INDEX

CERTIFIED ARTIFICIAL INTELLIGENCE PROFESSIONAL TRAINING COURSE

Table of contents

I. Detailed agenda of the training course

Day 1 Foundations of AI and data analysis

Day 2 Machine learning

Day 3 Deep Learning

Day 4 NLP, LLMs, RAG and AI Governance

- AI project management and implementation

Day 5 Certification exam

II. Standards cited in this training course

- ISO/IEC 22989:2022, Information technology — Artificial intelligence — Artificial intelligence concepts and terminology
- EU Artificial Intelligence Act

III. Bibliography

Section 3: Data analysis and visualization

[1] Provost, Foster, and Tom Fawcett. *Data Science for Business: What You Need to Know About Data Mining and Data-Analytic Thinking*. O'Reilly Media, 2013.

[2] Kirk, Andy. *Data Visualisation: A Handbook for Data Driven Design*. Sage Publications, 2019.

[3] Shmueli, Galit, Peter C. Bruce, Inbal Yahav, Nitin R. Patel, and Kenneth C. Lichtendahl Jr. *Data Mining for Business Analytics: Concepts, Techniques, and Applications in R*. John Wiley & Sons, 2018.

[4] Crabtree, Matt, and Adel Nehme. "What is Data Analysis? An Expert Guide with Examples." DataCamp, November 10, 2024. https://www.datacamp.com/blog/what-is-data-analysis-expert-guide.

[5] Cooper, R, Donald, and Pamela S. Schindler. *Business Research Methods*, 12th ed. McGraw-Hill Companies, 2014.

[6] "Methods of Data Collection." RudderStack, n.d. https://www.rudderstack.com/learn/data-collection/methods-of-data-collection/.

[7] Rahm, Erhard, and Hong Hai Do. "Data Cleaning: Problems and Current Approaches." *IEEE Data Engineering Bulletin*, 2000.

[8] Jain, Mradul, and Oun Muhammad. "What Tools and Frameworks Can You Use for Data Cleaning and Preparation?" *LinkedIn*, October 5, 2023. https://www.linkedin.com/advice/1/what-tools-frameworks-can-you-use-data-cleaning.

[9] "Guide to Data Cleaning: Definition, Benefits, Components, And How To Clean Your Data." *Tableau*, n.d. https://www.tableau.com/learn/articles/what-is-data-cleaning.

[10] Tukey, W, John. *Exploratory Data Analysis*. Addison-Wesley, 1977.

[11] LIM, Sze Zhong. "Mastering Exploratory Data Analysis (EDA): Everything You Need to Know." *Medium*, April 15, 2024. https://medium.com/data-and-beyond/mastering-exploratory-data-analysis-eda-everything-you-need-to-know-7e3b48d63a95.

[12] Murphy, K. Kevin. *Machine Learning: A Probabilistic Perspective*. The MIT Press. 2012.

[13] Ghadaaissa. "The Basics of Model Evaluation." *Medium*, February 15, 2024. https://medium.com/@ghadaaissa16/the-basics-of-model-evaluation-8356b986c963.

[14] "10 Data Analysis Tools and When to Use Them." *Coursera*, n.d. https://www.coursera.org/articles/data-analysis-tools.

[15] Shook, Jim, Robyn Smith, and Alex Antonio. "Transparency and Fairness in Machine Learning Applications." *Transparency and Fairness in Machine Learning Applications* 4 (2018): 443–63. https://doi.org/10.37419/jpl.v4.i5.2.

[16] Marini, Alice, Alexis Kateifides, Joel Bates, Gabriela Zanfir Fortuna, Michelle Bae, and Gargi Sen. “GDPR v. CCPA - The Future of Privacy Forum.” European Commission, n.d. https://ec.europa.eu/futurium/en/system/files/ged/gdpr_ccpa_comparison-guide.pdf.

[17] Ferrara, Emilio. “FAIRNESS AND BIAS IN ARTIFICIAL INTELLIGENCE: A BRIEF SURVEY OF SOURCES, IMPACTS, AND MITIGATION STRATEGIES.” USC Viterbi School of Engineering, 2023.

[18] Kotthoff, Lars, and Joaquin Vanschoren. “Automated Machine Learning.” Edited by Frank Hutter. SpringerLink, April 18, 2024. https://link.springer.com/book/10.1007/978-3-030-05318-5.

[19] Goyle, Kartikay, Quin Xie, and Vakul Goyle. *DataAssist: A Machine Learning Approach to Data Cleaning and Preparation*, July 17, 2023. https://doi.org/https://arxiv.org/pdf/2307.07119.

[20] Liu, Binsheng, Xiaolu Lu, and J. Shane Culpepper. “Strong Natural Language Query Generation.” Information Retrieval Journal 24 (July 15, 2021): 322–46. https://doi.org/10.1007/s10791-021-09395-3.

[21] Chen, Chun-houh, Wolfgang Härdle, and Antony Unwin. *Handbook of Data Visualization*. Springer, 2008.

[22] Unwin, A. Why Is Data Visualization Important? What Is Important in Data Visualization? *Harvard Data Science Review*, 2020. https://doi.org/10.1162/99608f92.8ae4d525.

[23] Noorain, Afra. 2024. “Fundamentals of Data Visualization Explained.” *Sprintzeal*, June 21, 2024. https://www.sprintzeal.com/blog/fundamentals-of-data-visualization.

[24] Knaflic, Cole Mussbaumer. *Storytelling With Data: A Data Visualization Guide For Business Professionals*. John Wiley & Sons, 2015.

[25] “What Is a Bar Chart?” *Jaspersoft*, n.d. https://www.jaspersoft.com/articles/what-is-a-bar-chart.

