverse Stakeholders." CFA Institute Research and Policy Center. https://rpc.cfainstitute.org/research/reports/2025/explainable-ai-in-finance.

CFA Institute Research Foundation

CFA Institute Research & Policy Center

Find *AI in Asset Management: Tools, Applications, and Frontiers* and more on the CFA Institute Research and Policy Center website: https://rpc.cfainstitute.org.

Read the full chapter