

ELDR: Expert-Locality-Aware Decode Routing for PD-Disaggregated MoE Serving

Sangjin Choi*
KAIST
Daejeon, Korea

Sukmin Cho*
KAIST
Daejeon, Korea

Yifan Xiong
Microsoft Research
Beijing, China

Ziyue Yang†
Shanghai Xingyunzhili
Artificial Intelligence Institute
Shanghai, China

Youngjin Kwon‡[✉]
KAIST
Daejeon, Korea

Peng Cheng
Microsoft Research
Redmond, WA, USA

Abstract

In prefill–decode (PD) disaggregated LLM serving, each request is assigned to a decode worker after prefill. Existing decode routers balance only load; for mixture-of-experts (MoE) models this is incomplete: equally loaded workers can differ in latency, since each decode step loads the weights of every distinct expert its batch activates. We present ELDR, an expert-locality-aware decode router for PD-disaggregated MoE serving. From a request’s prefill expert activations, ELDR builds an expert signature predicting the experts it will activate during generation. Offline, balanced K -means partitions signature space across decode workers; online, locality-band routing sends each request to the least-loaded worker among those best matching its signature. A signature cache, co-indexed with the KV cache at KV-block granularity, keeps signatures exact under prefix caching. Implemented in vLLM and evaluated on deployments of up to 40 GPUs, ELDR reduces median TPOT by 5.9–13.9% over the strongest of four load-balancing baselines across three MoE models and two workloads, with model outputs unchanged.

1 Introduction

Large Language Model (LLM) serving is moving toward Prefill-Decode (PD) disaggregation, which runs prompt processing (prefill) and token generation (decode) on separate worker pools [39]. The phases mismatch: prefill runs the prompt in parallel and is compute-bound, while decode is sequential and latency-sensitive, so colocating them lets long prefills stall decode and hurts both Time-to-First-Token (TTFT) and Time-per-Output-Token (TPOT). An $xPyD$ deployment instead provisions x prefill and y decode workers separately [7, 25, 30, 32], turning routing into the central serving decision: each request is assigned a prefill worker for its prompt and, once its KV state is materialized, a decode worker for its generation. Prefill-side routing has drawn most of the attention, where cache-affinity policies reuse KV state across prompts. Decode-side routing has not: existing

policies balance load and otherwise treat decode workers as interchangeable. For dense models they are, since equal-load workers do equal feed-forward work. For mixture-of-experts (MoE) models they are not.

For MoE models, equal load does not mean equal latency. Decode is memory-bandwidth bound, and at batched decode its cost is set by the *union* of distinct experts the batch loads from HBM each step. Here sparsity inverts. The property that makes MoE cheap, routing each token to only a few experts [9, 28, 35], fragments the decode batch across experts and destroys the weight reuse a dense batch enjoys, where one weight load amortizes over every token. A step pays for every expert any of its tokens selects, so the union, not the token count, governs latency. The gap is large: on Qwen3-30B-A3B, growing the active-expert count from 16 to 128 raises MoE-layer latency $4.7\times$ at fixed batch size, while batch size at fixed active-expert count barely moves it (§3). Because the union depends on which requests share a worker, expert composition is a first-order latency knob that load-based routing cannot see—a second routing axis: *expert locality*.

Expert locality is exploitable because it is structured. The MoE gate¹ picks each token’s experts from its hidden representation, so requests from the same domain activate overlapping experts: across tasks, code, math, medical, and legal prompts exercise distinct expert regions, and multilingual traffic separates by language. Expert choice is thus correlated across related requests, not just sparse within a token. Concretely, we observe same-domain decode batches activate 17–21% fewer distinct experts per step than mixed batches on task workloads (§3), so a router that colocates similar requests shrinks each worker’s per-step union while a load-only router scatters them.

Structure helps only if the router can see it in time. Placement happens at the prefill-to-decode handoff, before any output token, so decode-time expert choices are not yet observable. But prefill has already pushed the prompt through the same gates that will route decode, and the two agree: per-expert prefill and decode activation correlate at 0.70 to

*Work done during an internship at Microsoft Research.

†Work partially done while the author was at Microsoft Research.

‡Corresponding author.

¹We call the MoE router the MoE gate to avoid confusion with the PD router.

0.92 across three MoE models (§3). Expert locality is therefore not only structured but visible at exactly the moment the router must act.

We present ELDR, an expert-locality-aware decode router for PD-disaggregated MoE serving. ELDR turns each request’s prefill activations into a compact *expert signature* whose distances predict decode-time expert overlap, so requests with nearby signatures share experts at decode. It selects the signature representation by a single criterion, how faithfully signature proximity predicts that overlap, decoupled from any downstream clustering or routing.

Routing must then satisfy two objectives that need different information. Locality, which experts colocated requests share, is an aggregate property visible only across many requests; load, which decode’s non-MoE work imposes per colocated request, is instantaneous and visible only at the decision. Pure locality overloads popular domains; pure load balancing scatters expert-similar requests. ELDR splits the decision along that information boundary. Offline, *balanced K-means* partitions signature space into one locality region per decode worker, capturing structure without inheriting workload skew. Online, *locality-band routing* keeps the workers whose centroids fall within a similarity band of the request’s best match and routes to the least loaded among them: the band enforces locality, the load choice inside it absorbs live skew, and one decision serves both.

One case stresses the online path: prefix caching. A cache hit skips prefill for the shared prefix, so the gate never runs there and the signature is incomplete—worst exactly for the cache-hit requests that prefix-aware routing concentrates. ELDR therefore keeps an expert-signature cache co-indexed with the KV cache at block granularity: each block carries its tokens’ expert footprint, and summing cached and freshly computed blocks recovers the full signature—coherent across partial hits, evictions, and reuse, with no recomputation and no change to caching.