[26] CFI Team. “Bar Charts (Data Visualization and Technical Analysis).” *Corporate Finance Institute*. https://corporatefinanceinstitute.com/resources/business-intelligence/bar-charts-data-visualization-and-technical-analysis/.

[27] “What Is a Line Chart?” *Jaspersoft*, n.d. https://www.jaspersoft.com/articles/what-is-a-line-chart.

[28] Cisneros, Mike. “What Is a Line Graph? — Storytelling With Data.” *Storytelling With Data*. July 24, 2025. https://www.storytellingwithdata.com/blog/2020/3/24/what-is-a-line-graph.

[29] Jain, Anik, Batsheva Yanir, and Durgesh Kewat. “How Can You Use Visual Hierarchy to Improve Data Visualization Readability?” How Visual Hierarchy Improves Data Visualization Readability, *LinkedIn*. February 11, 2024. https://www.linkedin.com/advice/0/how-can-you-use-visual-hierarchy-improve-data-visualization 1mr5e#:~:text=Visual%20hierarchy%20is%20the%20principle,details%20of%20your%20data%20visualization.

[30] Salman, Mahnoor. “Color Theory for Data Visualizers.” *Medium*, July 17, 2024. https://medium.com/@mahnoorsalman96/color-theory-for-data-visualizers-782f59e87394.

[31] Treesak, Yanisa. “10 Common Data Visualization Mistakes and How to Avoid Them.” *Medium*, July 9, 2024. https://medium.com/agoda-engineering/10-common-data-visualization-mistakes-and-how-to-avoid-them-e3896fe8e104.

[32] Habib, JABERI Mohamed. “Choosing the Right Tool: Matplotlib vs. Pandas vs. Seaborn.” *Medium*, May 25, 2024. https://medium.com/@jaberi.mohamedhabib/choosing-the-right-tool-matplotlib-vs-pandas-vs-seaborn-4468b6b3c0a5.

[33] Biswal, Avijeet. “Tableau vs Power Bi: Which Is Better Data Visualization Tool.” Simplilearn, February 26, 2025. https://www.simplilearn.com/tutorials/power-bi-tutorial/power-bi-vs-tableau.

Section 4: Foundations of data science and machine learning

[1] IBM. "What Is Data Science?". Accessed March 26, 2025. https://www.ibm.com/topics/data-science.

[2] Cady, Field. *The Data Science Handbook.* John Wiley & Sons, Inc., 2017.

[3] Mitchell, Tom M. Machine Learning. McGraw-Hill, 1997.

[4] Géron, Aurélien. *Hands-On Machine Learning with Scikit-Learn, Keras, & TensorFlow*. 2nd ed. O’Reilly Media Inc., 2019.

[5] Burkov, Andriy. *The Hundred-Page Machine Learning Book*. 2019

[6] Voleti, Ritvik. “Unfolding the Evolution of Machine Learning and its Expediency” International Journal of Computer Science and Mobile Computing 10 no. 1, (2021): 1-7. 10.47760/ijcsmc.2021.v10i01.001.

[7] Russell, Stuart, and Peter Norvig. *Artificial Intelligence: A Modern Approach*. 4th ed. Pearson Education Limited, 2022.

[8] Hudson, Eddie. "Machine Learning Development vs Traditional Programming: A Comparison." *Divelement*, April 29, 2024. https://divelement.io/blog/machine-learning-development-vs-traditional-programming/.

[9] Bradley University. "What's the Difference Between Descriptive and Inferential Statistics?" Accessed March 26, 2025. https://onlinedegrees.bradley.edu/blog/whats-the-difference-between-descriptive-and-inferential-statistics.

[10] Anderson, David R., and Thomas A. Williams. "Random Variables and Probability Distributions." Encyclopaedia Britannica. Last modified February 1, 2025. https://www.britannica.com/science/statistics/Random-variables-and-probability-distributions.

[11] Orloff, Jeremy, and Jonathan Bloom. *Conditional Probability, Independence and Bayes’ Theorem*. 2022.

[12] Wood, Roy. *Business Applications of Artificial Intelligence and Machine Learning*. Oklahoma City. 2024

[13] Newsom, James. "Lecture 1: Types of Scales & Levels of Measurement.“ Last accessed March 26, 2025. https://web.pdx.edu/~newsomj/pa551/lecture1.htm.

[14] Jansen, Derek. "Nominal, Ordinal, Interval & Ratio Data." Grad Coach, November 2020. https://gradcoach.com/nominal-ordinal-interval-ratio/.

[15] GraphPad. "What Is the Difference Between Ordinal, Interval, and Ratio Variables? Why Should I Care?“ Last accessed March 26, 2025. https://www.graphpad.com/support/faq/what-is-the-difference-between-ordinal-interval-and-ratio-variables-why-should-i-care/.

[16] Chavan, Rahul. "Understanding Train, Test, and Validation Split in Simple Quick Terms." Medium, June 24, 2023. https://medium.com/@rahulchavan4894/understanding-train-test-and-validation-dataset-split-in-simple-quick-terms-5a8630fe58c8.

Section 5: Machine learning workflow

[1] Géron, Aurélien. *Hands-On Machine Learning with Scikit-Learn, Keras, & TensorFlow*. 2nd ed. O'Reilly Media, Inc., 2019.

[2] Cady, Field. *The Data Science Handbook*. John Wiley & Sons, Inc., 2017.

[3] Team Businesspedia. “Forecasting: Meaning, Nature, Planning and Forecasting, Importance and Limitations.” *Businesspedia*. February 16, 2025. https://businesspedia.in/forecasting-meaning-nature-planning-and-forecasting-importance-and-limitations/.

[4] Li, Jamie. “Data Cleaning With Python Pandas.” *OSEDEA*. n.d. https://www.osedea.com/insight/data-cleaning-with-python.

