



LIVING GLOSSARY FOR AI AND MEDICINE

FOREWORD

Artificial Intelligence (AI) has had an incredible impact on health and healthcare around the world. Researchers and clinicians are using it to enhance hospital operations and improve patient outcomes. One of the challenges of new innovations is that they give rise to new words and terminology. Since AI draws from fields such as computer science, health, and data science, frequently used terminology can be especially confusing to learners from different disciplines.

This glossary is the product of contributions and deep discussions by an international consortium of AI in medicine experts from some of the world's leading universities and research centres. It is our collective attempt to capture and define current terms for a global audience and include context and examples in the definitions that are relevant to healthcare professionals. Because of the ever-changing nature of AI, this PDF is the first stage of an evolving project to create an online multimodal living glossary that will be regularly updated as AI in medicine advances in the future.

The International AI in Medicine Education Working Group hopes that this glossary inspires and educates AI health learners around the world. Special thanks also go to the hard-working staff of the Temerty Centre for AI Research and Education in Medicine (T-CAIREM) at the University of Toronto for organizing this ambitious project for the global community.



Laura C. Rosella, PhD

Education Lead, Temerty Centre for Artificial Intelligence Research and Education in Medicine (T-CAIREM), University of Toronto

Thank you to all the members of the working group for their contributions



TABLE OF CONTENTS

PAGE 5

Active Learning
Adaptive Boosting (AdaBoost)
Agentic AI
Automation Bias
Algorithmic Bias
Alluvial Plot
Anomaly Detection (Outlier Detection or Novelty Detection)
Artificial Intelligence (AI)

PAGE 6

Artificial Neural Network (ANN)
Attention Models
Augmented Intelligence
Bagging
Bag of Words (BoW)
Base-Rate Neglect Bias
Batch Learning (Offline Learning)

PAGE 7

Bayesian Networks
Bias-Variance Tradeoff
Big Data
Blockchain
Box Plot
Causal Discovery
Chatbot
Clinical Decision Support Systems (CDSS)

PAGE 8

Cloud Computing
Computational Phenotyping
Computer Vision
Confirmation Bias
Confounding Bias
Convolutional Neural Network (CNN)
Cross-Validation
Curriculum Learning

PAGE 9

Cybersecurity
Data Augmentation
Dataset Drift
Dataset Shift
Data Governance
Data Wrangling
Data Leakage

PAGE 10

Deployment
Digital Therapeutics
Decision Trees
Deep Learning
Demographic Bias
Deep Reinforcement Learning
Digital Biomarkers
Edge AI

PAGE 11

Elastic Net Regression
Electronic Health Record (EHR)
Embedding
Ensemble methods
Ethics
Explainability
Explainable AI (XAI)
False Negative
False Positive

PAGE 12

F1 Score
Fairness in AI
Features
Feature Engineering
Feature Importance
Federated Learning
Few-Shot Learning
Foundation Models

PAGE 13

Framing Bias
Gender Bias
Generative Adversarial Network (GAN)
Generative AI
Generalizability
Generalized Linear Models (GLMs)
Gradient Boosting

PAGE 14

Gradient Descent
Graph Neural Networks
Hallucination
Health Chatbots
Historical Data Bias
Human-in-the-Loop (HITL)
Hyperparameter Tuning
Image Analysis

PAGE 15

Imbalanced Data
Imputation
Internet of Things (IoT)
Interpretability
Knowledge Graphs
K-Nearest Neighbors (KNN)
Labelling
Language Bias

PAGE 16

Large Language Models (LLMs)
LASSO (Least Absolute Shrinkage and Selection Operator)
Light Model
One-Hot Encoding
Online Learning
Ordinary Least Squares (OLS) Regression
Machine Learning

PAGE 17

Measurement Bias
Multimodal Learning
Multi-Omics Data Integration
Natural Language Processing (NLP)
Naive Bayes
Negative Predictive Value (NPV)
Neural Networks
Node (Neural Network Node)

PAGE 18

Omitted Variable Bias
Overfitting
Patient Digital Twin
Pipeline
Precision
Precision Medicine
Predictive Analytics
Prescriptive Analytics

PAGE 19

Preprocessing
Primacy Effect Bias
Privacy
Quantum Computing in Medicine
Random Forest
Recall
Recurrent Neural Networks (RNNs)
Reinforcement Learning (RL)

PAGE 20

Remote Patient Monitoring (RPM)
Regulatory AI Compliance
Robotics
Sampling Bias
Self-attention
Self-supervised Learning
Semi-supervised Learning
Sensitivity

PAGE 21

Sentiment Analysis
Shapley Additive exPlanations (SHAP)
Small Language Models (SLMs)
Specificity
Supervised Learning
Swarm Intelligence
Synthetic Data
Telemedicine

PAGE 22

Temporal Bias
Training
Training Dataset
Transferability
Transfer Learning
Transformer Models
Transparency
Trust

PAGE 23

Unsupervised Learning
Validation
Validation Dataset
Variance
Violin Plot
Word Cloud
Zero-Shot Learning

WHO WE ARE

As a consortium of representatives from universities and health centres worldwide, the AI in Medicine International Education Working Group is at the forefront of artificial intelligence (AI) education in medicine. Our group provides a forum for relevant discussion and collaboration and is dedicated to producing AI in medicine education material that the global community can use.

WHY WE ARE WRITING THIS

Our glossary is more than just a collection of definitions. It's a powerful tool designed to enhance communication, foster learning, and promote collaboration, especially in the interdisciplinary context of AI and medicine. As a living glossary, it's a dynamic and evolving resource that can be updated and revised to reflect the collective voices and changes in the field. It's a centralized reference point that adapts to the evolution of AI and medicine, empowering you with the latest knowledge. **In the future, our goal is to create a multimodal living glossary that will be regularly updated with new terms as the field of AI health evolves.**

HOW TO USE THIS GLOSSARY

This glossary is designed as a practical and accessible resource for anyone engaging with AI in medicine, from clinicians and researchers to policymakers and students. Each term includes a clear definition and a relevant application to medicine, helping to bridge the gap between technical AI concepts and their real-world impact in healthcare. Use this glossary to improve interdisciplinary communication, enhance your understanding of AI's role in medicine, and ensure clarity in discussions, research, and education. As a living resource, we encourage you to contribute to the evolution of this glossary by suggesting new terms or refinements using [this form](#).



HOW TO CITE THIS GLOSSARY

APA Style

International AI in Medicine Education Working Group. (Year). *AI and Medicine Glossary*. [URL]

If citing a specific term from the glossary, please include the term name (e.g., reinforcement learning) and retrieval date (if applicable). Example:

International AI in Medicine Education Working Group. "Reinforcement Learning." *Living Glossary for AI and Medicine*, Year, [URL]. Accessed [Date].

Active Learning

An iterative machine learning approach in which the algorithm selectively identifies the most informative unlabeled data points for annotation by human experts. This strategy improves model performance with fewer labeled examples, reducing annotation costs. In healthcare, active learning is particularly useful when expert labeling is expensive or limited, such as in medical imaging or rare disease classification.

Adaptive Boosting (AdaBoost)

An ensemble learning algorithm that improves classification accuracy by fitting a sequence of weak learners on repeatedly modified versions of the data to create a strong classifier. AdaBoost assigns higher weights to misclassified instances in each iteration, focusing the model's learning on difficult cases to improve overall accuracy. In medicine, AdaBoost is used in disease classification, patient risk prediction, and medical image analysis.