We implement ELDR in vLLM [16] as a thin layer over an existing PD-disaggregated stack: prefill-time signature capture, an offline fit, locality-band routing, and the block-granular signature cache. ELDR changes only which decode worker serves a request, leaving the model, gate decisions, kernels, and batching untouched, so expert selections—and hence outputs—are identical to standard top- k gating. Across three MoE models (Qwen3-30B-A3B [35], GPT-OSS-120B [28], Gemma-4-26B-A4B [9]) on task and language workloads, ELDR cuts median TPOT by 7.0–13.9% (task) and 5.9–10.0% (language) over the best load-balancing baseline with a signature cache under 1% of the KV cache. ELDR further generalizes to a 235B deployment under expert parallelism (§6.4).

ELDR makes three contributions:

- **Expert locality as a predictable routing axis.** Decode latency in PD-disaggregated MoE serving turns on the distinct experts colocated requests activate; this locality is

structured by domain; and, the enabler, it is readable from prefill before decode begins.

- **Expert-locality-aware decode routing.** ELDR builds a quality-selected expert signature, partitions it across decoders with balanced K -means, and routes with a locality band that balances expert locality against live load.
- **Prefix-cache-coherent signatures.** A block-granular signature cache co-indexed with the KV cache keeps signatures exact across partial hits, full hits, and evictions, so ELDR composes with prefix caching at negligible cost.

2 Background

2.1 Mixture-of-Experts Large Language Models

Mixture-of-Experts (MoE) has become a primary architecture for scaling frontier Large Language Models (LLMs) because it decouples model capacity from per-token computation. Instead of applying the same dense feed-forward network (FFN) to every token, an MoE layer replaces the FFN with multiple expert FFNs and an MoE gate that activates a small subset of experts for each token. As a result, the model can contain many more parameters than a dense model while using only a fraction of them during each forward pass. This sparse activation makes MoE attractive for efficient scaling: it increases model capacity and enables richer expert specialization without proportionally increasing per-token FLOPs. Recent architectures push this idea further with fine-grained experts, splitting FFN capacity across more, smaller experts so that each token can be routed to a more specialized combination of experts while keeping per-token FLOPs bounded.

2.2 Prefill-Decode Disaggregated Serving

Prefill-Decode (PD) disaggregation [39] is a practical architecture for latency-sensitive LLM serving. Inference has two phases: *prefill* processes the prompt in parallel and is compute-bound, while *decode* generates tokens autoregressively and is latency-sensitive and memory-bandwidth-bound. Colocating them causes phase interference—long prefills stall latency-sensitive decode steps—and forces one resource allocation onto two mismatched workloads. PD disaggregation separates the phases onto independent pools, typically an $xPyD$ configuration of x prefill and y decode workers, scaling capacity per phase.

This makes routing a central serving decision. A *PD router* assigns each request a prefill worker for the prompt and a decode worker for generation, determining where its KV state transfers and which worker serves its tokens. Existing policies optimize system-level objectives—cache-aware routers prefer prefillers holding reusable KV blocks [30], load-balancing policies (round-robin, join-shortest-queue, power-of-two-choices) spread requests—but decide only *where* a request runs, agnostic to the model-internal expert activations that set MoE decode’s memory-bandwidth cost. This

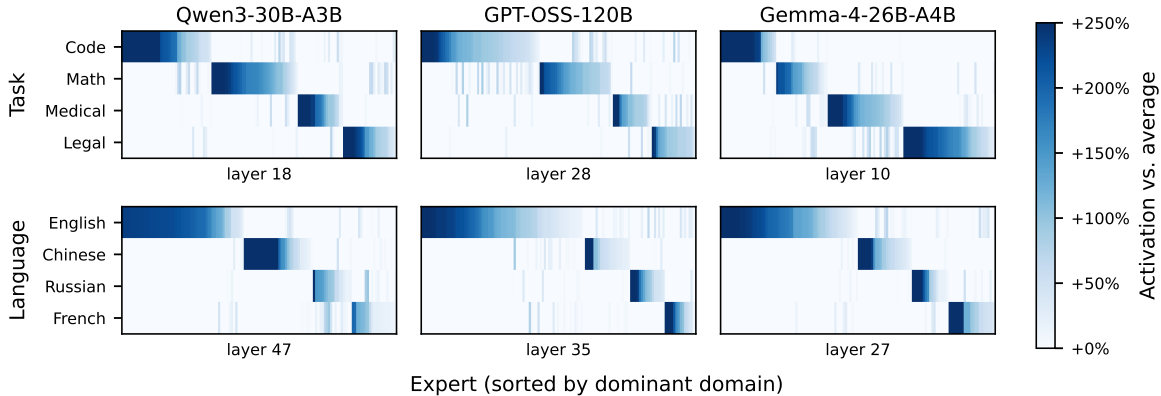


Figure 1. Decode-phase per-expert activation relative to the cross-domain mean, for three MoE models along task (top) and language (bottom, WildChat [37]) domains at each model’s most discriminative layer. Darker is above-average (below-average clipped to white); experts are reordered per panel into contiguous per-domain blocks. Each domain over-activates a distinct subset of experts. Task domains: Code [1, 4, 17, 40], Math [5, 11, 13, 19, 20], Medical [12, 14, 15], and Legal [3].




Figure 2. MoE layer latency scales with active experts, not batch size (single MoE layer, one MI300X).

leaves the activated-expert cost each request imposes on its decode worker unmodeled.