[5] The Ahn, Pham. “Feature Engineering: A Comprehensive Guide to Building Better Machine Learning Models”. *Medium*. January 8, 2025. https://medium.com/funny-ai-quant/feature-engineering-a-comprehensive-guide-to-building-better-machine-learning-models-84077458df23.

[6] Chawla, Avi. 2024. “Reduce Trees in Random Forest Model.” *Daily Dose of Data Science*. August 3, 2024. https://blog.dailydoseofds.com/p/reduce-trees-in-random-forest-model.

[7] Harsh, Manish, and Judy Lee. 2023. “Scaling AI With MLOps and the NVIDIA Partner Ecosystem.” *NVIDIA Technical Blog*, March 8, 2023. https://developer.nvidia.com/blog/scaling-ai-with-mlops-and-the-nvidia-partner-ecosystem/.

Section 6: Supervised learning

[1] Géron, Aurélien. *Hands-On Machine Learning with Scikit-Learn, Keras, & TensorFlow*. 2nd ed. O'Reilly Media, Inc., 2019.

[2] Montgomery, Douglas. C., Elizabeth A. Peck, and Geoffrey Vining. *Introduction to Linear Regression Analysis*. 6th ed. Wiley. 2021.

[3] Hosmer, David. W., Stanley Lemeshow, and Rodney X. Sturdivant. *Applied Logistic Regression*. 3rd ed. Wiley. 2013.

[4] Murphy, K. Kevin. *Machine Learning: A Probabilistic Perspective*. The MIT Press. 2012.

[5] Raghavan, Hemalatha. "K-Nearest Neighbor (KNN).” Last accessed June 4, 2024. https://cloud2data.com/k-nearest-neighbour/.

[6] Abi. "k-Nearest Neighbors." Machine-Learning-Course 1.0 Documentation. Last accessed March 26, 2025. https://machine-learning-tutorial-abi.readthedocs.io/en/latest/content/supervised/knn.html.

[7] Tech & Tales. “Navigating the Path to Understanding Decision Trees in Machine Learning.” *Medium*. September 9, 2023. https://techntales.medium.com/navigating-the-path-to-understanding-decision-trees-in-machine-learning-63459dc19ef3.

[8] Gunay, Deniz. “Random Forest.” *Medium*, September 16, 2023. https://medium.com/@denizgunay/random-forest-af5bde5d7e1e.

[9] Deng, Haowen, Youyou Zhou, Lin Wang, and Cheng Zhang. "Ensemble Learning for the Early Prediction of Neonatal Jaundice with Genetic Features." BMC Medical Informatics and Decision Making 21 (2021): Article 10.1186/s12911-021-01701-9. https://doi.org/10.1186/s12911-021-01701-9.

[10] Mateusz Buda, Atsuto Maki, and Maciej A Mazurowski. A systematic study of the class imbalance problem in convolutional neural networks. Neural Networks, 106:249–259, 2018. https://arxiv.org/pdf/1710.05381.

Section 7: Unsupervised learning

[1] Géron, Aurélien. *Hands-On Machine Learning with Scikit-Learn, Keras, & TensorFlow*. 2nd ed. O’Reilly Media, Inc., 2019.

[2] Mindy Support. “Unlocking the Power of Unlabeled Data in Machine Learning.” Last modified March 5, 2024. https://mindy-support.com/news-post/unlocking-the-power-of-unlabeled-data-in-machine-learning/.

[3] Little Dino. “Introduction to K-means clustering.” *Medium*, April 29, 2022. https://medium.com/@24littledino/introduction-to-k-means-clustering-976dc34fc4d5.

[4] Ain, Supriyo. “What is Association Rule Learning?” Medium, July 06, 2023. https://medium.com/@ainsupriyofficial/what-is-association-rule-learning-a6ef399fdc01.

[5] Kumar, Sanjay. “Evaluating Clustering Algorithms: A Comprehensive Guide to Metrics“ Linkedin, January 16, 2024. https://www.linkedin.com/pulse/evaluating-clustering-algorithms-comprehensive-guide-mba-ms-phd-mrvrc.

Section 8: Advanced ML and broader applications

[1] Sutton, Richard S., and Andrew G. Barto. *Reinforcement Learning: An Introduction*. 2nd ed. Cambridge, MA: MIT Press, 2018.

[2] Van Loon, Ronald. “Machine Learning Explained: Understanding Supervised, Unsupervised, and Reinforcement Learning.“ *LinkedIn*, January 15, 2018. https://www.linkedin.com/pulse/machine-learning-explained-understanding-supervised-ronald-van-loon.

[3] Lapan, Maxim. *Deep Reinforcement Learning Hands-On: A Practical and Easy-to-Follow Guide to RL from Q-Learning and DQNs to PPO and RLHF*. 3rd ed. Birmingham, UK: Packt Publishing, 2024.

[4] "Use AWS DeepRacer to Explore Reinforcement Learning - AWS DeepRacer." n.d. https://docs.aws.amazon.com/deepracer/latest/developerguide/deepracer-is-a-learning-environment-for-reinforcement-learning.html.
[5] Murphy, K. Kevin. *Machine Learning: A Probabilistic Perspective*. The MIT Press. 2012.
[6] NVIDIA. "What Is a Recommendation System?" *NVIDIA Glossary*. Accessed March 26, 2025. https://www.nvidia.com/en-us/glossary/recommendation-system/.
[7] Salesforce Marketing Cloud. Predictive Intelligence Benchmark Report. 2014. https://brandcdn.exacttarget.com/sites/exacttarget/files/deliverables/etmc-predictiveintelligencebenchmarkreport.pdf.
[8] Intel Corporation. "PayPal Solves Fraud Challenges." Accessed March 26, 2025. https://www.intel.com/content/www/us/en/customer-spotlight/stories/paypal-customer-story.html.
[9] Paypal Editorial Staff. "Harnessing Machine Learning Fraud Detection Technologies." *PayPal*, November 5, 2024. https://www.paypal.com/us/brc/article/payment-fraud-detection-machine-learning.
[10] "How Amazon Is Using Generative AI to Improve Product Recommendations and Descriptions." About Amazon, September 19, 2024. https://www.aboutamazon.com/news/retail/amazon-generative-ai-product-search-results-and-descriptions.
[11] "Top 30 Machine Learning Case Studies [2025]." DigitalDefynd. Accessed March 26, 2025. https://digitaldefynd.com/IQ/machine-learning-case-studies/.

