Comparisons and Roadmaps

Guide to AI and Data Courses

The big table | Often-confused pairs | Which course? | From interest to course | Roadmaps

Sercan KÜLCÜ · Giresun University · Computer Engineering

© Sercan Külcü. All rights reserved. · Interactive HTML version of these slides:
https://sercankulcu.github.io/files/ai_courses/slides/Chapter_40_Comparisons_and_Roadmaps.html

The Big Table: I. Core

Course	Approach	Input	Output	Success measure
Artificial Intelligence	search, logic, probability and learning combined	problem definition: states, actions, goal	action sequence, plan or decision	reaching the goal, cost, rationality
Machine Learning	statistical learning	labeled or unlabeled examples	a model that makes predictions	error on unseen data
Neural Networks	network of simple weighted units	numeric feature vector	class or numeric prediction	loss and accuracy
Deep Learning	multilayer representation learning	raw data: images, text, audio	learned representation and prediction	accuracy, loss, scaling

The Big Table: II. Branches of AI

Course	Approach	Input	Output	Success measure
NLP	linguistics, statistics and deep learning	text	label, translation, summary, answer	accuracy, BLEU, human evaluation
Reinforcement Learning	trial and error with rewards	environment: state, reward	policy: state → action	total expected reward
Pattern Recognition	statistical decision theory	measurement or feature vector	class label	classification error
Fuzzy Logic	partial membership and linguistic rules	measurements, expert rules	control signal or decision	control quality, expert agreement
Evolutionary Optimization	population-based heuristic search	objective, solution encoding	good, not proven optimal, solution	quality, evaluation count
Expert Systems	symbolic logic and rules	facts and expert rules	conclusion with justification	expert agreement, explainability
Graphical Models	probability and graphs combined	variables, dependencies, data	conditional probability, best explanation	likelihood, inference accuracy
Generative Models	learning the data distribution	huge, mostly unlabeled data	new samples: images, text, audio	realism and diversity (FID etc.)
Multi-Agent Systems	interacting autonomous agents	agents, goals, protocols	joint behavior, equilibrium	social welfare, stability
Explainable AI & Ethics	human-centered evaluation	model, its data and context of use	explanation, audit report, principles	interpretability, fairness, accountability

The Big Table: III. Signals, Images and Audio

Course	Approach	Input	Output	Success measure
Signal Processing	mathematical transforms and filters	sampled signal	filtered signal, spectrum	signal-to-noise ratio, frequency response
Image Processing	pixel-level transformations	image	processed image or regions	visual quality, PSNR
Computer Vision	extracting meaning from images	image or video	label, box, depth, 3D structure	accuracy, mAP, IoU
Speech and Audio	acoustic signal and language model	audio waveform	text, speaker, sound event	word error rate (WER)
Medical Imaging	image processing plus clinical knowledge	CT, MRI, X-ray, ultrasound	segmentation, measurement, diagnostic suggestion	sensitivity, specificity, Dice
Remote Sensing	physical sensing and image analysis	multispectral satellite imagery	land cover map, change	map accuracy, kappa
Computer Graphics	geometry and light simulation	scene: models, lights, camera	image or animation	realism, frame rate
Robotics	sense-plan-act loop	sensor measurements	motor commands, path	task success, safety
Biometrics	matching and verification	face, fingerprint, iris, voice	identity decision	FAR, FRR, EER

The Big Table: IV. Data

Course	Approach	Input	Output	Success measure
Data Mining	discovery-driven analysis	large tables or transaction data	rules, clusters, outliers	interestingness, support, confidence
Data Science	question, data, analysis, story	raw, messy data	report, visualization, model	impact on decisions, accuracy
Big Data	distributed storage and computation	huge, fast-flowing data	processed data, query results	scalability, latency, cost
IR & Recommenders	matching and ranking	query, documents, user history	ranked result list	precision, recall, NDCG
Time Series	time-dependent statistical models	timestamped measurements	forecast, trend, seasonality	forecast error (MAE, RMSE)
Social Network Analysis	graph theory and statistics	nodes and links	centrality, communities, diffusion	modularity, interpretability