3 Motivation

3.1 Active Expert Count Drives MoE Decode Latency

MoE decode latency is governed by the number of distinct experts activated at each decode step. Sparsity reduces computation but amplifies decode’s memory-bandwidth bottleneck by fragmenting the batch across experts. In a dense model, every token in a decode batch executes the same FFN, so one weight load amortizes over the whole batch. In an MoE model, the gate partitions the batch among experts, and weight reuse occurs only among tokens sharing an expert; as experts grow more numerous and finer-grained, fragmentation lowers arithmetic intensity further. MoE decode cost is therefore dominated by expert weight access: each step fetches the weights of every expert the batch activates.

To isolate this effect, we benchmark the single-layer MoE expert computation on one MI300X, sweeping decode batch size and the number of distinct active experts for Qwen3-30B-A3B, GPT-OSS-120B, and Gemma-4-26B-A4B. Across all models, latency tracks active experts far more strongly

than batch size (Fig. 2): on Qwen3-30B-A3B, growing active experts from 16 to 128 raises latency 4.7× at batch size 64, while batch size at a fixed expert count barely moves it.

This result shows that MoE decode efficiency depends on expert reuse within a batch. However, exploiting this requires the serving system to know which requests are likely to activate overlapping experts. We next show that this overlap is not random. Instead, expert usage is structured by request domain, making expert locality a predictable property that a serving system can exploit.

3.2 Experts Specialize by Domain

The MoE gate is input-dependent by design. The gate scores experts using the token’s hidden representation and dispatches the token to the top-scoring experts. Because the hidden representation encodes contextual features, domains such as task and language are natural axes along which expert usage may differ.

Fig. 1 shows decode-phase expert activation across three models: Qwen3-30B-A3B, GPT-OSS-120B, and Gemma-4-26B-A4B. The figure evaluates two sources of request structure: task and language. The task axis uses a mixture of code, math, medical, and legal benchmarks. The language axis uses multilingual WildChat requests, a corpus of one million real-world user–ChatGPT interactions spanning 68 languages such as English, Chinese, Russian, and French. In each heatmap, rows correspond to domains and columns correspond to experts, with color indicating activation relative to the average for that expert across domains. Experts are sorted by their dominant domain, so experts most activated by the same domain appear contiguously.

Across models, expert activation is strongly structured by both task and language. Code, math, medical, and legal requests over-activate different expert subsets. The same pattern appears on the multilingual axis: English, Chinese,




Figure 3. Prefill expert activation predicts decode activation. Each point is one expert (normalized prefill x vs. decode y , pooled over domains); points near the diagonal are experts used about equally in both phases.

Russian, and French requests activate different expert regions rather than uniformly exercising the full expert pool. This shows that expert selection is correlated across related requests, not just sparse at the individual-token level.

This specialization makes active-expert reduction actionable for serving: same-domain requests are more likely to share experts at decode time, so co-locating them on a worker yields a smaller per-step active-expert set than mixed batches. The active-expert count therefore depends on which requests are placed together, not only on how much work each worker receives. This structure is a property of the model’s gating networks—stable across requests of the same domain and observable wherever the gating networks are exercised. To exploit it for placement, a serving system must identify each request’s expert footprint at the moment of prefill→decode handoff, before any decode token has been generated.

3.3 Prefill Predicts Decode Expert Activation

Expert locality can guide a decode routing policy only if it is visible before decode begins. In an $xPyD$ deployment, a request first runs on a prefill worker and is then assigned to one of the decode workers, where it will be batched with other active requests. The decode routing policy must make this placement decision before any output tokens are generated, so the request’s decode-phase expert usage is not yet observable. The prefill phase, however, has already processed the prompt through the same MoE layers, exposing the expert choices made by the model’s gating networks.

Fig. 3 compares prefill and decode expert activation across the task and language domains studied above. Each point corresponds to one expert: the x -axis is its normalized activation during prefill, and the y -axis is its normalized activation during decode. Points near the diagonal indicate that experts heavily used while reading the prompt are also heavily used while generating the response. The correlation is strong for Qwen3-30B-A3B and Gemma-4-26B-A4B, and substantial for GPT-OSS-120B.

This correlation makes decode expert locality available before decode begins. Although the exact decode tokens are unknown at placement time, the prefill phase has already exposed a request-specific expert footprint that predicts the

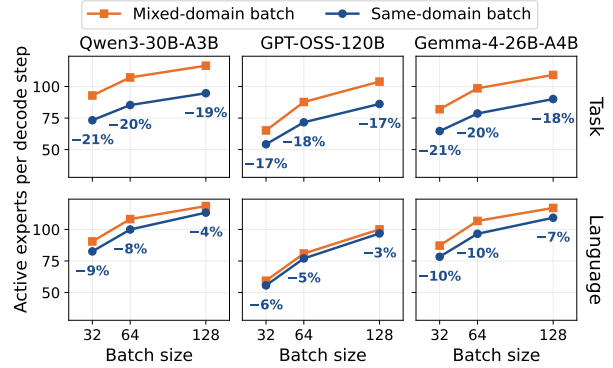


Figure 4. Same-domain batches (blue) activate fewer experts per decode step than mixed-domain (orange), for task (top) and language (bottom) across three MoE models.

experts likely to be reused during generation. Thus, in an $xPyD$ deployment, the system observes an expert-locality signal precisely at the boundary where the request leaves the prefill worker and must be assigned to a decode worker.

3.4 Opportunity: Expert-Locality-Aware Decode Routing

In PD-disaggregated MoE serving, decode-step latency is governed not only by per-worker load but by request composition: requests that dispatch to overlapping experts activate fewer distinct experts per step than an equally sized batch of unrelated requests. Fig. 4 quantifies this: same-domain batches activate 17–21% (task) and 3–10% (language) fewer experts per step than mixed-domain batches across batch sizes 32–128.

