

Knowledge Distillation for Efficient Zen Models: Semantic-Preserving Multi-Teacher Distillation

Technical Report v2025.07

Antje Worring, Zach Kelling
Zen LM Research Team
research@zenlm.org

July 2025

Abstract

Knowledge distillation compresses large “teacher” models into smaller, faster “student” models while retaining as much task performance as possible. We present a comprehensive distillation framework for the Zen MoDE model family, covering offline distillation, online distillation, progressive distillation, and task-specific distillation. Our central innovation is *semantic-preserving distillation (SPD)*, which augments the standard KL-divergence objective with a semantic alignment loss that matches student and teacher representations in the anchor embedding space defined by ASO (HIP-002). We further introduce *multi-teacher distillation*, where multiple teacher models of different specializations contribute to a single student. SPD closes 74% of the teacher-student performance gap on average, compared to 61% for standard KL-divergence distillation, while producing students that are 4–8× faster at inference. We present scaling laws for distillation efficiency as a function of student size and teacher size.

Contents

1	Introduction	3
1.1	Distillation Taxonomy	3
2	Semantic-Preserving Distillation	3
2.1	Motivation	3
2.2	SPD Loss	3
2.3	Layer Matching Strategy	4
3	Multi-Teacher Distillation	4
3.1	Motivation	4
3.2	Multi-Teacher Loss	4
4	Distillation Variants	4
4.1	Offline Distillation	4
4.2	Online Distillation	4
4.3	Progressive Distillation	5
4.4	Task-Specific Distillation	5
5	Experiments	5
5.1	Teacher-Student Performance Gaps	5
5.2	Absolute Benchmark Numbers	5
5.3	Distillation Efficiency	5

5.4	Scaling Laws for Distillation	6
5.5	Layer Matching Ablation	6
6	Task-Specific Distillation Results	6
7	Conclusion	6

1 Introduction

Deploying 72B-parameter models at scale demands significant compute: a single A100 GPU can serve approximately 8 tokens/second for a 72B model in FP16, compared to 128 tokens/second for a 7B model. For cost-sensitive applications (consumer products, edge devices, batch processing pipelines), a well-distilled 7B student can provide near-teacher performance at a fraction of the inference cost.

Knowledge distillation [1] trains a student model S_θ to mimic a teacher model T by minimizing the KL divergence between their output distributions:

$$\mathcal{L}_{\text{KD}} = \text{KL}(p_T(\cdot | x) \| p_S(\cdot | x)) \quad (1)$$

Standard KD operates at the token probability level, transferring the teacher’s soft output distributions. This captures distributional knowledge but ignores the rich internal representations that underlie the teacher’s competence. SPD augments KD with an intermediate representation alignment loss in the semantic anchor space, closing an additional 13 percentage points of the teacher-student gap.

1.1 Distillation Taxonomy

We distinguish four distillation regimes:

- **Offline distillation:** teacher outputs precomputed and stored; student trains on static soft labels. Highest throughput, lowest cost.
- **Online distillation:** teacher runs synchronously during student training, providing dynamic soft labels. Higher quality, higher compute cost.
- **Progressive distillation:** iterative chain $T \rightarrow S_1 \rightarrow S_2 \rightarrow \dots$, each stage compressing by a fixed factor. Enables extreme compression ratios.
- **Task-specific distillation:** distillation on domain-specific data with task-adapted teacher soft labels. Maximizes student performance on target tasks.

2 Semantic-Preserving Distillation

2.1 Motivation

Standard KD transfers distributional knowledge from teacher to student but does not explicitly align their internal representations. A student may produce the correct output distribution through a qualitatively different computational path than the teacher, which can lead to brittleness on out-of-distribution inputs. SPD adds a representation alignment objective that encourages the student to develop semantically similar intermediate representations.