Section 9: Foundations of deep learning

[1] Nielsen, Michael. "Using neural nets to recognize handwritten digits ." In *Neural Networks and Deep Learning*, 2019
[2] Aggarwal, Charu C. "An Introduction to Neural Networks." *Neural networks and deep learning: A textbook*. New York: Springer, 2019
[3] Ardiansyah, Sendy. "Artificial Intelligence vs. Human Brain — A Love-Hate Relationship." Medium, August 7, 2023. https://sendyardiansyah.medium.com/artificial-intelligence-vs-human-brain-a-love-hate-relationship-bc97090a66ee
[4] Ronaghan, Stacey. "Deep Learning: Overview of Neurons and Activation Functions." Medium, July 26, 2018. https://srnghn.medium.com/deep-learning-overview-of-neurons-and-activation-functions-1d98286cf1e4
[5] Murphy, Kevin P. Machine learning: A probabilistic perspective. Cambridge: The MIT Press, 2012
[6] Gad, Ahmed. "A Comprehensive Guide to the Backpropagation Algorithm in Neural Networks." neptune.ai, August 22, 2023. https://neptune.ai/blog/backpropagation-algorithm-in-neural-networks-guide

[7] James, Gareth, Daniela Witten, Trevor Hastie, and Robert Tibshirani. “Statistical Learning.” Essay. In *An Introduction to Statistical Learning with Applications in R*. Springer, 2013

[8] Myers, Chris. “AI with Deep Learning.” Cornell University, June 2023. https://cvw.cac.cornell.edu/DeepLearning/intro/hardware#:~:text=Deep%20learning%20methods%20can%20be,targeted%20specifically%20for%20deep%20learning

[9] Goodfellow, Ian, Yoshua Bengio, and Aaron Courville. *Deep Learning*. Cambridge: MIT Press, 2016.

Section 10: Core deep learning architectures

[1] Géron, Aurélien. “Introduction to Artificial Neural Networks with Keras.” *Hands-on Machine Learning with Scikit-Learn, Keras & TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems*. O’Rielly Media, 2019

[2] Aggarwal, Charu C. “An Introduction to Neural Networks” *Neural networks and deep learning*: A textbook. New York: Springer, 2019.

[3] “Convolutional Neural Network Basics: What You Need to Know.” Grammarly, September 6, 2024. https://www.grammarly.com/blog/ai/what-is-a-convolutional-neural-network/.

[4] Haque, Kh. Nafizul. “What Is Convolutional Neural Network - CNN (Deep Learning).” LinkedIn, April 3, 2023. https://www.linkedin.com/pulse/what-convolutional-neural-network-cnn-deep-learning-nafiz-shahriar

[5] Kumar, Ajitesh. “Different Types of CNN Architectures Explained: Examples.” Analytics Yogi, December 4, 2023. https://vitalflux.com/different-types-of-cnn-architectures-explained-examples/ .

[6] Huang, Gao, Zhuang Liu, Laurens van der Maaten, and Kilian Q Weinberger. Densely Connected Convolutional Networks, January 18, 2018. https://doi.org/arXiv:1608.06993

[7] Gad, Ahmed. “A Comprehensive Guide to the Backpropagation Algorithm in Neural Networks.” neptune.ai, August 22, 2023. https://neptune.ai/blog/backpropagation-algorithm-in-neural-networks-guide.

[8] Kotlar, Milos, Marija Punt, Zaharije Radivojevic, Milos Cvetanovic, and Veljko Milutinovic. “Novel Meta-Features for Automated Machine Learning Model Selection in Anomaly Detection.” IEEE Access 9 (2021). https://doi.org/10.1109/access.2021.3090936.

[9] Schmitt, Marc. “Explainable Automated Machine Learning for Credit Decisions: Enhancing Human Artificial Intelligence Collaboration in Financial Engineering.” *SSRN Electronic Journal*, 2024. https://doi.org/10.2139/ssrn.4716847.

[10] Sun, Yuan, Qiurong Song, Xinning Gui, Fenglong Ma, and Ting Wang. “Automl in the Wild: Obstacles, Workarounds, and Expectations.” *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, April 19, 2023. https://doi.org/10.1145/3544548.3581082.

[11] Jain, Ojasvi, and Kaushlesh Kumar. “AutoML: Automating the Machine Learning Pipeline.” Mphasis, n.d. https://www.mphasis.com/content/dam/mphasis-com/global/en/home/innovation/next-lab/thoughtleadership/auto-ml-whitepaper.pdf

[12] Bengio, Yoshua, Patrice Simard, and Paolo Frasconi. "Learning Long-Term Dependencies with Gradient Descent Is Difficult." *IEEE Transactions on Neural Networks* 5, no. 2

[13] Hochreiter, Sepp, and Jürgen Schmidhuber. "Long Short-Term Memory." *Neural Computation* 9, no. 8

[14] Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. "Attention Is All You Need." In *Advances in Neural Information Processing Systems* 30. Curran Associates, 2017.

[15] Radford, Alec, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. "Improving Language Understanding by Generative Pre-Training." OpenAI Technical Report, 2018.