Expert locality offers a second routing axis for PD-disaggregated MoE serving. Existing decode routing policies balance per-worker load but ignore expert overlap. Placing expert-similar requests on the same decoder shrinks the per-step active-expert set; placing dissimilar requests together expands it. The signal that enables overlap-aware routing is observable at the prefill → decode handoff: prefill expert activations predict decode-time expert usage. We call the per-request representation of this signal the *expert signature*.

3.5 Challenges

Turning this opportunity into a working routing policy raises three challenges.

3.5.1 Designing the expert signature. The expert signature predicts a request’s decode-time expert usage from its prefill activations, so that requests with similar predicted usage can be colocated to share experts. It is the building block of ELDR’s design: every routing decision derives from it, and routing quality is bounded above by signature quality. The challenge is to turn raw prefill data into a vector space suited




Figure 5. WildChat [37] request volume is heavily skewed: English and Chinese alone are ~75% of requests.

to clustering—one in which proximity between two signatures reflects how much their requests overlap in decode-time expert activation. Only such a space lets a clustering group requests by genuine expert affinity rather than by spurious similarity. The design space is large: raw activation counts, gate logits, layer masks, and reweighting schemes all yield candidate representations, each inducing a different geometry. Selecting among them requires a criterion that measures how faithfully signature proximity predicts decode-time expert overlap, independent of any particular clustering or routing policy built on top.

3.5.2 Reconciling locality and load. Locality alone is not enough. Non-MoE compute—attention, QKV projection, normalization—scales with the number of active requests on a worker, not with how many distinct experts they share. A locality-only policy concentrates traffic on the workers whose profiles match the dominant request domains; in WildChat, the top two languages alone account for ~75% of requests (Fig. 5). Overloaded workers see inflated non-MoE work and rising tail TPOT. The opposite extreme—a balance-only policy like round-robin or pure JSQ—keeps queues even but scatters expert-similar requests, erasing the locality benefit. Both objectives must hold at once.

Combining the two is hard because they draw on different information. Locality is purely *aggregate*: which signatures group together is visible only across many requests, never from one. Balancing load needs that aggregate view too—structural skew, where some clusters carry far more traffic than others, lets the router anticipate the distribution instead of reacting request by request—but also an *instantaneous* view: runtime variance from Poisson arrivals, variable output lengths, and prior routing decisions that no aggregate can predict, so the router can track each worker’s live load at the moment of decision. No single mechanism provides both: aggregate-only is blind to live load, instantaneous-only discards the cross-request structure. The challenge is to bring both views into the routing pipeline.

3.5.3 Maintaining signature coherence across prefix cache hits. The expert signature is produced during prefill, but prefill is often partially or fully skipped. With a prefix cache hit, the system reuses cached KV states for the shared prompt prefix without re-running the gating networks on the cached tokens. The expert footprint for those tokens is

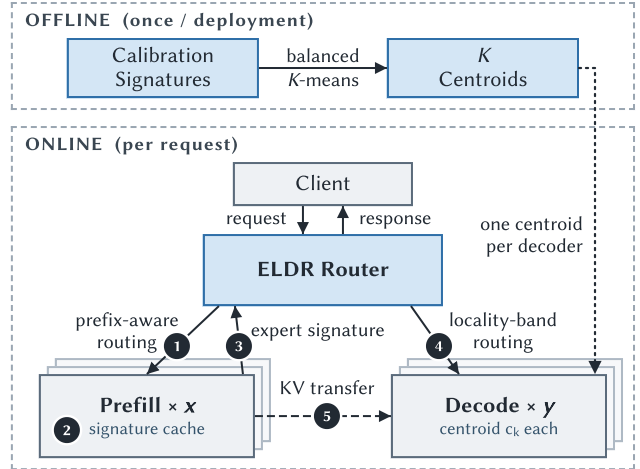


Figure 6. ELDR architecture: offline fitting of one centroid per decode worker over expert signatures, then online routing at the prefill–decode handoff by signature similarity, subject to load.

therefore absent from this request’s prefill—even though it was produced (and discarded) when an earlier request first populated the cache. Without a mechanism to recover this footprint, the signature for a cache-hit request is incomplete and misroutes the request.

A natural answer is to cache the per-request signature. This works on full prompt matches but breaks on *partial hits*: a request can share its leading KV blocks with a previous one and diverge on the rest, and no whole-prompt signature describes this kind of overlap. KV blocks are also dynamically evicted, and recycled across requests, so any signature mechanism must remain coherent under these transitions. The challenge is to keep the signature coherent with the cache across these states—partial hits, evictions, and reuse—without disabling the cache or re-running prefills.

4 Design

We present ELDR, an Expert-Locality-aware Decode Routing policy for PD-disaggregated MoE serving. Instead of placing decode requests by load alone, ELDR uses each request’s prefill expert activation to choose a decode worker where it is likely to share experts with other requests. This makes expert locality a practical routing signal for reducing TPOT.

4.1 Overview

Fig. 6 shows the ELDR architecture. ELDR integrates into an existing PD-disaggregated serving stack as a thin layer: a request router in front of the prefill–decode workers and a prefill-time hook that records per-block (1) *expert signatures* alongside the KV cache, with offline routing state loaded once at startup. The model, kernels, batching, and the prefill/decode engines are left unmodified. We implement and

evaluate ELDR on vLLM [16], and the design ports directly to other stacks such as SGLang [38]. ELDR runs in two stages: an (2) *offline stage* that groups requests by their prefill-time signatures to reduce expert activation, and an (3) *online stage* that uses this profile to route each request at the prefill-decode handoff.

Expert Signature. ELDR builds on the observation that a request’s prefill expert activations are highly correlated with its decode expert activations (§3.3). To exploit this, we design a per-request representation, the *expert signature* (§4.2), that summarizes a request’s prefill-time expert activations and predicts which experts it will use at decode time. The signature lets us estimate how much two requests’ decode-time expert usage overlaps directly from prefill information, via the cosine similarity between their signatures. ELDR therefore groups requests by signature similarity, so that the distinct experts activated across a group are fewer than under naive, similarity-agnostic routing.

