

How Big Models Teach Small Models, and Why It Looks Like School

Julien Laurenceau | Published 2026-08-06 | Updated 2026-10-22

<https://pepitedata.com/knowledge-distillation-teacher-student-models/>

PepiteData | LinkedIn: Julien Laurenceau | X: @Julien_Lau



Knowledge distillation trains another model from a teacher's behavior. It does not copy every capability.

Knowledge distillation trains a student model using signals from a teacher model. A smaller student can become useful on a defined task without reproducing the teacher's full capabilities. The practical questions are what signal you can access, how you verify the student, and whether the complete operating cost improves.

The school analogy is useful up to a point: examples, detailed feedback and practice can teach different things. A model does not understand a lesson like a person, and a convincing answer is not evidence that the lesson was learned. I would judge the result on held-out work, not on how closely it sounds like its teacher.

What does knowledge distillation actually transfer?

Knowledge distillation transfers behavior through a training objective, rather than moving the teacher's weights into a smaller container. The teacher can provide generated answers, probability distributions or internal representations. The student updates its own weights to learn from that signal. Distillation is often used for model compression, although the teacher and student need not always differ in size.

The [2024 survey by Xu et al.](#) describes a broad family of LLM distillation methods. Treat that family as a set of choices about supervision and training, not as a single recipe that makes any small model equivalent to a larger one.

Quantization reduces the numerical precision used to represent a model. Pruning removes selected components. Distillation trains a model from a teacher's behavior. These techniques can be combined; describing distillation as the opposite of compression confuses the mechanism with the goal.


Which teacher signals can you access?

The available teacher interface determines which distillation method is feasible. Generated text can support supervised training, probability matching requires suitable output scores, and internal feature matching requires access to selected activations. These signals are not interchangeable. A hosted interface that returns text does not automatically expose full token distributions or hidden states.

Training signal	Teacher access needed	What to check
Generated examples	Prompts and returned text	Answer correctness, coverage and data quality
Probability or logit matching	Compatible scores for candidate outputs or tokens	Score availability, vocabulary alignment and objective
Internal feature matching	Selected hidden representations	Layer alignment, dimensions and training access

In [Hinton, Vinyals and Dean's 2015 formulation](#), soft targets convey relative probabilities across classes. A higher softmax temperature makes that distribution less concentrated during training. This can expose relationships that a single winning label discards. The probabilities describe model output behavior; they are not a guarantee of calibrated confidence or correctness.

For language models, a sequence of teacher-generated tokens supplies a different training signal from the full next-token distribution. Both can be useful, but a dataset of answers does not contain every alternative token score. Similarly, an embedding returned by a service is an output; using it does not by itself mean you accessed an internal layer.



Choose a training signal the teacher actually exposes. Text access alone does not provide scores or hidden activations.

Why practice on the student's own attempts?

On-policy distillation trains on sequences sampled from the student, with supervision supplied by the teacher. Off-policy training uses sequences obtained elsewhere, such as teacher-generated demonstrations. Student-generated sequences expose states and mistakes the student actually encounters. The distinction concerns the source of training trajectories, not whether one method is universally better.

The [Thinking Machines on-policy distillation experiments](#) use teacher feedback at the token level on student trajectories. This targets a practical weakness of learning only from teacher demonstrations: after an early mistake, the student can enter contexts absent from those demonstrations.

Dense feedback requires the appropriate scoring interface and additional teacher computation. A text-only request asking whether an answer looks right is not automatically equivalent to that objective. I would compare methods under the same evaluation criteria and include teacher sampling and scoring in the training budget.

Can the student outperform its teacher?

Teacher performance is not a universal mathematical ceiling on student performance. A student's result depends on its starting model, supervision, additional data, optimization and evaluation task. A smaller specialist may perform well within a narrow scope while remaining less capable elsewhere. Matching a benchmark does not establish equivalent reliability on a production workload.

The [DeepSeek-R1 research](#) demonstrates that reasoning behavior from a large model can help train smaller models. That supports the usefulness of distillation; it does not establish that an arbitrary compact model can replace the teacher on every task.

A comparison also needs a clear baseline. Measure the student before distillation, the teacher and the trained student on the same untouched evaluation set. Otherwise, improvements from a strong starting model, additional training or an easier test can be misattributed to the teacher.

For a related example of why training-stage comparisons need careful scope, see [what the Qwen3 post-training cost comparison does and does not show](#).

What can transfer without being intended?

Teacher-generated training data can transmit unwanted behavior as well as useful task performance. Filtering visible references to an unwanted trait does not prove that all training signals associated with that trait are absent. This is a reason to evaluate behavior independently after training, rather than treating a clean-looking synthetic dataset as sufficient assurance.

In [Cloud et al., published in Nature in April 2026](#), students acquired behavioral traits through data without explicit references to those traits, including number sequences. The reported effect depended on the teacher and student sharing the same or behaviorally matched base models. It is a specific experimental finding, not proof that every teacher trait transfers to every student.

For a production evaluation, I would include unwanted behavior, refusal handling, out-of-scope requests and error patterns, alongside the intended task. If the teacher grades its own generated data, add checks that do not depend solely on that teacher's judgment.