Agentic AI

A form of artificial intelligence capable of autonomous, goal-directed decision-making with minimal human supervision. In medicine, agentic AI can dynamically respond to patient data in real time, enabling applications such as personalized treatment recommendations, robotic-assisted surgery, and adaptive clinical decision support. Although agentic systems act independently, they are designed to align with clinical standards, ethical principles, and regulatory requirements.

Automation Bias

The propensity for humans to overly trust and rely on suggestions from automated decision-making systems, potentially overlooking or dismissing contradictory information obtained without automation, even when such information is correct. In healthcare, this can occur when clinicians overly depend on AI-generated diagnoses or recommendations without sufficient critical assessment.

Algorithmic Bias

Systematic errors in AI outputs that arise from biased training data, flawed assumptions, or inadequate model design. In healthcare, algorithmic bias can lead to unequal treatment or misdiagnosis, disproportionately impacting certain patient populations. Addressing algorithmic bias is critical to ensuring fairness, safety, and equity in clinical AI applications.

Alluvial Plot

A type of flow diagram used to visualize how categories or groupings evolve over time or across different conditions. It displays relationships between categorical variables using ribbon-like flows that connect nodes or strata. In healthcare, alluvial plots are useful for tracking patient transitions between disease states, treatment pathways, or care settings over time.

Anomaly Detection (Outlier Detection or Novelty Detection)

A machine learning technique used to identify data points or patterns that deviate significantly from the norm. In healthcare, anomaly detection can flag unusual clinical events, rare disease presentations, equipment malfunctions, or data entry errors by modeling typical patterns and identifying deviations. It supports early warning systems, quality assurance, and fraud detection in clinical and operational settings.

Artificial Intelligence (AI)

The development of computer systems that can perform tasks typically requiring human intelligence, such as learning, reasoning, problem-solving, and decision-making. In medicine, AI may enhance diagnostics, treatment planning, and healthcare operations by analyzing complex data, recognizing patterns, and supporting clinical decision-making.

Artificial Neural Network (ANN)

A type of machine learning algorithm that is inspired by the human brain. ANNs can learn complex patterns from data, and they are used in various applications, such as image recognition, speech recognition, and natural language processing. ANNs are used in medicine for disease diagnosis, medical imaging analysis, drug discovery, and personalized treatment by detecting patterns in complex healthcare data.

Attention Models

A class of deep learning mechanisms that enable neural networks to focus on the most relevant parts of input data when making predictions. By dynamically weighting different inputs, attention models improve performance in tasks requiring contextual understanding, such as natural language processing and image recognition. In medicine, attention models enhance AI applications in medical imaging, clinical text analysis, and decision support. For example, they improve diagnostic accuracy in radiology by highlighting key regions in scans, assist in summarizing electronic health records by prioritizing critical information, and refine disease prediction models by focusing on the most important patient data.

Augmented Intelligence

A term used to emphasize AI's role in enhancing, rather than replacing, human intelligence in (medical) decision-making. Some argue for using "Augmented Intelligence" terminology over "Artificial Intelligence," for this reason.

Bagging

A method to reduce overfitting and improve generalizability by using subsets of data to train multiple models and average predictions. This method, applied in medical AI, can improve diagnostic accuracy and robustness in prediction tasks like cancer detection and risk estimation.

Bag of Words (BoW)

A simple text representation method in natural language processing (NLP) where a document is converted into a set of words, disregarding grammar and word order while maintaining word frequency. BoW is commonly used for text classification, sentiment analysis, and information retrieval. In medicine, BoW is applied in analyzing clinical notes, medical literature, and patient feedback. It helps in tasks such as disease classification based on electronic health records, identifying key topics in research articles, and detecting trends in patient-reported symptoms or experiences.

Base-Rate Neglect Bias

A cognitive bias where the underlying prevalence (base rate) of a condition is overlooked when interpreting probabilities or making decisions. In medical AI, this bias can lead to overestimating the likelihood of rare diseases or underestimating common ones, potentially resulting in inappropriate diagnostic or treatment decisions.

Batch Learning (Offline Learning)

A machine learning paradigm in which models are trained on a complete, pre-collected dataset. Unlike online learning, batch learning does not adapt continuously; instead, the model is trained once (or at scheduled intervals) and then deployed for use. In healthcare, batch learning is used in scenarios where models are periodically updated using large datasets, such as hospital-wide EHR records or imaging archives.

Bayesian Networks

Probabilistic graphical models that use directed acyclic graphs to represent variables and their conditional dependencies. They are used to model uncertainty and reason about probabilistic relationships. In medicine, Bayesian networks support clinical decision-making by estimating the likelihood of a disease based on symptoms, risk factors, and patient history, helping to capture complex interdependencies in diagnostic reasoning.

Bias-Variance Tradeoff

A fundamental concept in machine learning that describes the balance between bias (error due to overly simplistic models that underfit the data) and variance (error due to overly complex models that overfit to training data). The goal is to find an optimal model complexity that minimizes total prediction error. In medicine, the bias-variance tradeoff is critical when developing AI models for diagnostics or risk prediction. A high-bias model may overlook important patient features, leading to inaccurate diagnoses, while a high-variance model may perform well on training data but fail in real-world clinical settings.

Big Data

Extremely large and complex datasets generated from various digital sources, including electronic health records, medical imaging, wearable devices, and genomics. In medicine, big data enables AI models to identify patterns, make predictions, and support clinical decision-making, ultimately improving diagnostics, treatment, and healthcare management.

Blockchain

A secure, distributed database that can be used to store and share health information while protecting patient privacy. Used for secure patient data sharing and decentralized clinical trials.

Box Plot

A graphical visualization of a dataset's distribution that displays the median, interquartile range (IQR), and potential outliers. In healthcare research, box plots are commonly used to compare variables such as biomarker levels or treatment outcomes across patient groups, offering a concise summary of variability and central tendency.

Causal Discovery

The use of AI and statistical methods to identify causal relationships between variables in complex datasets. In medicine, causal discovery helps figure out how different things, like genes, lifestyle, and exposures to the environment, affect the risk of disease and how well a patient does, so that doctors can make more targeted treatments.

Chatbot

An AI-powered conversational system designed to interact with users through text or speech, providing automated responses and assistance based on predefined rules or machine learning models. Chatbots can range from simple scripted systems to advanced models using natural language processing (NLP) and deep learning. In medicine, chatbots are used for patient engagement, symptom checking, mental health support, and administrative tasks. They can assist with answering common health-related questions and supporting telemedicine services.

Clinical Decision Support Systems (CDSS)

Software tools that assist healthcare professionals by delivering patient-specific information and evidence-based recommendations to support clinical decision-making. In medicine, CDSS can flag potential drug interactions, suggest diagnostic pathways, or prioritize high-risk patients.

Cloud Computing

A model for delivering computing resources, such as storage, processing power, and software, on demand over the internet. In healthcare, cloud computing enables the development and deployment of AI applications, including machine learning, natural language processing, and large-scale data analytics. It supports scalable data infrastructure, facilitates real-time collaboration, and enhances access to computational tools for clinical and research purposes.