Offline. Before serving, ELDR runs a small calibration set through prefill and collects each request’s expert signature. ELDR then clusters these signatures (§4.3.1) along two axes: locality, so signatures grouped on one worker share experts, and balance, so clusters are evenly sized across decode workers (one centroid per worker). Re-fitting is cheap: the balanced K -means clustering over the captured signatures completes in under 10 s on CPU and updates only the router’s centroid table, leaving the model and decode workers running. A PD scale-up or scale-down, or a shifted workload, then needs only this re-fit at the new decode-worker count, reassigning one centroid per worker. We use a single offline fit in this work.

Online. During serving, ELDR leaves the standard PD pipeline intact and adds two lightweight steps: signature capture during prefill and a routing decision at the prefill-decode handoff:

1. The router sends the request to a prefill worker using prefix-aware routing.
2. The prefill worker runs prefill while ELDR captures the request’s *expert signature* (§4.2) in a signature cache co-indexed with the KV cache (§4.4).
3. At handoff, the request’s expert signature—summed over its expert signature cache blocks—is returned to the router with the KV-transfer metadata.
4. The router invokes ELDR’s *locality-band routing* (§4.3.2): it compares the signature against the decoder centroids and routes to the least-loaded worker within the locality band of the closest centroid.
5. The selected decoder pulls the KV cache and generates on the unchanged decode engine; co-locating similar-signature requests shrinks the per-step active-expert union, reducing decode latency.

4.2 Expert Signature

ELDR summarizes each request’s prefill-time expert activations into a compact per-request representation, the **expert signature**. The raw material is cheap to collect at prefill: a per-layer histogram of how many tokens each expert receives, available as a by-product of routing. The signature is designed so that the distance between two signatures predicts how much the two requests overlap in their decode-time expert usage—requests with nearby signatures activate similar experts at decode, requests with distant signatures activate different ones. The clustering layer (§4.3.1) then groups requests whose signatures are close, so each decode worker loads few distinct experts per step.

Turning the raw histogram into a good signature still takes care. The per-layer counts are the right raw signal, but two effects keep raw count distance from tracking decode-time overlap as closely as it could: a few generalist experts fire on nearly every request, inflating the norm and masking the rare specialists that distinguish workloads, and many layers carry little signal that separates requests. The signature therefore weights the counts and selects which layers to keep, on top of what to count. We proceed in three steps: we make the design goal precise (§4.2.1)—a signature is good when its distances rank request pairs the same way their decode-time expert overlap does—then construct the signature from three principles justified against that goal (§4.2.2), and validate each choice empirically (§4.2.3).

4.2.1 Goal of Signature Design. The signature design aims to make decode-time expert overlap predictable from prefill alone, so clustering on signatures groups requests that load few distinct experts per step and thereby bounds decode latency. A signature meets this goal when requests with nearby signatures activate similar experts at decode time, and distant signatures activate different ones.

To quantify it, we encode each calibration request i ’s decode-time behavior as a per-step *activation probability* $p_i(\ell, e) \in [0, 1]$: the fraction of i ’s decode steps in which the gate at layer ℓ selected expert e . Flattened across all (ℓ, e) pairs and L2-normalized, $p_i \in \mathbb{R}^{LE}$ is the decode-time pattern the signature must predict. Distance in p -space is the natural quantification of decode-time expert overlap. Signature quality is then the rank correlation between signature pair-distance and decode-pattern pair-distance, over random calibration pairs:

$$\rho = \text{Spearman}(\cos\text{-dist}(s_i, s_j), \cos\text{-dist}(p_i, p_j)). \quad (1)$$

We use Spearman (rank) correlation rather than Pearson: signature distance lives in \mathbb{R}^d while decode-pattern distance lives in \mathbb{R}^{LE} , so their numerical values are not commensurable. Only the ordering of pairs is, and ordering is exactly what the clustering layer consumes when deciding which requests to colocate. Rank agreement also captures the goal property: high ρ means nearby signatures correspond to

similar decode-time expert activation, and distant signatures to different activation.

4.2.2 Signature Design. Three principles guide s_r :

- **Discrete, not continuous.** The target $p_i(\ell, e)$ is the normalized count of top- k selections at each (ℓ, e) over decode, so the signature should measure the same quantity. Only the discrete top- k loads experts—sub-threshold scores are computed but never fetch one—so a discrete prefill count is a sparse, same-form estimator of which experts fire. Continuous gate scores instead spread softmax mass over all E experts, assigning magnitude to experts the top- k never loads and smoothing away the contrast that distinguishes requests.
- **Downweight common experts.** A handful of generalist experts fire on almost every prefill. They carry little information about request identity, yet their large counts drown out the rare specialist experts that actually distinguish workloads. We multiply each (ℓ, e) count by its inverse document frequency (IDF) over the calibration corpus, which shrinks the common experts and amplifies the rare ones. This reweighting works at the granularity of (ℓ, e) cells, denoising the signature *before* we choose a layer mask.
- **Keep informative layers.** Layers contribute unequally: some carry routing signal that maps prefill to decode, others add dimensions that don’t separate requests and dilute the signature direction. The mask $\mathcal{S} \subseteq \{1, \dots, L\}$ is the subset we keep, chosen per deployment to maximize ρ (Eq. 1) on calibration data.

Four steps build s_r . The IDF table and layer mask are fit once offline on calibration data; counting, IDF reweighting, layer masking, and normalization run online at the prefill \rightarrow decode handoff.

