




Taxonomy of Machine Learning


Learning paradigms | One map of every paradigm along four axes (supervision signal, data access, model type, objective), what each needs, when to use which, and a decision key from problem to paradigm

#LearningParadigms #Taxonomy #ProblemFraming #DecisionKey

Why a taxonomy, and how to read it



One idea at four depths: a paradigm is a set of answers to four questions




What this sheet assumes

Recall in one line, not re-taught

The four classic signals (labels: supervised; none: unsupervised; labels cut from the data: self-supervised; reward: reinforcement) are taught in machine-learning sheet 01. This sheet adds every other paradigm around them and organises them on four axes.

The big map: four axes



Every paradigm sits on four independent axes. Pick one value per axis to describe a project; sheet numbers of this folder in brackets

Signal	What the learner gets	Paradigm	Example
Full, exact labels	y for every x	supervised	fraud labels from charge-backs
Incomplete	y for a few x, many x without y	semi-supervised, active, PU	500 labelled scans, 50 000 unlabelled
Inexact	a coarse label for a group	multiple-instance, weak	slide labelled 'cancer', no region marked


Inaccurate	noisy or heuristic y	weak supervision, noisy-label learning	keyword rules, crowd labels
Demonstrations	expert (state, action) pairs	imitation learning	driving logs for behaviour cloning
Reward	a score after actions	reinforcement learning, bandits	click or no click
Derived from x	a hidden part of x	self-supervised	next token, masked patch
None	x only	unsupervised	customer segments

Axis 1, the supervision signal, from most to least help. Zhou (2018) splits weak supervision into incomplete, inexact and inaccurate

Axis	Values	Question that decides
Data access	batch · online (stream) · continual (task sequence) · active (learner queries) · federated (data stays on sites)	Does all the data exist before training, and can I choose what gets labelled?
Task relation	single task · transfer / domain adaptation · multi-task · meta-learning · few-shot, zero-shot, in-context	Is there a related task or a pretrained model whose knowledge I can reuse?
Model type	generative $p(x, y)$ vs discriminative $p(y x)$ · parametric (fixed size) vs non-parametric (grows with n) · instance-based vs model-based	How much data, and do I need $p(x)$ (generate, detect outliers, handle missing inputs)?
Objective / output	one label · number · ranking · structured output (sequence, tree, graph) · density · policy	Are the outputs independent, ordered, or parts that constrain each other?
Inference set	inductive (rule for any future x) vs transductive (labels for a known set)	Will new, unseen inputs arrive after training?

Axes 2-4: data access, task relation and output

What each paradigm needs and gives




How much each paradigm needs of each resource. 0 = not needed, 3 = essential. Read a row to see the bill before you choose

Paradigm	Bets that	Fails when
Semi-supervised	classes form clusters separated by low-density gaps	unlabelled data has other classes or a different distribution
Weak supervision	several cheap noisy sources are better than chance and roughly independent	sources share the same blind spot
Transfer	source and target share useful features	domains differ too much: negative transfer
Multi-task	tasks share structure	task gradients conflict
Meta-learning	test tasks come from the same task distribution as training tasks	a new task is unlike any training task
In-context learning	the pretrained model already holds the skill; examples only select it	the task needs knowledge or a mapping absent from pretraining
Online / continual	recent data reflects the future	old tasks matter and are overwritten (forgetting)
Active learning	the model's uncertainty points at informative items	the seed model is poor or queries are biased

Each paradigm in one line: the assumption it bets on and how it fails

From problem to paradigm



Decision key: answer the questions in order. Leaves name the paradigm and the sheet that covers it

Worked example: naming the paradigms of one real project

PROBLEM A hospital wants to flag pneumonia on chest X-rays. It has 120 000 unlabelled X-rays, 800 radiologist-labelled ones, free-text reports for 60 000 of them, a public model pretrained on natural images, and a radiologist who can label 50 images a week. Which paradigms, in which order?