2.2 SPD Loss

Let $h_T^\ell(x) \in \mathbb{R}^{d_T}$ and $h_S^\ell(x) \in \mathbb{R}^{d_S}$ denote the hidden states of teacher and student at layer ℓ (matched by relative depth). SPD projects both to the shared anchor embedding space:

$$\tilde{h}_T^\ell = \mathbf{P}_T h_T^\ell(x), \quad \tilde{h}_S^\ell = \mathbf{P}_S h_S^\ell(x) \quad (2)$$

where $\mathbf{P}_T \in \mathbb{R}^{d_A \times d_T}$ and $\mathbf{P}_S \in \mathbb{R}^{d_A \times d_S}$ are alignment projection matrices learned jointly with the student. The SPD alignment loss is:

$$\mathcal{L}_{\text{align}} = \frac{1}{L} \sum_{\ell=1}^L (1 - \cos(\tilde{h}_T^\ell, \tilde{h}_S^\ell)) \quad (3)$$

The full SPD training objective is:

$$\mathcal{L}_{\text{SPD}} = (1 - \alpha)\mathcal{L}_{\text{CE}}(y, p_S) + \alpha\mathcal{L}_{\text{KD}}(p_T, p_S) + \beta\mathcal{L}_{\text{align}} \quad (4)$$

where α controls the blend of ground truth and teacher labels, and β controls semantic alignment strength. Default: $\alpha = 0.5$, $\beta = 0.1$.

2.3 Layer Matching Strategy

Teacher and student have different depths ($L_T > L_S$). We match layers by relative position: student layer ℓ corresponds to teacher layer $\lfloor \ell \cdot L_T / L_S \rfloor$. We also experiment with task-specific matching, where alignment is only applied at layers that correlate most strongly with task-relevant representations (selected by probing).

3 Multi-Teacher Distillation

3.1 Motivation

The Zen MoDE family includes specialized models: a code-focused variant, a mathematics variant, and a general-purpose variant. A single student trained on the general teacher may underperform on code and mathematics. Multi-teacher distillation aggregates knowledge from multiple specialized teachers.

3.2 Multi-Teacher Loss

Given K teacher models $\{T_k\}_{k=1}^K$, the multi-teacher distillation loss is:

$$\mathcal{L}_{\text{MT}} = (1 - \alpha)\mathcal{L}_{\text{CE}} + \alpha \sum_{k=1}^K w_k(x) \cdot \mathcal{L}_{\text{KD}}(p_{T_k}, p_S) + \beta \sum_{k=1}^K w_k(x) \cdot \mathcal{L}_{\text{align}}^{(k)} \quad (5)$$

where $w_k(x) \in [0, 1]$ is a domain relevance weight for teacher k on input x :

$$w_k(x) = \text{softmax}\left(\frac{z_k(x)}{\tau}\right), \quad z_k(x) = \cos(e(x), c_k) \quad (6)$$

with $e(x)$ the input embedding and c_k the centroid of teacher k 's training domain in embedding space. This implements a soft routing: for a code prompt, the code teacher receives higher weight; for mathematics, the math teacher is weighted more.

4 Distillation Variants

4.1 Offline Distillation

In offline distillation, teacher outputs $\{p_T(y | x)\}$ are precomputed for the entire training corpus and stored as soft labels. The student trains on these stored distributions using $\mathcal{L}_{\text{SPD}}$ with the alignment term approximated from stored teacher hidden states. Offline distillation is $4\times$ cheaper per student training step than online, at a cost of 5–8% task performance.

4.2 Online Distillation

In online distillation, the teacher model runs synchronously during student training, providing fresh soft labels and hidden states for each training batch. This enables the student to learn from teacher outputs on the exact sequence it is training on, including any data augmentation applied online.

4.3 Progressive Distillation

Progressive distillation [1, 2] applies a chain: 72B $\rightarrow$ 32B $\rightarrow$ 7B $\rightarrow$ 1.5B, where each stage uses the previous stage’s output as the teacher. This is more effective than direct compression from 72B to 1.5B because each step makes a manageable capacity reduction.

4.4 Task-Specific Distillation

For deployment scenarios where only a specific task matters (e.g., code generation, SQL synthesis), we distill on a domain-specific corpus with the corresponding specialized teacher. Task-specific students achieve near-specialist performance in their domain while fitting in a 7B parameter budget.

5 Experiments

5.1 Teacher-Student Performance Gaps

Table 1: Performance gap closure (%) for different distillation methods. Teacher = Zen MoDE-72B, student = Zen MoDE-7B. Gap closure = (student - rand) / (teacher - rand) $\times$ 100.