1. *Count.* At each layer ℓ , count how many prefill tokens are routed to each expert, giving $c_r(\ell) \in \mathbb{N}^E$.
2. *Downweight common experts.* Multiply each cell by its inverse document frequency, $w(\ell, e) = \log((|C|+1)/(df(\ell, e)+1))$, where $df(\ell, e)$ counts calibration requests C in which expert e fires at least once at layer ℓ . The reweighted count is $\tilde{c}_r(\ell, e) = c_r(\ell, e) \cdot w(\ell, e)$.
3. *Select layers.* The mask \mathcal{S} is fit offline by greedy layer selection on ρ : starting from an empty mask, the layer whose inclusion most increases cumulative ρ is appended at each step, producing an ordering of all L layers. We keep the first N^* layers in this ordering, where N^* is the size at which cumulative ρ peaks (the \star in Fig. 8). At runtime, the per-request reweighted counts on \mathcal{S} stack to $x_r = [\tilde{c}_r(\ell)]_{\ell \in \mathcal{S}} \in \mathbb{R}^{N^* \cdot E}$.
4. *Normalize.* Divide by the vector’s length so similarity reflects the shape of expert usage, not the request’s token count:

$$s_r = x_r / \|x_r\|_2. \quad (2)$$




Figure 7. Signature quality ρ (Eq. 1) for six candidate transformations T . Bars are the mean across six cells (3 models \times 2 workloads); whiskers span the per-cell min/max.




Figure 8. Cumulative ρ (Eq. 1) versus the number of layers kept under greedy layer selection. One panel per model; task (blue) and language (orange) shown separately. The star marks the peak N^* chosen by ELDR’s offline fit.

The signature s_r is the building block consumed by the decode clustering and routing layers (§4.3).

4.2.3 Validation. We validate each design choice by its effect on ρ (Eq. 1), measured on 1,000 calibration requests per cell across three models (Qwen3-30B-A3B, GPT-OSS-120B, Gemma-4-26B-A4B) and two workloads (task, language).

Discrete beats continuous. Fig. 7 compares six choices of T on the full layer mask: count and count-idf (Steps 1–2 of §4.2.2), plus four ablations— $\sqrt{\text{count}}$; *gate prob* and *gate logit*, the summed softmax and raw gate scores over prefill tokens; and *binary*, the indicator that an expert appears in the top- k at least once. The three count-based variants beat both continuous ones by >0.035 in mean ρ : continuous scores put mass on experts the decode top- k never loads, and that mass cannot appear in p_i .

IDF denoises. Among the discrete variants, count-idf raises mean ρ over plain count by 1.5pt and its worst case ($\min_c \rho$) by 4.9pt (0.63 \rightarrow 0.68 on GPT-OSS language), with no loss on the best case. The gain concentrates on GPT-OSS, where a few generalist experts dominate the raw count and inflate the ℓ_2 norm; IDF down-weights them and surfaces the rare experts that distinguish workloads (+8.9pt on GPT-OSS task).

Layer mask compresses. Running greedy layer selection on count-idf signatures, Fig. 8 plots cumulative ρ against layers kept. Every cell peaks at $N^* < L$, exceeding the all-layer $\rho(L)$ by 0.005 to 0.032: extra layers add dimensions that don’t separate requests and dilute the signature. ELDR’s offline fit uses this per-cell peak N^* .

4.3 Decode Clustering and Routing

Given an expert signature, ELDR routes the request to one of K decode workers under two objectives. 1) *Locality*—co-located requests should share experts—is what the signature is built to enable. 2) *Load balance* is the constraint routing cannot ignore: non-MoE compute (attention, QKV projection, normalization) scales with a worker’s batch size, so an overloaded worker spikes tail TPOT (§3.5.2). ELDR splits these objectives across two stages by the information each needs. Locality is visible offline from aggregate calibration data, so the offline stage (§4.3.1) clusters signatures into K locality regions of equal calibration volume, one per worker. Load depends on the instantaneous per-worker state at the prefill → decode handoff, so the online stage (§4.3.2) routes each request within its region toward the less-loaded worker.

4.3.1 Offline clustering. ELDR produces one centroid per decode worker by clustering the calibration signatures into K groups such that signatures within a group are close (locality) and groups are balanced in size. K -means [21] is the standard algorithm for locality-driven clustering: it minimizes the mean cosine distance between each signature and its assigned centroid. This objective optimizes locality but has no size penalty—a single centroid in a tight, dense region serves all those points cheaply because they are close to it, while a centroid covering a sparse region serves few points but at similar per-point cost. On skewed data, this produces uneven cluster sizes, and since each cluster maps to one decode worker, the imbalance translates directly to uneven decoder load at runtime, causing tail latency spikes. The missing piece is a size constraint: ELDR uses **Hungarian-balanced K -means** [22], which replaces nearest-neighbor assignment with a globally optimal assignment: each centroid takes at most $\lceil N/K \rceil$ points, and the assignment minimizing total cosine distance is found by the Hungarian algorithm. Online routing (§4.3.2) uses cosine similarity to the resulting K unit centroids to identify candidate decoders. Because K is the number of decode workers, scaling the decode pool up or down can be absorbed by a re-cluster at the new K —a cheap offline re-fit (under 10 s, §6.2) that needs no re-deployment; the same re-fit could track workload drift.