Section 11: Generative architectures

[1] Géron, Aurélien. “Introduction to Artificial Neural Networks with Keras.” *Hands-on Machine Learning with Scikit-Learn, Keras & TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems*. O’Rielly Media, 2019

[2] Aggarwal, Charu C. “An Introduction to Neural Networks” *Neural networks and deep learning*: A textbook. New York: Springer, 2019.

[3] “Convolutional Neural Network Basics: What You Need to Know.” Grammarly, September 6, 2024. https://www.grammarly.com/blog/ai/what-is-a-convolutional-neural-network/.

[4] Haque, Kh. Nafizul. “What Is Convolutional Neural Network - CNN (Deep Learning).” LinkedIn, April 3, 2023. https://www.linkedin.com/pulse/what-convolutional-neural-network-cnn-deep-learning-nafiz-shahriar

[5] Kumar, Ajitesh. “Different Types of CNN Architectures Explained: Examples.” Analytics Yogi, December 4, 2023. https://vitalflux.com/different-types-of-cnn-architectures-explained-examples/ .

[6] Huang, Gao, Zhuang Liu, Laurens van der Maaten, and Kilian Q Weinberger. Densely Connected Convolutional Networks, January 18, 2018. https://doi.org/arXiv:1608.06993

[7] Gad, Ahmed. “A Comprehensive Guide to the Backpropagation Algorithm in Neural Networks.” neptune.ai, August 22, 2023. https://neptune.ai/blog/backpropagation-algorithm-in-neural-networks-guide.

[8] Kotlar, Milos, Marija Punt, Zaharije Radivojevic, Milos Cvetanovic, and Veljko Milutinovic. “Novel Meta-Features for Automated Machine Learning Model

Selection in Anomaly Detection.” IEEE Access 9 (2021). https://doi.org/10.1109/access.2021.3090936.
[9] Schmitt, Marc. “Explainable Automated Machine Learning for Credit Decisions: Enhancing Human Artificial Intelligence Collaboration in Financial Engineering.” *SSRN Electronic Journal*, 2024. https://doi.org/10.2139/ssrn.4716847.
[10] Sun, Yuan, Qiurong Song, Xinning Gui, Fenglong Ma, and Ting Wang. “Automl in the Wild: Obstacles, Workarounds, and Expectations.” *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, April 19, 2023. https://doi.org/10.1145/3544548.3581082.
[11] Jain, Ojasvi, and Kaushlesh Kumar. “AutoML: Automating the Machine Learning Pipeline.” Mphasis, n.d. https://www.mphasis.com/content/dam/mphasis-com/global/en/home/innovation/next-lab/thoughtleadership/auto-ml-whitepaper.pdf
[12] Goodfellow, Ian, Yoshua Bengio, and Aaron Courville. *Deep Learning*. Cambridge: MIT Press, 2016.

[13] Kingma, Diederik P., and Max Welling. "Auto-Encoding Variational Bayes." In *Proceedings of ICLR 2014*. arXiv:1312.6114

[14] Ho, Jonathan, Ajay Jain, and Pieter Abbeel. "Denoising Diffusion Probabilistic Models." In *Advances in Neural Information Processing Systems* 33 (2020). arXiv:2006.11239.

[15] Bond-Taylor, Sam, Adam Leach, Yang Long, and Chris G. Willcocks. "Deep Generative Modelling: A Comparative Review of VAEs, GANs, Normalizing Flows, Energy-Based and Autoregressive Models." *IEEE Transactions on Pattern Analysis and Machine Intelligence*

Section 12: Computer vision

[1] Szeliski, Richard. *Computer Vision: Algorithms and Applications*. 2nd ed. Springer, 2022.
[2] Solem, Jan Erik. *Programming Computer Vision with Python*. 2012.
[3] Elgendy, Mohamed. *Deep Learning for Vision Systems*. Manning Publications Co. New York. 2020.
[4] “Tesla Vision: Tesla Self-Driving ADAS features.” *Car ADAS* (blog), April 1, 2025. https://caradas.com/tesla-vision-adas-features/.
[5] “Computer Vision at Tesla for Self-Driving Cars.” *Think Autonomus*, September 15, 2022. https://www.thinkautonomous.ai/blog/computer-vision-at-tesla/.
[6] Buolamwini, Joy. “Gender Shades.” Massachusetts Institute of Technology. Last accessed March 27, 2025. https://www.media.mit.edu/projects/gender-shades/overview/.

[7] He, Kaiming, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. "Masked Autoencoders Are Scalable Vision Learners." In *Proceedings of CVPR 2022*. arXiv:2111.06377.

Section 13: Deep learning and future directions

[1] "Unveiling Deep Learning: Exploring Its Real-World Applications and Overcoming Key Challenges." LinkedIn, December 6, 2024. https://www.linkedin.com/pulse/unveiling-deep-learning-exploring-its-real-world-6ktnc.

[2] "What Is Deep Learning?." Google, n.d. https://cloud.google.com/discover/what-is-deep-learning?hl=en.

[3] Shetty, Shravya. "How AI Supports Early Disease Detection in India." Google, March 19, 2024. https://blog.google/technology/health/google-ai-india-early-disease-detection/.

[4] "AI in Mammography: Improving Breast Cancer Screening with Artificial Intelligence." Google Health, n.d. https://health.google/caregivers/mammography/.

[5] "AI & Robotics." Tesla, n.d. https://www.tesla.com/AI.

[6] "Introducing ChatGPT". 2022. https://openai.com/index/chatgpt.

[7] Esteva, Andre, Brett Kuprel, Roberto A. Novoa, Justin Ko, Susan M. Swetter, Helen M. Blau, and Sebastian Thrun. "Dermatologist-Level Classification of Skin Cancer with Deep Neural Networks."

Section 14: Natural language processing

[1] Lane, Hobson, and Maria Dyshel. *Natural language processing in action.* Shelter Island, NY: Manning Publications, 2023.

