

Key Terms and Clinical Thresholds

Master reference card, "Healthcare Analytics and AI: Building Systems That Actually Work"

by [Richard J. Young, Ph.D.](#)

richard@deepneuro.ai

Compact glossary: the terms and acronyms you must know, each with a one-line definition and the number to remember.

PART 1: Clinical and Business

Medical and Clinical

Sepsis: life-threatening organ dysfunction from a dysregulated response to infection; **5–7% of inpatients**, the marquee failure case (Epic model AUROC 0.63).

Readmission: return to hospital within a window, usually **30 days**; CMS penalizes excess rates; the most-studied ML target.

Comorbidity: two or more chronic conditions at once; quantified by the **Charlson** or **Elixhauser** index.

EHR: Electronic Health Record; structured (labs, dx, meds) plus unstructured (notes) data; the primary source for clinical AI.

ICD-10: WHO diagnosis code set; **~72,000 codes**; a core feature in most pipelines.

CPT: AMA procedure code set used for billing and reimbursement.

SNOMED CT: clinical terminology of **>350,000 concepts**; target ontology for NLP normalization.

Diabetic retinopathy: diabetes complication causing blindness; first autonomous FDA-cleared AI (IDx-DR), **>90% sensitivity**.

Atrial fibrillation (AFib): irregular atrial rhythm; **5× stroke risk**, often silent; a wearable-screening target.

Survival analysis: methods for time-to-event data with censoring (readmission, progression, death).

Palliative care / Hospice: comfort-focused care; hospice for prognosis **≤6 months**; AI raises acute prognostic ethics.

Substance use disorder (SUD): harmful substance use; records get extra protection under **42 CFR Part 2**.

SDOH: non-medical drivers of health; account for **30–55%** of outcome variation; proxy-bias risk.

Insurance, Payment, and Operations

Prior authorization: plan pre-approval before a service; high-volume AI automation target.

Claims adjudication: insurer evaluation of a claim against policy, coding, and rates to set payment.

Fraud, waste, and abuse (FWA): improper billing spanning intent to inefficiency; costs **\$100B–\$300B/yr** in the U.S.

Capitation: fixed per-member-per-month payment; shifts risk to the provider.

Fee-for-service (FFS): pay per service; volume incentive; historically dominant.

Value-based care (VBC): payment tied to outcomes and cost, not volume.

DRG: inpatient stays grouped for prospective payment; Medicare's base since **1983**.

HCC: Hierarchical Condition Category risk model; **V28 = 115 categories**; coding intensity adds **~\$84B/yr** excess (MedPAC).

Risk adjustment: adjusts payment for patient health status; drives MA capitation and AI coding.

Medicare Advantage (MA): privatized Medicare; **54% of beneficiaries** (**~35.2M, 2026**); risk-adjusted capitation.

Dual eligible: qualifies for Medicare and Medicaid; **~12.5M**, 19% of Medicare pop but 35% of its spend.

Star Rating (CMS): 1–5 MA quality score; **≥4 stars** earns a 5% (or 10%) quality bonus.

Medical loss ratio (MLR): share of premium spent on care; ACA floor **80–85%**.

Upcoding: assigning a more severe code than documentation supports; a form of abuse.

Core ML and Evaluation

AUROC: discrimination (ranking) metric; **0.5 chance, 1.0 perfect**; blind to calibration.

AUPRC: ranking on the rare class; baseline = prevalence; preferred under imbalance.

Brier score: proper score of calibration + discrimination; **0 perfect**, 0.25 random at balance.

Calibration: predicted probability matches observed rate; "30%" should see ~30 events per 100.

Sensitivity (recall): share of true positives caught; prioritized for must-not-miss screening.

Specificity: share of true negatives cleared; trades off against sensitivity.

F1 score: harmonic mean of precision and recall; better than accuracy under imbalance.

Cross-validation: rotate train/test folds; use **temporal** splits for clinical data.

External validation: test on a new site/time/population; minimum bar before deployment.

Overfitting: learning noise; strong on train, weak on new data.

Regularization: L1/L2, dropout, early stopping to constrain complexity.

Gradient boosting (XGBoost): sequential trees fit to residuals; top performer on tabular clinical data.

Random forest: bagged decision trees; overfit-resistant with built-in importance.

Logistic regression: interpretable linear baseline for binary outcomes.

K-means: partition into k clusters by centroid distance; patient segmentation.

SMOTE: synthesizes minority examples by interpolation; can generate clinically implausible records.

