



Deep Learning for Finance en 25 artículos
for the working analyst





Esta colección de **25 artículos** fue diseñada para acompañar el curso *Deep Learning for Finance*, una experiencia formativa avanzada en el uso de redes neuronales aplicadas a la información financiera moderna. Nuestra curaduría apunta a consolidar el conocimiento práctico, técnico y conceptual de los analistas cuantitativos que desean integrar técnicas de deep learning en la construcción de modelos para series temporales, detección de fraude, análisis textual, pricing de derivados, y finanzas descentralizadas.

A diferencia de nuestras otras selecciones, aquí **no se abordan fundamentos generales de estadística, modelos clásicos de ML o estructuras financieras tradicionales**. En cambio, nos enfocamos exclusivamente en los avances recientes y arquitecturas especializadas que han revolucionado el análisis financiero.

1.

Black, F., & Scholes, M. (1973): The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3), 637–654.

2.

Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019): Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear PDEs. *Journal of*

3.

Goodfellow, I., Bengio, Y., & Courville, A. (2016): *Deep Learning*. MIT Press.

4.

Hochreiter, S., & Schmidhuber, J. (1997): Long short-term memory. *Neural Computation*, 9(8), 1735–1780.

5.

Lim, B., et al. (2021): Temporal Fusion Transformers for interpretable multi-horizon time series forecasting. *International Journal of Forecasting*, 37(4), 1748–1764.

6.

Vaswani, A., et al. (2017): Attention is all you need. *NeurIPS*.

7.

Bahdanau, D., Cho, K., & Bengio, Y. (2015): Neural machine translation by jointly learning to align and translate. *ICLR*.

8.

Zhang, A., et al. (2020): A survey on deep learning applications in financial services. *Journal of Financial Data Science*, 2(2), 10–37.

9.

Sirignano, J., & Cont, R. (2019): Universal features of price formation in financial markets: perspectives from deep learning. *Quantitative Finance*, 19(9), 1449–1459.



10.

Krauss, C., Do, X. A., & Huck, N. (2017): Deep neural networks, gradient-boosted trees, random forests: Statistical arbitrage on the S&P 500. *European Journal of Operational Research*, 259(2), 689–702.

11.

Hamilton, J. D. (1994): *Time Series Analysis*. Princeton University Press.

12.

Scarselli, F., et al. (2009): The Graph Neural Network model. *IEEE Transactions on Neural Networks*, 20(1), 61–80.

13.

Wu, Z., et al. (2020): A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 32(1), 4–24.

14.

Zhang, Y., & Zohren, S. (2022): Deep reinforcement learning for trading: A review. *arXiv preprint arXiv:2107.09600*.

15.

Brown, T., et al. (2020): Language models are few-shot learners. *NeurIPS*.

16.

Lewis, P., et al. (2020): Retrieval-augmented generation for knowledge-intensive NLP tasks. *NeurIPS*.

17.

Zhao, W., et al. (2021): A survey of RAG: Retrieval-Augmented Generation for deep NLP. *arXiv preprint*.

18.

Zhou, Z., et al. (2023): FinGPT: Large language models for financial natural language processing. *arXiv preprint arXiv:2306.14852*.



19.

Mnih, V., et al. (2016): Asynchronous methods for deep reinforcement learning. ICML.

20.

Schulman, J., et al. (2017): Proximal Policy Optimization Algorithms. arXiv preprint.

21.

Jiang, Z., & Liang, J. (2017): Cryptocurrency portfolio management with deep reinforcement learning. arXiv preprint.

22.

Chen, T., et al. (2019): Applications of deep learning to risk management and financial forecasting. Journal of Financial Data Science, 1(1), 10–24.

23.

Borovkova, S. (2021): Machine Learning for Asset Managers. Risk Books.

24.

Hiransha, M., et al. (2018): NSE stock market prediction using deep-learning models. Procedia Computer Science, 132, 1351–1362.

25.

LeCun, Y., Bengio, Y., & Hinton, G. (2015): Deep learning. Nature, 521(7553), 436–444.





colegio-bourbaki.com