Fig. 9 projects the calibration signatures onto their first two principal components and overlays the balanced K -means centroids. Task domains (Code/Math/Medical/Legal) and WildChat languages (English/Chinese/Russian/French) occupy distinct regions across all three models—the signature space has genuine semantic structure for the clustering to exploit. The centroids spread across the distinct regions rather than crowding onto the densest one, which is exactly the property the balance constraint buys: equal cluster size is only achievable when centroids cover all the distinct regions, not when they collapse onto the most populated one. Some centroids consequently land in sparser regions of the

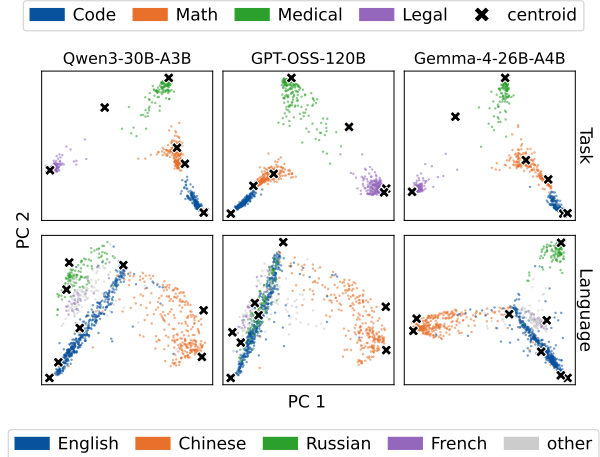


Figure 9. PCA of calibration signatures with Hungarian-balanced centroids ($K=8$ for legibility), colored by task domain (top) or top-4 WildChat [37] language (bottom).

projection—vanilla K -means would draw them onto the densest modes, but the balance constraint pulls them outward to absorb their $\approx N/K$ share.

4.3.2 Online routing. Pure top-1 routing—sending each request to its single nearest centroid—spikes tail latency because it ignores instantaneous load: a burst of requests near one centroid creates load imbalance at that worker while others sit idle. ELDR’s online routing rule, **locality-band routing**, favors locality when the signature singles out one worker, and falls back to load-aware balancing when several workers are nearly tied.

For each request, ELDR computes cosine similarity between the request’s expert signature and each of the K centroids; let $s^* = \max_k s_k$ denote the highest similarity. Among workers with $s_k \geq s^* - \tau$ (the *locality band*), ELDR selects the one with the smallest load—the number of in-flight decode requests on that worker, tracked at the proxy and updated as requests are dispatched and complete. The parameter $\tau \in [0, 1]$ interpolates between pure top-1 routing ($\tau=0$) and pure shortest-queue routing ($\tau=1$).

The locality band adapts to signature confidence: a confident signature that strongly favors one centroid puts few workers in the band (locality preserved), while an ambiguous signature near several centroids puts more workers in the band (load balance applied). ELDR uses $\tau=0.1$ throughout the paper; §6.3 examines this choice.

4.4 Prefix Cache Coherence

Prefix caching reuses KV blocks across requests that share a prompt prefix, skipping the prefill of the shared portion entirely. This creates a problem for ELDR. Routing requires the full prompt’s expert signature, but the expert activations for a cached prefix were produced by an earlier request and are not computed at this request’s prefill.

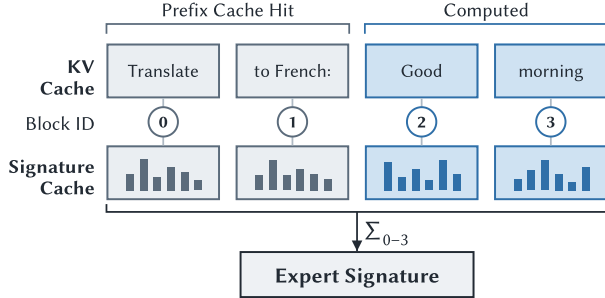


Figure 10. ELDR stores expert signatures at KV cache block granularity: the signature cache is co-indexed with KV cache.

A naive approach is to cache each request’s expert signature and reuse it when an identical prompt arrives. But a request can interact with the prefix cache in several ways: a prompt may miss the cache entirely, hit on some leading blocks and compute the rest (a partial hit), or hit on the entire prompt and skip prefill (a full hit); a cached block may also be evicted and reused for a different request. A per-request signature cache only matches on full hits, and any mismatch degrades the expert locality routing is meant to exploit.

The prefix cache manages KV blocks one at a time, so ELDR tracks the expert signature at the same granularity: one row per KV block, indexed by the same block id. As prefill writes a token’s K/V into a block, the top- k experts selected for that token at each MoE layer are written to the matching signature row. A request’s KV cache spans a block set $\mathcal{B}(r)$, so its signature is the sum over those blocks, $s_r = \sum_{b \in \mathcal{B}(r)} \text{sig}[b]$, exact regardless of which request populated each block: on a partial hit (Fig. 10), cached blocks contribute rows from the earlier request that first prefilled them and fresh blocks contribute this request’s rows, and the sum matches a cold prefill of the same prompt. The signature cache inherits the KV cache’s lifecycle—a reused block resets its signature, a cached prefix carries its signature, and an eviction reclaims both stores together—so ELDR keeps no prefix-tree, eviction, or sharing state of its own.

ELDR’s signature cache is a preallocated GPU tensor with one slot per KV-cache block, using one int8 per (block, MoE layer, expert); its total size is bounded by $\text{num_gpu_blocks} \cdot L \cdot E$ bytes. For the models we evaluate ($L \leq 48$, $E=128$), each block’s row is at most 6 KiB, keeping the whole cache well under 1% of the KV cache it shadows. Each forward accumulates the new tokens’ expert picks into the matching rows in a single GPU pass; retrieving a request’s signature sums its blocks’ rows into its $[L, E]$ signature, batched into the post-step device-to-host copy. Indexed by block id and tied to the KV cache’s block lifecycle, the cache adds no allocator, no eviction state, and no synchronization beyond what the KV cache already provides.

5 Implementation

ELDR is implemented on top of vLLM [16] as three components: a routing proxy that dispatches each request to a decoder by signature locality, engine-side hooks that emit per-request expert signatures during prefill, and a one-shot offline pipeline that clusters captured signatures into per-decoder centroids—roughly 2,000 lines of new Python.