[2] Bird, Steven, Ewan Klein, and Edward Loper. "Analyzing the Meaning of Sentences." Essay. In *Natural Language Processing with Python*. O'Rielly Media, 2009.

[3] Tunstall, Lewis, Leonardo von Werra, and Thomas Wolf. "Hello Transformers." Essay. In *Natural Language Processing with Transformers: Building Language Applications with Hugging Face*. O'Reilly Media, 2022.

[4] Aggarwal, Charu C. "Recurrent Neural Networks." *Neural networks and deep learning: A textbook.* New York: Springer, 2019.

[5] Jurafsky, Daniel, and James H. Martin. *Speech and Language Processing*. 3rd ed. draft. Stanford University, 2024. https://web.stanford.edu/~jurafsky/slp3/.

[6] Costa-jussà, Marta R., James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Heffernan, Elahe Kalbassi, et al. "No Language Left Behind: Scaling Human-Centered Machine Translation." arXiv:2207.04672, 2022.

Section 15: Large language models

[1] Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. "Attention Is All You Need." In *Advances in Neural Information Processing Systems*30. Curran Associates, 2017.

[2] Radford, Alec, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. "Language Models Are Unsupervised Multitask Learners." OpenAI Technical Report, 2019.

[3] Wei, Jason, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, et al. "Emergent Abilities of Large Language Models." *Transactions on Machine Learning Research*, 2022. arXiv:2206.07682.

[4] Kaplan, Jared, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Ramesh Barak, Jack Clark, et al. "Scaling Laws for Neural Language Models." arXiv:2001.08361, 2020.

[5] Radford, Alec, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. "Language Models Are Unsupervised Multitask Learners." OpenAI Technical Report, 2019.

[6] Wei, Jason, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. "Chain-of-Thought Prompting Elicits Reasoning in Large Language Models." In *Advances in Neural Information Processing Systems* 35 (2022). arXiv:2201.11903.

[7] Min, Sewon, Xinxi Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. "Rethinking the Role of Demonstrations: What Makes In-Context Learning Work?" In *Proceedings of EMNLP 2022*. arXiv:2202.12837.

[8] https://hai.stanford.edu/ai-index/2025-ai-index-report/technical-performance

Section 16: Retrieval-Augmented generation

[1] Lewis, Patrick, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, et al. "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks." In *Advances in Neural Information Processing Systems* 33 (2020). arXiv:2005.11401.

[2] Karpukhin, Vladimir, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. "Dense Passage Retrieval for Open-Domain Question Answering." In *Proceedings of EMNLP 2020*. arXiv:2004.04906.

[3] Gao, Yunfan, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, et al. "Retrieval-Augmented Generation for Large Language Models: A Survey." arXiv:2312.10997, 2023.

[4] OpenAI, " Morgan Stanley uses AI evals to shape the future of financial services”, https://openai.com/index/morgan-stanley/

Section 17: Agentic AI systems

[1] Wang, Lei, Chen Ma, Xueyang Feng, Zeyu Zhang, Hao Yang, Jingsen Zhang, Zhiyuan Chen, et al. "A Survey on Large Language Model Based Autonomous Agents." arXiv:2308.11432, 2023.

[2] Yao, Shunyu, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shfran, Karthik Narasimhan, and Yuan Cao. "ReAct: Synergizing Reasoning and Acting in Language Models." In *Proceedings of ICLR 2023*. arXiv:2210.03629.

[3] Yao, Shunyu, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shfran, Karthik Narasimhan, and Yuan Cao. "ReAct: Synergizing Reasoning and Acting in Language Models." In *Proceedings of ICLR 2023*. arXiv:2210.03629.

[4] OpenAI. “Function Calling.” OpenAI API Documentation. https://developers.openai.com/api/docs/guides/function-calling

[5] Yao, Shunyu, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shfran, Karthik Narasimhan, and Yuan Cao. "ReAct: Synergizing Reasoning and Acting in Language Models." In *Proceedings of ICLR 2023*. arXiv:2210.03629.

[6] Batchelder, Eli, Ben Hylak, and collaborators. "Dive into Claude Code: The Design Space of Today's and Future AI Agent Systems." arXiv:2604.14228, 2026.

[7] Salesforce. "Salesforce Unveils Agentforce — What AI Was Meant to Be." Salesforce News, September 12, 2024. https://www.salesforce.com/news/press-releases/2024/09/12/agentforce-announcement/

Section 18: AI security

[1] “OWASP Top 10 for LLM Applications 2025.” Creative Commons. 2025.

[2] Mustak, Mekhail, Joni Salminen, Matti Mäntymäki, Arafat Rahman, and Yogesh K. Dwivedi. “Deepfakes: Deceptions, mitigations, and Opportunities.” *Journal of Business Research*, 154, (2023): https://doi.org/10.1016/j.jbusres.2022.113368

[3] “Unusual CEO Fraud via Deepfake Audio Steals US$243,000 From UK Company.” *Trend Micro*, September 05, 2022. https://www.trendmicro.com/vinfo/mx/security/news/cyber-attacks/unusual-ceo-fraud-via-deepfake-audio-steals-us-243-000-from-u-k-company.

[4] Stupp, Catherine. “Fraudsters Used AI to Mimic CEO’s Voice in Unusual Cybercrime Case.” *The Wall Street Journal*. Last modified August 30, 2019. https://www.wsj.com/articles/fraudsters-use-ai-to-mimic-ceos-voice-in-unusual-cybercrime-case-11567157402.

Section 19: AI ethics

[1] *Microsoft*. Microsoft Responsible AI Standard. June, 2022.

[2] “Ethics Guidelines for Trustworthy AI: High-Level Expert Group on Artificial Intelligence.” European Commission, April 08, 2019.