A student can inherit an unwanted pattern along with the useful one. Check what transferred, not just whether training completed.

What should pass before a deployment?

A distilled model should pass an evaluation that represents its intended workload and the consequences of failure. Define acceptance criteria before selecting training examples or tuning checkpoints. Keep final evaluation data separate from generation and training. Include task quality, operational behavior and failure handling; a single average accuracy number cannot establish deployment readiness.

Gate	Evidence to collect	Decision it supports
Task quality	Held-out examples, error costs and difficult cases	Whether the student meets the required quality
Scope and safety	Unseen inputs, unwanted behavior and escalation cases	When to reject, abstain or use a fallback
Serving performance	Latency distribution, throughput and memory at target concurrency	Whether the chosen hardware meets the service target
Lifecycle cost	Generation, filtering, training, serving and refresh costs	Whether savings survive realistic operating assumptions

A [February 2026 SMS threat-detection preprint by ElZemity et al.](#) reports an automated distillation workflow using multiple teachers and compact students. Its best reported accuracy was 94.31%, with 96.25% recall. Those are results from that study, not latency measurements or evidence that the models are ready for an arbitrary phone or security workflow.

The table is my evaluation checklist, not a benchmark result. Where labels are uncertain, retain expert review and document disagreement. For high-cost mistakes, define an escalation path before optimizing the number of requests handled without one.

When does a smaller model save money?

Distillation saves money only when the student's lifecycle cost is lower for an acceptable level of service quality. Count data generation, rejected samples, filtering, training, evaluation, hosting and future refreshes. Compare those costs with the alternative at realistic traffic and concurrency. A smaller parameter count alone does not establish a cheaper service.

A useful planning calculation is: preparation and refresh cost, plus requests multiplied by measured serving cost, plus the operating cost over the same period. Include fallback requests to the teacher. Avoid counting hosting twice when a per-request estimate already includes it.

The opportunity is strongest when the task is stable, repeated often and testable. Low volume, frequent changes, broad open-ended work or expensive errors can favor keeping the stronger model in the request path. A hybrid system can route only suitable cases to the student.

I would make that decision within the broader [AI infrastructure and SRE operating model](#): capacity, observability, quality monitoring and rollback still matter after training. If you need to compare hosting options, my [capacity planning and platform sizing review](#) can test the assumptions before a hardware commitment.

FAQ: knowledge distillation in practice

Knowledge distillation is a training method, so its value depends on the teaching signal, the student's capacity and the evaluation workload. The questions below separate common deployment assumptions from what training actually guarantees. None of these choices removes the need to measure the resulting model on work it has not seen during training.

Is distillation the same as fine-tuning?

Distillation describes learning from a teacher's signals. Fine-tuning describes additional training of an existing model. Fine-tuning a student on teacher-generated examples can therefore be a form of distillation; not every fine-tuning dataset or objective uses a teacher.

Can I distill from a text-only interface?

Returned text can be used as training examples when the applicable terms and data permissions allow it. Text-only access does not provide the full output distribution or internal activations needed by other methods. Check the interface and the planned objective before building the dataset.

Does a distilled model need the teacher at inference?

A trained student can run independently of its teacher. A deployed application may still call the teacher for fallback cases, evaluation or new training data. Include those calls in the cost and dependency model.

Does a smaller model always need less hardware?

Fewer parameters can reduce weight storage, but deployment requirements also depend on precision, context length, cache memory, batch size, concurrency and the runtime. Measure memory and latency with the intended workload before choosing hardware.

Sources and scope

The sources support the training mechanisms and the specific research results discussed here. The evaluation checklist and cost decision framework are my engineering recommendations, not measurements from those papers. The cited experiments do not establish performance, safety or savings for a customer's deployment; those require workload-specific testing.

- [Hinton, Vinyals and Dean, Distilling the Knowledge in a Neural Network, 2015.](#)
- [Xu et al., A Survey on Knowledge Distillation of Large Language Models, 2024.](#)
- [Thinking Machines, On-Policy Distillation, research experiments and method.](#)
- [DeepSeek-AI, DeepSeek-R1, 2025, revised 2026.](#)
- [Cloud et al., Language models transmit behavioural traits through hidden signals in data, Nature, April 2026.](#)
- [ElZemity et al., Agentic Knowledge Distillation, preprint, February 2026.](#)

No affiliation with or endorsement by the cited research teams is implied.

Updated October 22, 2026. Revised the distinctions between teaching signals, qualified the teacher-ceiling and cost claims, and added deployment checks. The featured composition is retained with brighter lighting; the supporting illustrations are new. No customer data or customer-specific results are used.

Related posts

This article belongs to the AI infrastructure and SRE cluster and focuses on how teacher signals train a student model. The parent guide covers running models in production. The related articles below cover a local-model distillation workflow and the interpretation of training-cost comparisons, rather than duplicating the teaching mechanisms explained here.

- [AI infrastructure and SRE: running models in production](#)
- [Distilling frontier reasoning into a local model](#)
- [Interpreting frontier post-training cost comparisons](#)

PDF snapshot of the public article. The online version is authoritative; links remain clickable.