Computational Phenotyping

The use of machine learning and artificial intelligence techniques to automatically identify and characterize clinically meaningful patterns or traits (phenotypes) from large-scale health data, such as electronic health records (EHRs), genomic data, and medical imaging. This approach enables the discovery of novel disease subtypes, patient stratification, and more precise prediction of health outcomes, supporting personalized medicine and population health research.

Computer Vision

AI techniques that enable computers to interpret and analyze visual data. This involves using technology like sensors and convolutional neural networks (CNN). This is used in medical imaging for automated tumor detection, radiology analysis, and pathology slide interpretation.

Confirmation Bias

The cognitive tendency to favor information that confirms one's existing beliefs or expectations. In AI and medicine, confirmation bias can influence how clinicians interpret model outputs or how data is selected and labeled for training, potentially reinforcing pre-existing assumptions and limiting model objectivity.

Confounding Bias

Arises when an unmeasured or inadequately measured variable influences both the independent and dependent variables, leading to spurious associations between an exposure (e.g., treatment) and an outcome. Failure to account for confounders in the context of a causal analysis that uses AI or machine learning models can result in misleading conclusions about causality or effectiveness.

Convolutional Neural Network (CNN)

A specialized deep learning algorithm designed for processing and analyzing visual data. CNNs excel at recognizing patterns in images by using layers of convolutional filters to detect features such as edges, textures, and shapes. In medicine, CNNs are widely used in medical imaging applications, including radiology, pathology, and dermatology, where they assist in detecting abnormalities, segmenting structures, and supporting diagnostic decision-making.

Cross-Validation

A model evaluation technique that assesses how well a machine learning algorithm generalizes to unseen data. It involves partitioning the dataset into training and validation sets multiple times. Common approaches include the hold-out method (e.g., 70% training, 30% validation) and k-fold cross-validation, where the dataset is divided into k subsets, and the model is trained and tested k times. In healthcare, cross-validation is used to build robust models from limited data.

Curriculum Learning

Curriculum learning is a machine learning strategy where models are trained progressively, starting from simpler tasks or examples and gradually increasing complexity. Inspired by the way humans learn, this approach helps the model first grasp foundational concepts before tackling more difficult ones. Curriculum learning can lead to faster convergence, improved generalization, and more stable training outcomes.

Cybersecurity

The practice of safeguarding computer systems, networks, and data from unauthorized access, breaches, or malicious attacks. In healthcare, cybersecurity is critical for protecting sensitive patient data and ensuring the integrity of datasets used to train AI models. Strong cybersecurity practices help maintain trust in digital health systems and prevent compromised data from degrading model performance or violating privacy laws.

Data Augmentation

Techniques used to artificially expand training datasets by creating modified versions of existing data, improving AI model performance. Data augmentation is often used to increase the size and diversity of data to reduce overfitting and improve generalization.

Dataset Drift

A phenomenon where the statistical properties of a dataset change over time, potentially leading to decreased model performance. In medicine, dataset drift can affect AI models used for diagnostics or predictive analytics, as shifts in patient demographics, clinical practices, or disease prevalence may cause the model to make less accurate predictions over time. Monitoring and updating models regularly is essential to maintain reliability.

Dataset Shift

A shift in the statistical properties of a dataset over time, which can lead to reduced model accuracy and reliability. In healthcare, dataset drift may result from evolving patient demographics, changing clinical guidelines, or new diagnostic technologies. Continuous monitoring and model retraining are necessary to ensure sustained performance of AI systems in clinical practice.

Data Governance

The policies and processes that ensure the ethical and responsible use of healthcare data, including consent, privacy, and security issues. In medicine, strong data governance frameworks are essential for ensuring patient confidentiality, enabling secure data sharing for research and clinical decision-making, and maintaining compliance with regulatory standards. Effective data governance supports trustworthy AI applications in healthcare by ensuring data integrity.

Data Wrangling

The process of cleaning, transforming, and structuring raw data to prepare it for analysis. This includes handling missing values, correcting inconsistencies, standardizing formats, and merging data from multiple sources. In healthcare, data wrangling is essential for preparing datasets such as electronic health records (EHRs), medical imaging, and patient-reported outcomes, ensuring data quality and interoperability for AI modeling, clinical decision-making, and research.

Data Leakage

A modeling error where information from outside the training data, or from future or outcome-related variables, improperly influences the learning process, resulting in overly optimistic performance. In medicine, data leakage can undermine model reliability. For example, including post-treatment variables when predicting patient outcomes may falsely inflate accuracy but render the model useless in real-time decision-making.

Deployment

The process of integrating and implementing an AI model or system into a real-world environment where it can be used for decision-making, automation, or analysis. Deployment involves transitioning a model from development to operational use, ensuring it performs reliably in practice. In medicine, AI deployment includes integrating predictive models into electronic health records (EHRs) for clinical decision support, deploying AI-powered diagnostic tools in radiology, or using machine learning algorithms to optimize hospital operations. Successful deployment requires rigorous validation, bias or drift monitoring, and compliance with regulatory and ethical standards.

Digital Therapeutics

Evidence-based software interventions designed to prevent, manage, or treat medical conditions through digital platforms. Examples include mobile applications delivering cognitive behavioral therapy (CBT) for depression or anxiety, and digital programs that support diabetes self-management. Digital therapeutics are often regulated as medical devices and are integrated into clinical care pathways to improve health outcomes.

Decision Trees

A machine learning algorithm that uses a tree-like model of decisions to classify data or make predictions. Each node represents a feature, each branch represents a decision based on that feature, and each leaf represents an outcome. Decision trees are interpretable and can handle both classification and regression tasks. In medicine, decision trees are widely used for clinical decision support, risk stratification, and diagnostic algorithms. Their transparency makes them popular in healthcare applications.

Deep Learning

A technique where AI agents learn optimal actions through interactions with their environment, using rewards and penalties to guide behavior. In medicine, deep reinforcement learning is applied to personalized treatment planning, robotic surgery, and real-time decision-making systems, enabling adaptive strategies based on patient-specific feedback.

Demographic Bias

Arises when there are systematic discrepancies in how the model performs across sociodemographic, clinical, and/or geographic factors. One way in which this can arise is when the data inadequately represents diverse demographic groups. This can lead to outcomes that are not generalizable across all population segments and can compromise equity in care.

Deep Reinforcement Learning

A machine learning technique where an AI learns optimal actions through trial and error. In medicine, this can be used in applications such as robotic surgery and adaptive treatment planning, where AI learns the best strategies for precision care.

Digital Biomarkers

Health-related data points collected via digital devices (e.g., wearables, smartphones) that can be analyzed using AI to provide insights into patient health.

Edge AI

The deployment of artificial intelligence algorithms directly on local hardware devices—such as smartphones, medical wearables, or bedside monitors—rather than relying on centralized cloud infrastructure. Edge AI enables real-time processing and decision-making, which is especially valuable in time-sensitive healthcare applications like remote patient monitoring, emergency response, or mobile diagnostics, while also enhancing data privacy and reducing latency.

Elastic Net Regression

A regularized regression method that combines LASSO (L1) and Ridge (L2) penalties to enhance prediction accuracy and manage multicollinearity. Elastic Net is especially useful with high-dimensional or highly correlated data. In medicine, it supports tasks such as disease risk prediction, biomarker selection, and analysis of complex datasets like genomics or electronic health records (EHRs), balancing feature selection with model stability.