Cox proportional hazards: semi-parametric time-to-event model; outputs hazard ratios.

Kaplan-Meier: non-parametric survival curve honoring right-censoring.

Deep Learning, NLP, and LLMs

CNN: convolutional network for grid data; dominant in medical imaging.

U-Net: encoder-decoder with skip connections; the default for medical segmentation.

Transformer: self-attention architecture; backbone of modern LLMs and ViTs.

Attention mechanism: weighs every token against every other; enables context-dependent representations.

BERT: bidirectional encoder; clinical variants BioBERT, ClinicalBERT for NER and classification.

LLM: large language model for documentation, QA, and agentic reasoning; risks hallucination and sycophancy.

Foundation model: broadly pre-trained model adapted downstream (GPT-4, Med-PaLM).

RAG: retrieval-augmented generation; grounds output in

retrieved evidence to cut hallucination.

Fine-tuning: continue training a pre-trained model on task data.

Embedding: dense vector where geometric proximity = semantic similarity.

NER: named entity recognition; extracts meds, dx, procedures from notes.

SHAP: Shapley-value feature attributions; local and global explanations.

Explainability (XAI): methods making predictions understandable; SHAP, LIME, attention, counterfactuals.

Hallucination: fluent but fabricated model output; acutely dangerous for dosing or interactions.

Reinforcement learning (RL): learns sequential decisions by maximizing reward; sepsis, ventilator weaning.

Privacy, Federation, and Data

Federated learning: train across sites sharing only weight updates, never raw data.

FedAvg: canonical federated algorithm averaging local model updates centrally.

Differential privacy: bounds any one record's impact via calibrated noise; $\epsilon \leq 1$ **strong**, >10 **weak**.

k-Anonymity: each record indistinguishable from $\geq k - 1$

others on quasi-identifiers.

De-identification: remove PHI via Safe Harbor (18 identifiers) or Expert Determination.

Data drift: input distribution shifts between training and deployment; triggers monitoring.

Synthetic data: generated records (e.g., CTGAN) preserving statistics while protecting patients.

Regulatory and Legal

HIPAA: U.S. health-data law (1996); Privacy, Security, and Breach Notification rules.

PHI: protected health information; **18 identifier categories** must be safeguarded or removed.

HITECH (2009): strengthened HIPAA penalties; extended rules to business associates.

42 CFR Part 2: heightened protection for SUD records; consent stricter than HIPAA.

FDA 510(k): premarket pathway showing substantial equivalence to a predicate; common for AI devices.

SaMD: Software as a Medical Device; FDA category for most clinical AI (Class I–III by risk).

GDPR (2018): EU data-protection law; data-subject rights and heavy penalties.

EU AI Act: risk-tiered EU AI law (in force 2024); healthcare generally **high-risk**.

HL7 FHIR: RESTful interoperability standard; the expected EHR integration layer.

Model card: structured disclosure of intended use, subgroup performance, and limitations.

IRA (2022): \$2,000 annual Part D out-of-pocket cap (2025); CMS drug-price negotiation.

V28: 2026 CMS-HCC model; **189**→**115** categories, cuts MA pay ~\$10–12B/yr.

Ethics, Fairness, and Safety

Algorithmic bias: systematic unfair errors from biased data, labels, or proxies.

Equalized odds: fairness: equal TPR and FPR across groups; usually the defensible clinical choice.

Demographic parity: equal selection rate across groups; conflicts with calibration.

Fairness metric: parity measures that are mutually incompatible in general (**Chouldechova impossibility**).

Disparate impact: neutral practice that harms a protected group; **four-fifths (80%) rule**.

Proxy discrimination: a permitted feature (zip, language)

stands in for a protected attribute.

Deskilling: erosion of clinician competence from over-reliance on AI.

Alert fatigue: desensitization to CDS alerts; override rates >90%.

Sycophancy: LLM agreeing with a user regardless of accuracy; RLHF side effect.

Guardrail: external input/output safety layer enforcing scope and blocking harmful content.

Prompt injection: malicious instructions in input or retrieved context hijacking an LLM.

If you remember one thing. Numbers without context mislead. $\epsilon \leq 1$ is strong privacy only for the query it covers; AUROC 0.5 to 1.0 says nothing about calibration; “>90% override” is a workflow verdict, not a model verdict. Every threshold on this card is a starting point for a conversation, not a passing grade.

Richard J. Young, Ph.D. · richard@deepneuro.ai · Senior AI Research Scientist, UnitedHealth Group · UNLV Lee Business School · ORCID 0000-0002-1109-7552