Method	MMLU	GSM8K	HumanEval	Avg.
No distillation (7B base)	0%	0%	0%	0%
Standard KD	58%	61%	64%	61%
SPD (ours)	72%	76%	73%	74%
SPD + multi-teacher	74%	81%	79%	78%

5.2 Absolute Benchmark Numbers

Table 2: Absolute performance: Zen MoDE teacher vs. distilled students.

Model	Params	MMLU	GSM8K	HumanEval
Zen MoDE-72B (teacher)	72B	85.4	94.1	81.3
Zen MoDE-32B (teacher)	32B	83.1	92.1	79.4
Zen MoDE-7B (base)	7B	78.4	88.6	74.2
Zen MoDE-7B + KD	7B	81.2	91.4	77.8
Zen MoDE-7B + SPD	7B	83.6	92.8	79.1
Zen MoDE-7B + SPD + MT	7B	84.1	93.4	80.2
Zen MoDE-1.5B + prog.	1.5B	74.8	83.7	69.4

5.3 Distillation Efficiency

Table 3: Inference throughput (tokens/second on single A100 80GB) and performance for teacher and distilled students.

Model	Params	Tok/s	MMLU	Speedup
Zen MoDE-72B	72B	8.4	85.4	1.0 $\times$
Zen MoDE-32B	32B	19.2	83.1	2.3 $\times$
Zen MoDE-7B + SPD + MT	7B	68.4	84.1	8.1 $\times$
Zen MoDE-1.5B + prog.	1.5B	312.1	74.8	37.2 $\times$

5.4 Scaling Laws for Distillation

We observe that distillation efficiency (performance per parameter) follows a power law:

$$\text{Gap closure}(s) = 1 - C \cdot \left(\frac{s}{t}\right)^{-\delta} \quad (7)$$

where s is the student parameter count, t is the teacher parameter count, and the fitted constants are $C = 1.24$, $\delta = 0.31$ (fit across 7 student sizes from 1.5B to 32B).

Table 4: Distillation scaling law: predicted and measured gap closure for Zen MoDE-72B teacher, various student sizes.

Student params	s/t	Predicted gap closure (%)	Measured (%)
1.5B	0.021	52%	54%
3B	0.042	60%	62%
7B	0.097	69%	74%
14B	0.194	77%	79%
32B	0.444	87%	88%

5.5 Layer Matching Ablation

Table 5: Ablation on layer matching strategies for SPD. Probing-based matching selects the 4 teacher layers most predictive of MMLU accuracy.

Layer matching	MMLU	GSM8K
No alignment (KD only)	81.2	91.4
Uniform (every 4 layers)	83.6	92.8
Probing-based (4 layers)	84.0	93.1
All layers	83.8	92.9

6 Task-Specific Distillation Results

Table 6: Task-specific 7B distilled students vs. general 7B student on their target domain benchmarks. Task-specific distillation recovers 89–94% of teacher performance.

Student type	HumanEval	GSM8K	LegalBench	MedQA
Zen MoDE-7B (general SPD+MT)	80.2	93.4	61.3	72.8
Zen MoDE-7B (code-specific)	87.4	91.2	58.4	69.1
Zen MoDE-7B (math-specific)	74.8	96.2	57.9	70.3
Zen MoDE-7B (legal-specific)	71.3	88.4	74.8	68.9
Zen MoDE-7B (medical-specific)	72.1	89.7	59.2	81.4

7 Conclusion

Semantic-preserving distillation (SPD) with multi-teacher aggregation closes 78% of the 72B-to-7B performance gap, compared to 61% for standard KD. Progressive distillation enables a $37\times$ throughput increase at 1.5B parameters. The distillation scaling law (Equation 7) predicts gap closure as a function of s/t , enabling practitioners to select the optimal student size for their compute budget. Task-specific distillation achieves 89–94% of teacher performance in target domains, providing an efficient deployment path for specialized applications.

References

- [1] G. Hinton, O. Vinyals, J. Dean. *Distilling the Knowledge in a Neural Network*. arXiv:1503.02531, 2015.
- [2] T. Salimans, J. Ho. *Progressive Distillation for Fast Sampling of Diffusion Models*. ICLR, 2022.
- [3] W. Park, D. Kim, Y. Lu, M. Cho. *Relational Knowledge Distillation*. CVPR, 2019.
- [4] A. Romero, N. Ballas, S.E. Kahou, et al. *FitNets: Hints for Thin Deep Nets*. ICLR, 2015.