Electronic Health Record (EHR)

A digital record of a patient's health history, including medical diagnoses, treatments, and medications.

Embedding

A machine learning technique that transforms high-dimensional data—such as words, images, or categorical variables—into lower-dimensional continuous vectors that preserve meaningful relationships and structure. Embeddings capture semantic or contextual similarities, making them essential for processing complex data types. In medicine, embeddings are used to analyze electronic health records, represent medical concepts in natural language processing, and enhance diagnostic models by enabling more efficient and nuanced data interpretation.

Ensemble methods

A machine learning approach that combines predictions from multiple models to improve overall performance and robustness. By aggregating outputs from diverse models—such as decision trees, neural networks, or support vector machines—ensemble methods reduce bias, variance, and overfitting. In healthcare, they are used to enhance diagnostic accuracy, risk prediction, and outcome forecasting.

Ethics

The principles that guide our behavior. AI ethics is a field that is concerned with the ethical implications of AI. AI ethics scholars are working to develop guidelines for the responsible development and use of AI.

Explainability

The extent to which the inner workings and outputs of an AI model can be understood and interpreted by humans. In healthcare, explainability is critical for building trust in AI-driven diagnoses and treatment recommendations. Techniques such as SHAP values or attention maps help clinicians understand why a model made a particular prediction, enabling more informed and accountable decisions.

Explainable AI (XAI)

The development of AI models and algorithms that can be easily understood and interpreted by humans is particularly important in the context of healthcare decision-making. Critical for patient safety and regulatory approval in medicine.

False Negative

A classification error where a model incorrectly predicts a positive outcome for a case that is actually negative. In medicine, false positives can lead to unnecessary diagnostic tests, treatments, and patient anxiety, especially in screening scenarios such as cancer detection or infection surveillance.

False Positive

An error in a classification where a model incorrectly predicts a positive outcome for a case that is actually negative. In other words, the model identifies a condition or event that is not actually present. In medicine, false positives can lead to unnecessary tests, treatments, and patient anxiety.

F1 Score

A performance metric for evaluating the accuracy of an AI model, particularly in healthcare applications where both precision (positive predictive value) and recall (sensitivity) are critical. The harmonic mean of precision and recall provides a balanced measure that accounts for both false positives and false negatives.

Fairness in AI

The principle of ensuring that artificial intelligence systems make unbiased, equitable, and just decisions across different populations. Fair AI minimizes disparities in model performance and outcomes, particularly for historically marginalized or underrepresented groups. In medicine, fairness in AI is critical for developing diagnostic tools, treatment recommendations, and predictive models that work accurately across diverse patient populations. Bias in healthcare AI can lead to disparities in disease detection, access to care, and treatment effectiveness. Ensuring fairness requires diverse training data, bias detection techniques, and continuous evaluation to promote equitable healthcare outcomes..

Features

A measure of how much each input variable contributes to the predictive performance of a machine learning model. Feature importance helps identify which variables most influence the model's output. In healthcare, understanding feature importance supports interpretability and clinical relevance, and it can be calculated using techniques such as SHAP values, permutation importance, or model-specific weights.

Feature Engineering

The process of transforming raw data into meaningful inputs (features) that improve a machine learning model's performance. Effective feature engineering is critical in medicine, where domain knowledge is used to derive clinically relevant variables from complex datasets, enhancing model accuracy and interpretability.

Feature Importance

The measure of how much each input variable (feature) contributes to the predictive power of a machine learning model. It helps identify which variables (features) have the greatest impact on the model's predictions. There are various methods for calculating feature importance.

Federated Learning

A machine learning approach that allows models to be trained across decentralized datasets located at different institutions or devices, without transferring the data itself. In healthcare, federated learning enhances patient privacy by keeping sensitive data (e.g., EHRs or imaging) local while still enabling collaborative model development. This supports AI innovation across hospitals, research centers, and mobile health platforms without compromising confidentiality.

Few-Shot Learning

A machine learning approach that enables models to make accurate predictions with only a small number of labeled examples. This is especially valuable in healthcare scenarios like rare disease diagnosis, where labeled data are scarce. Few-shot learning improves model generalization in low-data contexts by leveraging prior knowledge or pre-trained models.

Foundation Models

Large-scale AI models trained on vast amounts of diverse data that can be adapted to a wide range of tasks with minimal fine-tuning. In medicine, foundation models power applications such as medical image analysis, clinical text processing, drug discovery, and personalized healthcare by leveraging general knowledge and adapting to specific medical domains.

Framing Bias

A cognitive bias where the order in which information or problems are presented influences decision-making. In healthcare AI, the initial framing of data or clinical questions can skew how algorithms interpret and prioritize information, affecting diagnostic or treatment decisions. Misleading framing can result in skewed diagnostic or treatment recommendations, highlighting the importance of careful data presentation and model input design.

Gender Bias

The inclination to treat individuals differently or to generate skewed outcomes based solely on their gender. In healthcare AI, this can occur if models are trained on data that underrepresents or mischaracterizes one gender, leading to inequitable care or misdiagnosis.

Generative Adversarial Network (GAN)

A machine learning framework consisting of two neural networks—a generator and a discriminator—that are trained together in a competitive process. The generator creates synthetic data, while the discriminator evaluates its authenticity. This adversarial training continues until the generated data becomes indistinguishable from real data. In medicine, GANs are used for generating realistic medical images, augmenting rare data, and enhancing image resolution in radiology and pathology.

Generative AI

A type of artificial intelligence that creates new content, such as text, images, or synthetic data, based on learned patterns from training data. In medicine, generative AI is used for tasks such as generating and summarizing medical reports, simulating patient data for research, and assisting in drug discovery, particularly in fields like proteomics and genomics.

Generalizability

The extent to which an AI model maintains its performance when applied to new, unseen healthcare data beyond the original training set. High generalizability is essential for ensuring that models remain accurate and reliable across different patient populations, clinical settings, and institutions.

Generalized Linear Models (GLMs)

A broad class of regression models that extend linear regression by allowing the response variable to follow distributions beyond the normal (e.g., binomial, Poisson). GLMs use a link function to relate the expected value of the outcome to a linear combination of predictors. In medicine, GLMs are widely used to model outcomes such as disease risk (logistic regression), event counts (Poisson regression), and skewed healthcare costs (gamma regression).

Gradient Boosting

A machine learning technique that builds predictive models by sequentially combining multiple weak learners (typically decision trees), each correcting the errors of the previous one to improve overall accuracy. By minimizing errors iteratively through gradient descent, gradient boosting creates highly effective models for complex datasets. In medicine, gradient boosting is widely used in predictive analytics, such as identifying patients at risk of developing chronic diseases, optimizing treatment plans based on patient data, and improving diagnostic accuracy by integrating different clinical and demographic features.

Gradient Descent

An optimization algorithm used to minimize the error of a model by iteratively adjusting its parameters in the direction that reduces the loss function. By computing gradients, the algorithm updates model weights step by step, reducing the loss function. In medicine, gradient descent is fundamental to training AI models for tasks such as medical image analysis and disease prediction. It enables deep learning models, like neural networks, to learn from complex healthcare data and improve their accuracy over time.