The Big Table: V. Mathematical Foundations

Course	Approach	Input	Output	Success measure
Linear Algebra	vector spaces and transformations	vectors, matrices	solution, decomposition, transformation	accuracy, numerical stability
Probability & Statistics	probability models and inference	sample data	estimate, confidence interval, test result	unbiasedness, power, significance
Optimization	mathematical model of objective and constraints	objective function and constraints	optimal or near-optimal solution	objective value, convergence rate
Information Theory	measuring information with probability	source and channel model	bounds, codes, compression rate	bits, channel capacity

The Big Table: VI. Systems and Applications

Course	Approach	Input	Output	Success measure
MLOps	software engineering meets ML	trained model, data pipeline	a running, monitored service	latency, uptime, data drift
Embedded AI	resource-constrained inference	trained model, device constraints	small, fast model	latency, energy, memory
GPU & Parallel Computing	data-parallel computation	parallelizable algorithm	accelerated program	speedup, memory bandwidth
AI in Cybersecurity	security knowledge plus ML	network traffic, logs, files	alert, class, risk score	detection rate, false alarm rate
Bioinformatics	biology, algorithms and statistics	DNA, protein sequences, expression data	alignment, structure, function prediction	biological validity

Math and Data Needs (1/3)

Course	Math load	Data needs
Artificial Intelligence	medium: logic, probability, graphs	little; often a model and knowledge suffice
Machine Learning	high: linear algebra, probability, optimization	moderate to large
Neural Networks	medium-high: derivatives, matrices	moderate
Deep Learning	high: derivatives, matrices, probability	large; needs GPUs
NLP	medium	large text corpora
Reinforcement Learning	high: probability, dynamic programming	generated through interaction
Pattern Recognition	medium-high: probability, linear algebra	labeled examples
Fuzzy Logic	low-medium: set theory	little; expert knowledge
Evolutionary Optimization	low-medium	none; evaluations suffice
Expert Systems	low: propositional and predicate logic	little; a knowledge base
Graphical Models	high: probability	moderate
Generative Models	high	very large
Multi-Agent Systems	medium: game theory	little; simulation

Math and Data Needs (2/3)

Course	Math load	Data needs
Explainable AI & Ethics	low-medium	the model's own data
Signal Processing	high: complex numbers, Fourier	no learning; the signal suffices
Image Processing	medium: convolution, Fourier	little; learning not required
Computer Vision	high: geometry, linear algebra	many labeled images
Speech and Audio	medium-high	many hours of recordings
Medical Imaging	medium-high	little; labels are expensive
Remote Sensing	medium	plenty; labels limited
Computer Graphics	high: linear algebra, geometry	none
Robotics	high: transformations, control	physical experiments and simulation
Biometrics	medium	few samples per person
Data Mining	medium	large
Data Science	medium: statistics	moderate to large
Big Data	low	very large

Math and Data Needs (3/3)

Course	Math load	Data needs
IR & Recommenders	medium: vector spaces	lots of interaction data
Time Series	medium-high	long history
Social Network Analysis	medium: graphs, matrices	network data
Linear Algebra	the theory itself	none
Probability & Statistics	the theory itself	a sample
Optimization	the theory itself	none
Information Theory	high: probability	none
MLOps	low	streaming production data
Embedded AI	medium	training data already exists
GPU & Parallel Computing	low-medium	none
AI in Cybersecurity	medium	imbalanced and constantly changing
Bioinformatics	medium-high	huge sequence datasets

Often-Confused Pairs (1/2)

Pair	Difference
Image Processing ↔ Computer Vision	In image processing both input and output are images; in computer vision the input is an image and the output is meaning: a label, a box, depth.
Artificial Intelligence ↔ Machine Learning	AI can find how to reach a goal by search or inference; ML learns the same behavior by generalizing from data.
Machine Learning ↔ Data Mining	They share algorithms; ML aims to predict new examples, data mining to discover patterns in data.
Data Science ↔ Data Mining	Data mining focuses on finding patterns in large data; data science covers the whole path from question to decision.
Neural Networks ↔ Deep Learning	Deep learning moves on to convolutional, recurrent and transformer architectures; this course builds the basic network and training they rest on.
Pattern Recognition ↔ Machine Learning	ML studies learning algorithms in general; this course focuses on the design from measurement to class decision and on the Bayes error.
NLP ↔ Speech and Audio	Speech recognition goes from the sound wave to text; NLP analyzes meaning once the text is available.