GIVEN

- labels: 800 exact, 60 000 reports (noisy text)
- unlabelled: 120 000 images
- pretrained ImageNet-style model
- budget: 50 expert labels per week

1 Signal axis

800 exact + 60 000 noisy + 120 000 none

Why: a mix of incomplete and inaccurate labels

2 Weak labels from reports

labelling functions on report text ('consolidation', 'no acute findings') → probabilistic labels

Why: turns 60 000 reports into training signal (sheet 02)

3 Reuse a model

fine-tune the pretrained network on weak + exact labels

Why: transfer: low-level visual features carry over (sheet 03)

4 Use the unlabelled pool

pseudo-label confident images, consistency on augmentations

Why: semi-supervised (sheet 02)

5 Spend the expert budget

each week, query the 50 most uncertain and diverse images

Why: active learning (sheet 07)

6 Check the deployment domain

test on scanners from other hospitals; adapt if accuracy drops

Why: domain shift (sheet 03)

ANSWER Weak supervision + transfer learning + semi-supervised learning, with an active-learning loop for the radiologist and a domain-shift check before each new site

CHECK ✓ Every data asset is used: 800 exact labels for validation and fine-tuning, reports via weak labels, the unlabelled pool via pseudo-labels, the expert via queries

Look-alikes and myths

Semi-supervised

few labels + many unlabelled x, same task

VS

Self-supervised

no human labels, pretext task from x

Tell: Are there any human labels for the target task?

Transfer learning

reuse a model trained on task A for task B

VS

Multi-task learning

train tasks A and B together

Tell: Are the tasks trained one after the other or at the same time?

Online learning

one stream, one task, update per example

VS

Continual learning

a sequence of tasks, keep the old ones

Tell: Must the model still solve earlier tasks?

Active learning

the model chooses which x get labelled

VS

Semi-supervised

the model labels x itself

Tell: Who produces the new labels: a human oracle or the model?

Inductive

rule for any future input

VS

Transductive

labels for a given unlabelled set only

Tell: Do new inputs arrive after training?

Few-shot learning

K examples per new class, weights may adapt

VS

In-context learning

examples in the prompt, weights frozen

Tell: Does any weight change?

Paradigm names that are easy to mix up

MYTH

✗ A project uses exactly one paradigm

FACT

✓ Most real systems chain three or four

Pretrain self-supervised, transfer, label actively, update online: each paradigm handles one resource.

✗ Unlabelled data always helps

✓ It helps only if its structure matches the classes

Semi-supervised learning can lower accuracy when the cluster assumption fails or unlabelled data comes from another distribution.

✗ Zero-shot means the model saw no related data

✓ It saw no labelled examples of the target classes

CLIP and LLMs saw huge related corpora; zero-shot refers to task labels only.

✗ Reinforcement learning is just supervised learning with a reward

✓ Actions change the data the learner sees

Delayed, sparse rewards and exploration make it a different problem (reinforcement-learning folder).

Misconceptions about paradigms

Interview question

'You have 1 000 labels and 1 M unlabelled items; what do you do?' Strong answer walks the axes: check for a pretrained model (transfer), use the unlabelled pool (self-training or self-supervised pretraining), spend any labelling budget actively, and keep a clean labelled test set that none of these methods touch.

✓ Recap and self-test

Check you can explain each

Tick a box only when you could explain the point to someone else



Four axes

supervision signal, data access, task relation, model type and output



Weak supervision has three forms

incomplete, inexact, inaccurate



Each paradigm bets on an assumption

clusters, shared features, task distribution, stable recent data



Use the decision key

labels? stream? pretrained model? expert on call?



Combine paradigms

one per resource you have: labels, unlabelled data, models, experts

Q1 Name the four axes of the taxonomy.

A Supervision signal, data access, task relation, model type (plus the output / objective).

Q2 Zhou's three kinds of weak supervision?

A Incomplete (few labels), inexact (coarse labels on groups), inaccurate (noisy labels).

Q3 Labels on bags, positive if any instance is positive: which paradigm?

A Multiple-instance learning.

Q4 Only positives labelled, the rest unlabelled: which paradigm?

A Positive-unlabelled (PU) learning.

Q5 Inductive vs transductive in one line?

A Inductive learns a rule for any future input; transductive only labels a given unlabelled set.

Q6 Online vs continual learning?

A Online: one stream, update per example; continual: a sequence of tasks where old ones must not be forgotten.

Q7 Which paradigm lets the model choose which items humans label?

A Active learning.

Q8 Few labels, a pretrained model of the domain and fewer than ~50 examples per class: first thing to try?

A Zero-shot or few-shot / in-context use of the pretrained model, then fine-tuning if it falls short.

Q9 Why can unlabelled data hurt a semi-supervised model?

A If classes do not form separated clusters, or the unlabelled data comes from another distribution, pseudo-labels reinforce errors.