Graph Neural Networks

A class of neural networks specifically designed to process data structured as graphs—collections of nodes and edges representing entities and their relationships. Unlike traditional neural networks that handle images or sequences, GNNs are well-suited for modeling complex relational data. In healthcare, they are used for drug interaction prediction, modeling biological networks, and analyzing relationships between clinical entities.

Hallucination

Also known as confabulation. Hallucination refers to the generation of false, misleading, or nonsensical information by an AI model, often presented as if it were factual. This occurs when the model extrapolates beyond its training data or misinterprets input patterns. In medicine, AI hallucinations can pose significant risks, such as incorrect diagnoses, inaccurate medical summaries, or fabricated clinical recommendations.

Health Chatbots

AI-powered chatbots that provide virtual assistance to patients, such as answering health-related questions and scheduling appointments. Certain chatbots might also assist clinicians, such as supporting charting and diagnosis/treatment planning.

Historical Data Bias

Systematic distortion in data that arises from societal, cultural, institutional, or individual biases present at the time the data was collected. This bias reflects the values, norms, and power structures of the past, and it can lead to unfair or inaccurate outcomes when historical data is used in modern decision-making systems, especially in machine learning or AI.

Human-in-the-Loop (HITL)

A framework in artificial intelligence where human expertise is integrated into the model development, training, or decision-making process to enhance accuracy, reliability, and ethical considerations. HITL systems leverage human oversight to refine AI outputs, correct errors, and improve model learning. In medicine, HITL ensures physicians validate AI-generated diagnoses or treatment recommendations. It also plays a key role in training AI models for medical imaging, ensuring that automated assessments align with expert interpretations.

Hyperparameter Tuning

The process of selecting the optimal values for hyperparameters, settings that govern how a machine learning algorithm learns (e.g., learning rate, tree depth, or regularization strength). Effective tuning significantly impacts model performance and is often achieved through techniques like grid search, random search, or Bayesian optimization.

Image Analysis

The use of computer vision, deep learning, and pattern recognition techniques to interpret and extract information from medical images. In healthcare, image analysis supports tasks such as detecting tumors in radiology scans, segmenting anatomical structures, and monitoring disease progression in imaging modalities like MRI, CT, and ultrasound.

Imbalanced Data

A situation where certain classes or outcomes (e.g., rare diseases) are underrepresented in a dataset, which can lead to biased predictions in AI models. In medicine, imbalanced data is common and requires specific strategies, such as resampling, class weighting, or specialized algorithms, to ensure that models do not overlook minority cases that may be clinically important.

Imputation

A data preprocessing technique where missing values in a dataset are estimated and filled in, so the data can be used for analysis or model training. In medicine, imputation helps fill in missing laboratory results or clinical data in electronic health records, ensuring AI models can make accurate predictions despite incomplete data.

Internet of Things (IoT)

A network of interconnected physical devices—such as sensors, wearables, and medical equipment—that collect, transmit, and exchange data over the internet. In healthcare, IoT enables real-time monitoring of patient health through devices like smartwatches, glucose monitors, and implantable sensors. This connectivity supports remote care, early detection of health issues, and data-driven decision-making in clinical and home settings.

Interpretability

The extent to which a human can understand and explain how an AI model arrives at a decision. In medicine, interpretability is critical for building trust, ensuring transparency, and enabling clinicians to validate AI-driven diagnoses and treatment recommendations. Tools such as SHAP values and feature importance scores are commonly used to support interpretability.

Knowledge Graphs

AI-driven data structures that connect medical concepts, diseases, treatments, and research to support medical decision-making. Used in AI-driven drug discovery, clinical decision-making, and medical literature analysis.

K-Nearest Neighbors (KNN)

A non-parametric machine learning algorithm used for classification and regression tasks. It predicts the class or value of a new data point based on the majority class (for classification) or average value (for regression) of its k closest neighbors in the feature space. In medicine, KNN is used for disease diagnosis by comparing patients with similar clinical features, predicting outcomes, and classifying medical images, predicting patient outcomes by comparing with similar cases, and medical image classification by analyzing patterns in labeled datasets.

Labelling

The process of annotating data with meaningful tags or categories to train AI models. In medicine, labelling is essential for supervised learning tasks, such as marking disease regions in medical images, categorizing symptoms in patient records, and structuring clinical notes for natural language processing models.

Language Bias

Errors or misinterpretations occur when language differences are not properly addressed in data collection, processing, or algorithm design. This bias can impact patient interactions, documentation, and the natural language processing components of healthcare AI, potentially leading to misunderstandings or misclassification.

Large Language Models (LLMs)

Advanced AI models trained on vast amounts of text to understand, generate, and analyze human language. In healthcare, LLMs are used for tasks such as summarizing medical records, generating clinical notes, assisting patient communication, and synthesizing scientific evidence.

LASSO (Least Absolute Shrinkage and Selection Operator)

A regression technique that performs both variable selection and regularization by applying an L1 penalty to shrink some coefficients to zero. This helps prevent overfitting and improves model interpretability by selecting only the most relevant features. In medicine, LASSO is widely used in predictive modeling and clinical research to identify key risk factors for diseases, optimize prognostic models, and refine feature selection in high-dimensional datasets, such as genomic data or electronic health records.

Light Model

A machine learning model that is designed to be small, fast, and computationally efficient while still maintaining good predictive performance. Light models are particularly useful when working with limited computational resources or deploying models on edge devices (e.g., mobile phones or medical devices)

One-Hot Encoding

A technique used to convert categorical variables into a binary format, where each category is represented by a unique binary vector. This transformation allows machine learning models to process categorical data effectively. In medicine, one-hot encoding is used to structure patient data for AI models, such as converting categorical variables

like disease types, medication names, or hospital departments into a numerical format. It enables predictive models to analyze patient demographics, treatment histories, and diagnostic categories while preserving the distinctiveness of each category.

Online Learning

A machine learning approach where the model updates continuously as it receives new data, typically one instance at a time. Unlike traditional batch learning, online learning is suited to dynamic or real-time environments such as continuous patient monitoring, wearable sensor data analysis, or adaptive clinical decision systems.

Ordinary Least Squares (OLS) Regression

A foundational statistical method used to estimate the linear relationship between one or more independent variables and a continuous outcome, by minimizing the sum of squared errors between observed and predicted values. In medicine, OLS regression is widely applied in epidemiological research to quantify associations—for example, between risk factors like smoking or physical activity and disease outcomes.

Machine Learning

A subfield of artificial intelligence (AI) focused on developing algorithms that can learn patterns from data and make predictions or decisions without being explicitly programmed. Machine learning encompasses various techniques, including supervised learning (trained on labeled data), unsupervised learning (identifying patterns in unlabeled data), and reinforcement learning (learning through trial and error). In medicine, machine learning powers diagnostic tools, predictive analytics, clinical decision support, and patient stratification.

Measurement Bias

A systematic error that arises from inaccuracies or inconsistencies in how data are collected, recorded, or measured. In healthcare, measurement bias can result from variations in diagnostic equipment, lab protocols, or observer practices, potentially misleading AI models and compromising prediction accuracy or fairness across clinical settings.

Multimodal Learning

Multimodal learning is a machine learning approach where models integrate and learn from multiple types (or modalities) of data—such as text, images, audio, and video—to capture richer context and achieve improved performance. By combining complementary information across modalities, multimodal learning enables more robust, accurate, and context-aware predictions.