[3] UNESCO. "Recommendation on the Ethics of Artificial Intelligence." *UNESCO*, November 23, 2021.

[4] Tuck, David. "'Move Fast And Break Things' Doesn't Apply To AI." *Forbes*, December 8, 2023, www.forbes.com/councils/forbestechcouncil/2023/12/08/move-fast-and-break-things-doesnt-apply-to-ai/

[5] OpenAI Reports. *Release Strategies and the Social Impacts of Language Models.* https://arxiv.org/abs/1908.09203.

[6] Ferrara, Emilio. *Fairness and Bias in Artificial Intelligence: A Brief Survey of Sources, Impacts, and Mitigation Strategies.* University of California. https://arxiv.org/pdf/2304.07683.

[7] Tulsiani, Dr. Ravinder. "Strategies To Mitigate Bias In AI Algorithms." *eLearning Industry,* July 12, 2024. https://elearningindustry.com/strategies-to-mitigate-bias-in-ai-algorithms.

[8] Google Cloud. "What is Human-in-the-Loop (HITL) in AI & ML?" Google. https://cloud.google.com/discover/human-in-the-loop?hl=en.

[9] Dastin, Jeffrey. "Insight - Amazon Scraps Secret AI Recruiting Tool That Showed Bias Against Women." *Reuters*, October 11, 2018. https://www.reuters.com/article/world/insight-amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK0AG/.

[10] Novelli, Claudio, Mariarosaria Taddeo, and Luciano Floridi. "Accountability in Artificial Intelligence: What It Is and How It Works." *AI & Society* 39, no. 4 (2024): 1871–1882. https://doi.org/10.1007/s00146-023-01635-y.

[11] European Commission. *White Paper on Artificial Intelligence: A European Approach to Excellence and Trust.* Brussels: European Commission, February 2020. https://commission.europa.eu/publications/white-paper-artificial-intelligence-european-approach-excellence-and-trust_en.

[12] Pillai, Vinayak. "Enhancing Transparency and Understanding in AI Decision-Making Processes." *ICONIC Research and Engineering Journals* 8, no. 1 (July 2024): 168-176. https://www.irejournals.com/formatedpaper/1706039.pdf.

[13] Kumar, Sowmya Ramesh. "Effective Communication of Data Science Results to Non-Technical Stakeholders." *International Journal of Science and Research* 12, no. 8 (2023): 2497–2504. https://www.ijsr.net/archive/v12i8/SR24302211252.pdf.

Section 20: AI governance and strategy

[1] "Smart Cyber How AI Can Help Manage Cyber Risk." Deloitte, 2019. https://www2.deloitte.com/content/dam/Deloitte/global/Documents/Risk/gx-ra-smart-cyber.pdf.

[2] Boncella, Robert. "AI and Management: Navigating the Alignment Problem for Ethical and Effective Decision-Making." *Issues In Information Systems*, 2024. https://doi.org/10.48009/4_iis_2024_116.

[3] "White Paper on AI Governance Leadership Insights and the Voluntary AI Safety Standard in Practice." Australian Government, n.d.

https://www.governanceinstitute.com.au/app/uploads/2024/09/GovInst-AI-Whitepaper.pdf.

[4] Hall, Aaron. "7 Best Strategies for Stakeholder Feedback Integration." Attorney Aaron Hall, September 29, 2024. https://aaronhall.com/7-best-strategies-for-stakeholder-feedback-integration/?utm_source=chatgpt.com.

[5] Walter, Tim B, Martin Bremer, Lars Kaminski, and Volker Stich. "A Strategic AI Procedure Model For Implementing Artificial Intelligence." CONFERENCE ON PRODUCTION SYSTEMS AND LOGISTICS CPSL 2021, 2021.

[6] Marr, Bernard. "How Walmart Is Using Machine Learning AI, IOT and Big Data to Boost Retail Performance." *Forbes*, August 29, 2019. https://www.forbes.com/sites/bernardmarr/2017/08/29/how-walmart-is-using-machine-learning-ai-iot-and-big-data-to-boost-retail-performance/.

[7] BALBAA, Muhammad Eid, and Marina Sagatovna ABDURASHIDOVA. "The Impact of Artificial Intelligence in Decision Making: A Comprehensive Review." *EPRA International Journal of Economics, Business and Management Studies*, February 9, 2024, 27–38. https://doi.org/10.36713/epra15747.

[8] Westland, Jason, and R. Max Wideman. *The Project Management Life Cycle.* Vancouver: AEW Services, 2007.

[9] Ferrara, Emilio. "Fairness and Bias in Artificial Intelligence: A Brief Survey of Sources, Impacts, and Mitigation Strategies." *Sci* 6, no. 1 (December 26, 2023). https://doi.org/10.3390/sci6010003.

[10] Sarmiento, Adelina. "Measuring the Impact of AI Projects." Keepler, August 28, 2024. https://keepler.io/2024/08/28/measuring-the-impact-of-ai-projects/.

[11] Musani, Parvez. "Decking the Aisles with Data: How Walmart's AI-Powered Inventory System Brightens the Holidays." Walmart Global Tech, October 25, 2023. https://tech.walmart.com/content/walmart-global-tech/en_us/blog/post/walmarts-ai-powered-inventory-system-brightens-the-holidays.html.

[12] Géron, Aurélien. *Hands-on Machine Learning with Scikit-Learn, Keras & TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems.* O'Rielly Media, 2019.

[13] Sharma, Swati. "The AI Feedback Loop: From Insights to Action in Real-Time." *Zonka Feedback*, January 29, 2025. https://www.zonkafeedback.com/blog/ai-feedback-loop.