Often-Confused Pairs (2/2)

Pair	Difference
Generative Models ↔ Deep Learning	Deep learning is the general toolkit and mostly makes predictions; this course uses those tools to generate new samples from a data distribution.
Computer Graphics ↔ Computer Vision	Computer vision tries to recover the scene from an image; computer graphics produces the image from the scene.
Machine Learning ↔ Probability & Statistics	Statistics emphasizes a model's meaning and uncertainty, ML its predictive performance.
IR & Recommenders ↔ Data Mining	Data mining discovers patterns in data; this course turns discovery into a ranked list for one user's current need.
Expert Systems ↔ Fuzzy Logic	Fuzzy logic lets rules be partly true; this course teaches representing knowledge and chaining rules in general.
Evolutionary Optimization ↔ Optimization	Optimization reaches a provable optimum using derivatives and convexity; this course teaches derivative-free search with no guarantee but wide applicability.
MLOps ↔ Big Data	Big data systems store and process data at scale; this course manages the lifecycle of models trained on that data.

Which Course? (1/2)

For each problem, decide which course it belongs to, then open the answer.

1. **Automatically outlining tumours in a hospital's MRI scans.**
   ▶ **Answer**
2. **Building a "customers who bought this also bought" list for an online shop.**
   ▶ **Answer**
3. **Controlling an oven's temperature with rules like "a bit hot" and "very hot".**
   ▶ **Answer**
4. **A drone learning to stay level in the wind by trying.**
   ▶ **Answer**
5. **Storing and querying per-second readings from a hundred thousand sensors.**
   ▶ **Answer**

Which Course? (2/2)

For each problem, decide which course it belongs to, then open the answer.

1. **Finding the shortest total route plan for school buses.**
   ▶ **Answer**
2. **Auditing whether a credit model treats some groups of applicants unfairly.**
   ▶ **Answer**
3. **Voice commands that work on a phone without an internet connection.**
   ▶ **Answer**
4. **Finding the genes linked to a disease from expression data of thousands of samples.**
   ▶ **Answer**
5. **Catching unusual connections in a bank's network traffic.**
   ▶ **Answer**

CHOOSING

From Interest to Course

Pick one or more interests; the courses that serve them are listed, ranked by how many of your interests they touch.

Images and video | Language and speech | Data and business | Robots and the physical world | Mathematics and theory | Systems and hardware | Trust, society and health | Games and decision making

Nothing selected yet.

Roadmaps (1/2)

Vision track Systems that make sense of images

Linear Algebra → Signal Processing → Image Processing → Computer Vision → Deep Learning → Medical Imaging

---

Language track Systems that work with text and speech

Probability & Statistics → Machine Learning → NLP → Deep Learning → Generative Models → IR & Recommenders

---

Data track Turning data into decisions and forecasts

Probability & Statistics → Data Science → Data Mining → Big Data → Time Series → Social Network Analysis

---

Each track runs from prerequisites to specialization; the tracks cross, and following more than one is normal.

Roadmaps (2/2)

Robotics track Sensing and acting in the physical world

Linear Algebra → Optimization → Robotics → Reinforcement Learning → Multi-Agent Systems → Embedded AI

Theory track Preparing for research

Linear Algebra → Probability & Statistics → Optimization → Information Theory → Graphical Models → Machine Learning

Production track Turning a model into a real product

Machine Learning → Deep Learning → GPU & Parallel Computing → MLOps → Embedded AI → AI in Cybersecurity

Each track runs from prerequisites to specialization; the tracks cross, and following more than one is normal.

Final Thoughts

A common language

Linear algebra, probability and optimization are shared by almost every course; with these in place, moving between courses is easy.

The hub

Machine learning sits at the centre of the map: the branches rest on it and the applications use it.

The question decides

Two courses may use the same algorithm; what separates them is the question they ask and how they measure success.

The field moves

Separate courses today may merge tomorrow; understanding the kinds of relation lasts longer than memorizing the map.

The End

Chapter 40: Comparisons and Roadmaps

← Previous chapter · Map of the field