Multi-Omics Data Integration

The use of AI and machine learning techniques to combine and analyze multiple layers of biological data, such as genomics, proteomics, transcriptomics, and metabolomics. This integrative approach supports precision medicine by uncovering complex molecular interactions, identifying disease mechanisms, and tailoring interventions to individual patients based on comprehensive biological profiles.

Natural Language Processing (NLP)

A field of artificial intelligence that enables computers to understand, interpret, generate, and respond to human language in a meaningful way. In healthcare, NLP is used to extract insights from clinical notes, automate documentation, analyze patient narratives, and process unstructured text in electronic health records (EHRs), supporting more efficient and informed clinical decision-making.

Naive Bayes

A probabilistic classification algorithm based on Bayes' Theorem, assuming that features are independent given the outcome. Despite its simplicity, Naive Bayes performs well in many applications, including medical text classification, disease prediction, and triaging patient notes, particularly when interpretability and speed are important.

Negative Predictive Value (NPV)

The probability that individuals with a negative test result truly do not have the condition or disease being tested for. In medical practice, a high NPV means that clinicians can confidently rule out a disease or condition when a test result is negative.

Neural Networks

A neural network is a computational model inspired by the structure and function of biological neurons, consisting of interconnected nodes (neurons) organized into layers. Each node processes input data using weights, biases, and activation functions, enabling the network to learn complex patterns, relationships, or representations within data. Neural networks form the basis of deep learning, widely used in tasks such as classification, regression, image recognition, and natural language processing.

Node (Neural Network Node)

A basic computational unit in a neural network, also known as a neuron. Each node receives input values, applies a mathematical operation using weights, a bias term, and an activation function, and outputs a value to subsequent nodes. Nodes are the building blocks of neural networks and enable the learning of complex patterns from data.

Omitted Variable Bias

A type of error that occurs when important variables influencing the outcome are excluded from a predictive model. This omission can lead to incorrect estimates of relationships between predictors and outcomes, compromising both the validity and generalizability of AI models. In healthcare, it can result in misleading predictions and flawed clinical decisions, especially in observational data.

Overfitting

A modeling error that occurs when a machine learning model learns patterns too closely tied to the training data, including noise or random fluctuations, rather than capturing the underlying general trends. In medicine, overfitting can lead to models that perform well on historical patient data but fail to generalize to new cases. For example, a predictive model trained on a small dataset of ICU patients may appear highly accurate but underperform when applied to a broader hospital population, limiting its clinical utility.

Patient Digital Twin

A virtual representation of an individual patient that uses AI to simulate disease progression, predict treatment responses, and personalized care. Used in precision medicine, simulating individualized responses to treatments. They can help clinicians test different therapeutic strategies, personalize drug dosages, and improve surgical planning by simulating patient-specific outcomes before interventions.

Pipeline

A structured sequence of steps in data processing and model development that transforms raw data into meaningful outputs. In AI and medicine, a pipeline typically includes stages such as data collection, preprocessing, feature engineering, model training, evaluation, and interpretation. Well-designed pipelines ensure reproducibility, efficiency, and accuracy in clinical decision support, diagnostics, and medical research applications.

Precision

Also called positive predictive value, measures the proportion of correctly identified positive cases out of all predicted positive cases. Statistically, this is calculated as $TP / (TP + FP)$, where TP = true positives and FP = false positives (TP+FP = all positive calls). In medicine, precision is crucial for AI models used in diagnostics, such as cancer detection, where a high precision rate ensures that most patients identified as having the disease actually do. False positives in medicine may lead to unnecessary treatments or interventions.

Precision Medicine

Precision medicine—also called personalized medicine—is an approach to disease prevention, diagnosis, and treatment that considers individual differences rather than a one-size-fits-all strategy. This can involve using biomarkers, genetics, environment, and lifestyle data. Precision medicine reflects tailored medical care to the unique characteristics of each patient or group of patients.

Predictive Analytics

The use of data, statistical algorithms, and machine learning techniques to identify the likelihood of future outcomes based on historical data. In healthcare, predictive analytics is used to detect disease early, anticipate complications, stratify patient risk, and personalize treatment plans.

Prescriptive Analytics

A type of data analysis that uses AI, machine learning, and optimization techniques to recommend the best course of action based on predictive insights. In medicine, prescriptive analytics helps guide clinical decision-making, optimize treatment plans, and improve healthcare resource allocation by suggesting actionable steps tailored to patient needs and system constraints.

Preprocessing

A critical step in the machine learning pipeline that involves cleaning, transforming, and structuring raw data before it is used to train a model. Preprocessing tasks may include handling missing values, normalizing inputs, encoding categorical variables, and removing noise. In medicine, preprocessing ensures that diverse clinical data sources, such as EHRs, lab results, and sensor outputs, are consistent, interpretable, and suitable for reliable model training.

Primacy Effect Bias

A cognitive bias in which information encountered early disproportionately influences perception or decision-making. In AI, particularly in large language models or time-sequenced data analysis, early inputs may overly influence predictions or recommendations. In healthcare, this can lead to misinterpretation if models give undue weight to initial symptoms, diagnoses, or data entries, underscoring the importance of input ordering and contextual balancing in algorithm design.

Privacy

The right of individuals to control their personal information, and in medicine, this includes personal health information. AI systems collect and process a lot of personal information. It is important to ensure that this information is used in a way that respects individual privacy and follows local regulations.

Quantum Computing in Medicine

The use of quantum mechanics-based computing to solve complex problems exponentially faster than classical computers. This computing could accelerate AI applications related to drug discovery and improve molecular modeling for personalized medicine.

Random Forest

An ensemble learning algorithm that builds multiple decision trees during training and outputs the average prediction (for regression) or majority vote (for classification). Trees are constructed using random subsets of the data and features, improving accuracy and reducing overfitting. In healthcare, random forests are widely used for disease classification, risk stratification, and identifying key predictors from large clinical datasets.

Recall

Also known as sensitivity, recall measures the proportion of actual positive cases that were correctly identified by an AI or machine learning model. It is important in healthcare applications where missing a positive case (e.g., failing to detect a disease) can have serious consequences.

Recurrent Neural Networks (RNNs)

A type of neural network designed to handle sequential or time-dependent data by maintaining memory of previous inputs. In healthcare, RNNs are particularly useful for analyzing longitudinal patient records, monitoring physiological signals over time (e.g., ECG or EEG), and detecting temporal trends in chronic disease progression or treatment response.

Reinforcement Learning (RL)

A type of machine learning where an AI agent learns by interacting with an environment and receiving rewards or penalties for its actions. In medicine, RL is used in areas like personalized treatment plans, robotic surgery optimization, and drug discovery simulations.

Remote Patient Monitoring (RPM)

Technology-enabled systems that use connected devices to collect and transmit patient health data in real-time, allowing clinicians to track conditions outside of clinical settings. AI-enhanced RPM can detect early warning signs, prompt timely interventions, and reduce hospital readmissions.

Regulatory AI Compliance

The process of ensuring that AI models used in healthcare adhere to legal, ethical, and safety standards set by regulatory bodies (e.g., FDA, EMA, HIPAA, Health Canada). This includes model validation, risk assessment, transparency, and data protection to ensure responsible deployment in clinical practice.