[14] Bourgais, Aurélien, and Issam Ibnouhsein. "Ethics-by-Design: The next Frontier of Industrialization." *AI and Ethics* 2, no. 2 (April 29, 2021). https://doi.org/10.1007/s43681-021-00057-0.

[15] Hernández, Ernesto Giralt. "Towards an Ethical and Inclusive Implementation of Artificial Intelligence in Organizations: A Multidimensional Framework." *Cornell University,* May 2, 2024. https://doi.org/10.48550/arXiv.2405.01697.

IV. List of acronyms

Adam: Adaptive Moment Estimation
AGI: Artificial General Intelligence
AI: Artificial Intelligence
AMF: Annual Maintenance Fee
ANAB: ANSI National Accreditation Board
ANNs: Artificial Neural Networks
API: Application Programming Interface
ARIMA: Autoregressive Integrated Moving Average
AutoML: Automated Machine Learning
BART: Bidirectional and Auto-Regressive Transformer
BEC: Business Email Compromise
BERT: Bidirectional Encoder Representations from Transformers
BiLSTM: Bidirectional Long Short-Term Memory
BoW: Bag of Words
CAIP: Certified Artificial Intelligence Professional
CBOW: Continuous Bag of Words
CCPA: California Consumer Privacy Act
CI/CD: Continuous Integration and Continuous Delivery/Deployment
CMMC: Cybersecurity Maturity Model Certification
CNN: Convolutional Neural Networks
CoT: Chain-of-Thought
CPD: Continuing Professional Development
CPU: Central Processing Unit
CRF: Conditional Random Fields
DBSCAN: Density-Based Spatial Clustering of Applications with Noise
DP: Deep Learning
DSRPAI: Dartmouth Summer Research Project on Artificial Intelligence
EDA: Exploratory Data Analysis
FDA: Food and Drug Administration
FN: False Negative
FNN: Feed-Forward Neural Network
FP: False Positive
FTC: Federal Trade Comission
GAN: Generative Adversarial Network
GDPR: General Data Protection Regulation
GloVe: Global Vectors for Word Representation
GPT: Generative Pre-trained Transformer
GPT-2: Generative Pre-trained Transformer 2
GPT-3: Generative Pre-trained Transformer 3
GPT-4: Generative Pre-trained Transformer 4
GPU: Graphics Processing Unit
GRU: Gated Recurrent Units

HCI: Human-Computer Interaction
HMM: Hidden Markov Models
HOG: Histogram of Oriented Gradients
HRI: Human-robot interaction
HTML: Hyper Text Markup Language
IBM: International Business Machines Corporation
IEC: International Electrotechnical Commission
IEEE: Institute of Electrical and Electronics Engineers
IMUs: Inertial Measurement Unit
IQ: Intelligence Quotient
IQR: Interquartile Range
ISO: International Organization for Standardization
KNN: K-Nearest Neighbors
KPI: Key Performance Indicators
LiDAR: Light Detection and Ranging
LightGBM: Light Gradient-Boosting Machine
LIME: Local Interpretable Model-agnostic Explanations
LLM: Large Language Models
LOOCV: Leave-One-Out Cross-Validation
LSTM: Long Short-Term Memory
MAE: Mean Absolute Error
MIT: Massachusetts Institute of Technology
ML: Machine Learning
MLE: Maximum Likelihood Estimation
MLP: Multilayer Perceptron
MRI: Magnetic Resonance Imaging
mRNA: Messenger Ribonucleic Acid
MSE: Mean Squared Error
NASA: National Aeronautics and Space Administration
NER: Named Entity Recognition
NIST: National Institute of Standards and Technology
NLP: Natural Language Processing
NLTK: Natural Language Toolkit
NMT: Neural Machine Translation
OECD: Organisation for Economic Co-operation and Development
OWASP: Open Worldwide Application Security Project
PCA: Principal Component Analysis
PECB: Professional Evaluation and Certification Board
POS: Part-of-Speech
QA: Question Answering
RAG: Retrieval Augmented Generation
RBM: Restricted Boltzmann Machines
R-CNN: Region-based Convolutional Neural Networks
ReLU: Rectified Linear Unit

ResNet: Residual Network
RFE: Recursive Feature Elimination
RGB: Red, Green, and Blue
RL: Reinforcement Learning
RLHF: Reinforcement Learning from Human Feedback
RMSE: Root Mean Square Error
RMSProp: Root Mean Square Propagation
RNN: Recurrent Neural Networks
ROI: Return on Investment
ROS: Robot Operating Systems
ROUGE: Recall-Oriented Understudy for Gisting Evaluation
SAS: Statistical Analysis System
SBOM: Secure Software Bill of Materials
SGD: Stochastic Gradient Descent
SHAP: SHapley Additive exPlanations
SIFT: Stop and Reflect, Investigate the Source, Find Reliable Coverage, and Trace the Claim
SLAM: Simultaneous Localization and Mapping
SMOTE: Synthetic Minority Over-sampling Technique
SMT: Statistical Machine Translation
SPSS: Statistical Package for Social Science
SQL: Structured Query Language
SSD: Single Shot Multibox Detector
SSL: Self-Supervised Learning
SURF: Speeded Up Robust Features
TF-IDF: Term Frequency-Inverse Document Frequency
TN: True Negative
T-NLG: Turing Natural Language Generation
TP: True Positive
TPU: Tensor Processing Unit
t-SNE: T-Distributed Stochastic Neighbor Embedding
UMAP: Uniform Manifold Approximation and Projection
UNESCO: United Nations Educational, Scientific and Cultural Organization
USD: United States Dollar
VAE: Variational Autoencoders
VGGNet: Visual Geometry Group Network
ViTs: Vision Transformers
VQA: Visual Question Answering
XAI: Explainable Artificial Intelligence
XGBoost: Extreme Gradient Boosting
XOR: Exclusive Or
YOLO: You Only Look Once