Robotics

The application of AI-powered machines and automation to perform tasks traditionally conducted by humans in healthcare settings. Medical robotics includes robotic-assisted surgery, rehabilitation support, automated medication dispensing, and AI-integrated diagnostic tools. These technologies enhance precision, reduce human error, and improve both clinical efficiency and patient outcomes.

Sampling Bias

Arises when the dataset used to train an AI model is not representative of a target population. This can lead to skewed predictions and poor generalizability, particularly when minority or high-risk groups are underrepresented.

Self-attention

A mechanism used by current transformer models and previously by encoder-decoder architectures to determine the most important components of an input. By doing this, an AI model can focus on the more important aspects of data rather than considering all components equally. This is important in medical applications of these models as they can determine the parts of data that should contribute the most to the output, whether that be in medical imaging, general diagnostics, or treatment planning.

Self-supervised Learning

Self-supervised learning is a machine learning paradigm in which models learn useful representations from unlabeled data by predicting hidden or masked parts of the input or by solving pretext tasks generated directly from the data itself. This method leverages intrinsic structure within the data to reduce dependency on labeled examples, thereby enabling models to learn robust features without human-annotated supervision.

Semi-supervised Learning

A machine learning method that combines a small amount of labeled data with a larger pool of unlabeled data to train models more efficiently. In healthcare, semi-supervised learning helps reduce annotation costs while improving performance in scenarios like rare disease detection, clinical text analysis, and imaging tasks where labeled data is limited.

Sensitivity

The ability of a model or test to correctly identify individuals with a condition (true positives). Sensitivity is especially important in medical screening and diagnostic tools, where missing a positive case could delay treatment or worsen outcomes. It is calculated as: $\text{Sensitivity} = \text{TP} / (\text{TP} + \text{FN})$, and is also referred to as Recall in machine learning contexts.

Sentiment Analysis

A natural language processing (NLP) technique used to determine the emotional tone or sentiment expressed in text data. It categorizes text as positive, negative, or neutral and can be enhanced with deep learning for more nuanced interpretation. In medicine, sentiment analysis is applied to analyze patient feedback, assess public health trends from social media, and gauge healthcare provider satisfaction. It helps hospitals and healthcare organizations improve patient experience by identifying concerns in reviews, monitoring mental health discussions, and understanding patient emotions from clinical notes or support chatbots.

Shapley Additive exPlanations (SHAP)

A model-agnostic method for interpreting machine learning predictions by quantifying the contribution of each input feature to the output. SHAP values represent how much each feature increases or decreases a specific prediction compared to the average model output. In healthcare, SHAP enhances transparency and trust by helping clinicians understand why an AI system made a particular diagnostic or treatment recommendation.

Small Language Models (SLMs)

A subset of language models that are more compact and computationally efficient than large language models (LLMs), often designed for specific domains or tasks. In medicine, SLMs are used for targeted applications such as clinical documentation assistance, medical chatbot interactions, and specialized natural language processing with lower computational requirements. Beneficial for mobile and embedded medical AI applications (edge computing).

Specificity

A performance metric that measures a model's ability to correctly identify true negative cases—those without the condition of interest. High specificity is crucial in medicine to reduce false positives, avoid unnecessary treatments, and minimize patient anxiety. Specificity is calculated as: $\text{Specificity} = \text{TN} / (\text{TN} + \text{FP})$, where TN = true negatives and FP = false positives.

Supervised Learning

A machine learning approach where models learn from labelled training data to identify relationships between data features and labels. This is often used in medical applications where the outcome (i.e. label) is clearly defined and captured, and the task is also known (e.g., prediction or classification).

Swarm Intelligence

A decentralized approach to problem-solving inspired by the collective behavior of natural systems, such as bird flocks or insect colonies. In healthcare, swarm intelligence has been applied to optimize hospital logistics, coordinate robotic surgical systems, and design adaptive scheduling algorithms that improve operational efficiency.

Synthetic Data

Artificially generated data that closely resembles real healthcare data while preserving patient privacy. In healthcare, synthetic data enables AI model training, validation, and research without compromising patient confidentiality and personal health information.

Telemedicine

The use of digital communication technologies to deliver healthcare services remotely. This includes virtual consultations, remote monitoring, and digital follow-ups, enabling timely access to care, particularly for patients in rural or underserved areas. Telemedicine enhances care continuity and reduces barriers to healthcare delivery.

Temporal Bias

A bias introduced when data used to train a model reflects outdated practices, disease patterns, or population characteristics. In healthcare, temporal bias can compromise model performance if the data do not account for evolving clinical guidelines, diagnostic criteria, or shifts in population health trends. Monitoring for and mitigating temporal bias is essential for maintaining model accuracy and relevance in real-time clinical settings.

Training

The process by which an AI model learns from labeled data by adjusting its internal parameters to minimize prediction errors. In healthcare, training typically involves large datasets such as patient records, medical images, or clinical text. The goal is to create accurate and generalizable models for tasks like disease classification, prognosis prediction, and clinical decision support.

Training Dataset

A subset of data used to train an AI model. In medical AI, a training dataset typically includes labeled examples, such as diagnosed cases in radiology images or patient outcomes, to help the model learn patterns before being tested on unseen data.

Transferability

The ability of a machine learning model to maintain performance when applied to different datasets, populations, or clinical settings beyond its original training environment. In medicine, assessing transferability is critical to ensure that AI tools developed in one institution (e.g., a hospital or health system) can generalize effectively to others, supporting broader, equitable deployment without retraining from scratch.

Transfer Learning

A machine learning technique in which a model developed for one task is reused or fine-tuned for a different but related task. This approach significantly reduces the amount of labeled data and training time needed for new applications. In medicine, transfer learning is commonly used in medical imaging, where models pre-trained on large image datasets (e.g., ImageNet) are adapted to detect diseases in radiology, pathology, or dermatology images with improved efficiency and accuracy.

Transformer Models

A deep learning architecture designed for processing sequential data, particularly in natural language processing (NLP). Transformers use self-attention mechanisms to weigh the importance of different parts of input data, allowing them to capture long-range dependencies and contextual meaning more effectively than traditional models like recurrent neural networks (RNNs). In medicine, transformer models are widely used for clinical text analysis, medical chatbot development, and biomedical research. Examples include models like BERT and GPT.

Transparency

Refers to the degree to which the inner workings, decision-making processes, and limitations of an artificial intelligence system are understandable and accessible to users, developers, and other stakeholders. Transparency is important for ensuring that AI systems are fair and accountable.

Trust

The confidence that an AI system will perform reliably, ethically, and in alignment with user expectations and clinical standards. In healthcare, trust is critical for adoption and use; it is built through model transparency, rigorous validation, ethical oversight, and ongoing human oversight in clinical decision-making.

Unsupervised Learning

Machine learning where the model finds patterns in data without labelled examples. Instead of being given explicit instructions, the model autonomously groups similar data points or detects hidden relationships through clustering, dimensionality reduction, or anomaly detection techniques. In medicine, unsupervised learning is used for patient segmentation, disease subtyping, anomaly detection in medical imaging and drug discovery.

Validation

A dataset used to assess a model's performance and generalizability after training. Internal validation uses a hold-out portion of the original dataset, while external validation uses entirely independent data from a different source or population. In healthcare, external validation is essential for evaluating how well a model performs across different settings, institutions, or patient groups.

Validation Dataset

A dataset used to assess a machine learning model's performance and generalizability. Internal validation involves using a reserved portion of the original dataset not seen during training. External validation uses data from a separate source, population, or institution, offering a more rigorous test of the model's ability to generalize to new, real-world scenarios—critical in healthcare applications for ensuring reliability across settings.

Variance

A type of error in machine learning that occurs when the algorithm is sensitive to small changes in the data. A number of factors, such as the size of the training data or the complexity of the algorithm, can cause variance.

Violin Plot

A plot that combines a box plot and a kernel density plot, illustrating data distribution and its probability density. In medicine, violin plots are useful for analyzing patient data distributions, such as comparing biomarker levels across different patient groups, visualizing the spread of treatment responses, or assessing variations in hospital stay durations. They help researchers and clinicians interpret complex data patterns more effectively than traditional box plots.

Word Cloud

A visual representation showcasing the frequency of words in a given dataset. In medicine, word clouds are used to analyze and summarize large volumes of unstructured health data, such as patient feedback, medical literature, or clinical notes. They can help identify common symptoms in patient reports, highlight key topics in public health discussions, or provide insights into emerging trends in medical research.

Zero-Shot Learning

A machine learning approach where a model can make predictions about classes or concepts it has never seen during training by leveraging semantic relationships or descriptive information. In medicine, zero-shot learning enables models to interpret rare diseases, novel symptoms, or emerging health threats without needing explicit examples. For instance, a model trained on general clinical data might recognize an unfamiliar condition described in text or flagged in clinical notes, supporting decision-making in scenarios where annotated datasets are limited or unavailable.

REFERENCES

Bard, (2023, May 3). Glossary of 15 to 40 words that are important to understand when looking at artificial intelligence in healthcare. Retrieved from bard.google.com

Chiolero A, Buckeridge D. [Glossary for public health surveillance in the age of data science](#). J Epidemiol Community Health 2020;74:612-616.

Patrishkoff, David, and Hoyt, Robert E.. No-Code Data Science: Mastering Advanced Analytics, Machine Learning, and Artificial Intelligence. 2023. [Lulu.com](https://lulu.com)

Rajpurkar, P., Chen, E., Banerjee, O. et al. [AI in health and medicine](#). Nat Med 28, 31–38 (2022).

Mooney, Stephen J. and Pejaver, Vikas. [Big Data in Public Health: Terminology, Machine Learning, and Privacy Annual Review of Public Health](#) 39:1, 95-112 (2018)

Wang, J., & Redelmeier, D. A. (2024). [Cognitive biases and artificial intelligence](#). NEJM AI, 1(12), Alcs2400639.

ACKNOWLEDGMENT

We gratefully acknowledge AI Educator and Researcher Dr. Orest B. Boyko, MD PhD, who raised the idea of the glossary and presented the initial concept to T-CAIREM. We are also grateful for this leadership and collaboration that catalyzed the concept of an AI Living Glossary for medical education.



THE USE OF AI IN GENERATING AND UPDATING THIS GLOSSARY

This glossary was developed using generative artificial intelligence (AI) tools to synthesize, structure, and refine definitions. AI-assisted drafting facilitated the initial compilation of terms, drawing from diverse sources to ensure comprehensive coverage of key concepts. However, subject matter experts reviewed, edited, and validated all entries to ensure accuracy, clarity, and contextual relevance.

FULL LIST OF WORKING GROUP MEMBERS WHO CONTRIBUTED TO THIS GLOSSARY

Laura Rosella, Education Lead, [Temerty Centre for Artificial Intelligence Research and Education in Medicine \(T-CAIREM\)](#), [University of Toronto](#)

Yindalon Aphinyanaphongs, Assistant Professor, [Department of Population Health](#), [New York University](#)

James Barry, Associate Professor, [Pediatrics-Neonatology](#), [School of Medicine](#), [University of Colorado Anschutz Medical Campus](#)

Orest Boyko, Professor of Clinical Sciences, [Roseman University College of Medicine](#)

Harvey Castro, Healthcare Advisory, [University of Texas at San Antonio](#)

Anthony Chang, Pediatric Cardiologist, [Children's Hospital of Orange County \(CHOC\)](#)

Shane Eaton, Director of Operations, [Artificial Intelligence in Medical Systems Society](#)

Bradley Erickson, Medical Director for AI, [Mayo Clinic](#)

Jamie Fairclough, Associate Dean, [Roseman College of Medicine](#) | Director, [Dartmouth College](#)

Diana Ferro, Health Care Research Officer and Data Scientist, [Ospedale Pediatrico Bambino Gesù \(OPBG\)](#), [IRCCS](#), [Rome, Italy](#)

Gregg Gascon, Advisor, [Data Science](#), [OhioHealth](#)

Alexander Gimson, Consultant Transplant Hepatologist, [Cambridge University Hospitals NHS Foundation Trust](#)

Robert Hoyt, Associate Clinical Professor, [Virginia Commonwealth University School of Medicine](#)

Nihal Haque, Faculty Affiliate, [Geriatrician and Hospitalist](#), [North York General Hospital](#), [Temerty Centre for Artificial Intelligence Research and Education in Medicine \(T-CAIREM\)](#), [University of Toronto](#)

Johanna Kim, Executive Director, [Centre for Artificial Intelligence in Medicine and Imaging](#)

Alfonso Limon, Senior Data Scientist, [Children's Hospital of Orange County \(CHOC\)](#)

Pietro Lio, Computer Scientist, [Cambridge University](#)

Muhammad Mamdani, Director, [Temerty Centre for Artificial Intelligence Research and Education in Medicine \(T-CAIREM\)](#), [University of Toronto](#)

Piyush Mathur, Staff Anesthesiologist and Critical Care Physician, [Cleveland Clinic](#)

Ryan McAdams, Professor, [Department of Pediatrics](#), [University of Wisconsin School of Medicine and Public Health](#)

Ehsan Misaghi, Director of Education & Research, [Artificial Intelligence in Medical Systems Society](#)

Abhishek Moturu, Education Trainee Co-Lead, [Temerty Centre for Artificial Intelligence Research and Education in Medicine \(T-CAIREM\)](#), [University of Toronto](#)

Hari Mudipalli, Data Scientist and Engineer, [Roseman College of Medicine](#)

Arwa Nada, Associate Professor, [The University of Tennessee Health Science Centre](#)

Oded Nov, Technology Management and Innovation Department | Chair, [New York University](#)

Gemma Postill, Education Trainee Co-Lead, [Temerty Centre for Artificial Intelligence Research and Education in Medicine \(T-CAIREM\)](#), [University of Toronto](#)

Mijanou Pham, Executive Director, [American Board of Artificial Intelligence in Medicine](#)

Zoryana Salo, Centre Manager, [Temerty Centre for Artificial Intelligence Research and Education in Medicine \(T-CAIREM\)](#), [University of Toronto](#)

Marianne So, Curriculum & Event Coordinator, [Temerty Centre for Artificial Intelligence Research and Education in Medicine \(T-CAIREM\)](#), [University of Toronto](#)

Ying Wan, Co-Lead Education, [Artificial Intelligence in Medical Systems Society](#)

Alaa Youssef, Post Doctoral Researcher, [Centre for Artificial Intelligence in Medicine and Imaging](#), [Stanford University](#)

Nathan Yung, Assistant Clinical Professor | Cybersecurity Project Lead, [University of California San Diego](#)