Signature path. A prefill hook accumulates per-layer per-expert activation counts at KV-block granularity, stored in an expert signature cache keyed by block id. At prefill completion, the prompt’s per-block signatures are summed and piggybacked onto the existing prefill–decode handoff—no new RPC channel, and no per-request state in the proxy.

Per-decoder EP load balancing. Following METRO [36]—which balances expert-parallel (EP) ranks by *activated-expert* load rather than token count—we use the calibration trace to estimate each cluster’s expected per-expert decode activation, then for every (decoder, layer) greedily assign the next most-active expert to the lightest-loaded rank. Since each decoder serves a distinct cluster, a single global EP layout would balance one cluster at the others’ expense; per-decoder placement matches each decoder to the workload it sees while holding the per-rank expert count fixed (uniform weight memory).

Offline fit. A single script turns the calibration trace into the artifacts loaded at startup: a routing artifact (layer mask, IDF weights, and K Hungarian-balanced centroids in the masked, IDF-weighted, L2-normalized space) and, for multi-GPU expert parallelism, an expert placement file. It runs once per (model, dataset, K); recalibration is needed only when the model, dataset, or prefill–decode topology changes.

6 Evaluation

Testbed. We evaluate ELDR on a 5-node cluster of AMD MI300X GPUs (8 GPUs per node, 192 GB HBM each), interconnected by 400 Gbps NDR InfiniBand with 8 ConnectX-7 HCAs per node (one per GPU). All experiments run on vLLM [16] 0.21.0rc1 / ROCm 7.2, with prefill–decode disaggregation served by vLLM’s NIXL [26] connector.

Models. We exercise four open MoE models. Three are served at $TP=1$: **Qwen3-30B-A3B** [35] (128 experts, top-8, bf16), **GPT-OSS-120B** [28] (128 experts, top-4, mxfp4), and **Gemma-4-26B-A4B** [9] (128 experts, top-8, bf16). For the large-MoE scaling study, **Qwen3-235B-A22B** [35] (128 experts, top-8, bf16) is served with $TP=4 / EP=4$ per instance.

Datasets. Each model is evaluated on two workloads. **Task** is an 11,668-prompt mix spanning four domains drawn from public benchmarks, with naturally unequal domain shares reflecting source-dataset sizes: 3,600 *legal* [3], 2,779 *code* [1, 4, 17, 40], 2,744 *math* [5, 11, 13, 19, 20], and 2,545 *medical* [12, 14, 15]² — a $1.41\times$ largest/smallest ratio that stresses load balance when routing by signature. **Language**

²We use the professional-medicine subset of MMLU [12].

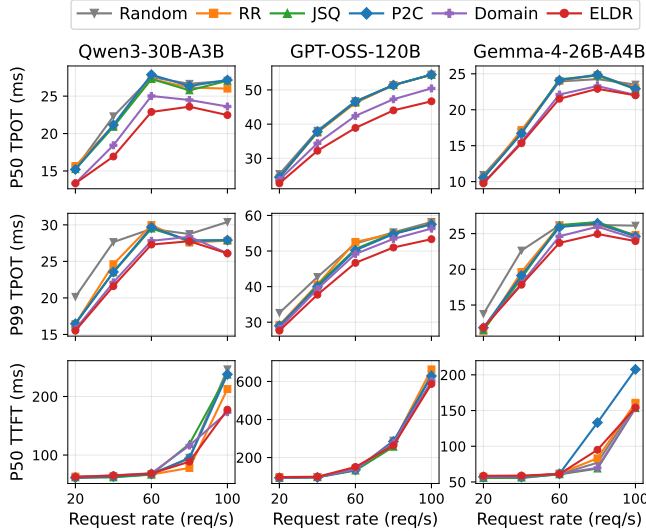


Figure 11. TPOT (median, p99) and median TTFT vs request rate on the **task** workload at 8P16D.

is a 14,000-prompt subset of WildChat [37] that inherits the heavy language skew shown in Fig. 5: English, Chinese, Russian, and French account for 87.6% of the volume. For each dataset, the 1,000-prompt offline calibration split is disjoint from the serving evaluation set.

Topologies and workload. Unless noted, the main evaluation uses an 8-prefiller / 16-decoder (8P16D) topology, with each instance at $TP=1$ so the 24 GPUs straddle 3 nodes. We additionally sweep 8P8D and 8P24D for topology generalization, and $2P8D \times (TP=4, EP=4)$ for the 235B run, the largest deployment in this study at 40 GPUs spanning all five nodes of the cluster. Requests are Poisson-distributed at offered rates from 20 to 100 qps (24–56 qps for 235B), each rate held for 120s with a 30s warmup. We cap output at 512 tokens and enforce `ignore_eos`.

Routers and baselines. Our serving stack follows the prefill/decode policy split adopted by vLLM [16], SGLang [38], and NVIDIA Dynamo [25]: cache- and hash-affinity policies on the prefill side, where the prefiller holds cross-request KV state, and load-balance policies on the decode side, where each request receives a fresh per-request KV transfer and the decoder holds no cross-request state. We adopt PrefixHash—from the affinity family—at our prefill, with a least-loaded fallback when the prefix-matched prefiller is above a $1.25\times$ average-load threshold. On the decode side we compare ELDR against four load-balance baselines—Random, Round-Robin (RR), Join-Shortest-Queue (JSQ) [34], and Power-of-Two-Choices (P2C) [23]—and Domain, a naïve locality-aware baseline that mirrors ELDR but routes on a given oracle domain label instead of the expert signature: it splits the decoders among the domains in proportion to the calibration domain mix, and sends each request to the decoders assigned

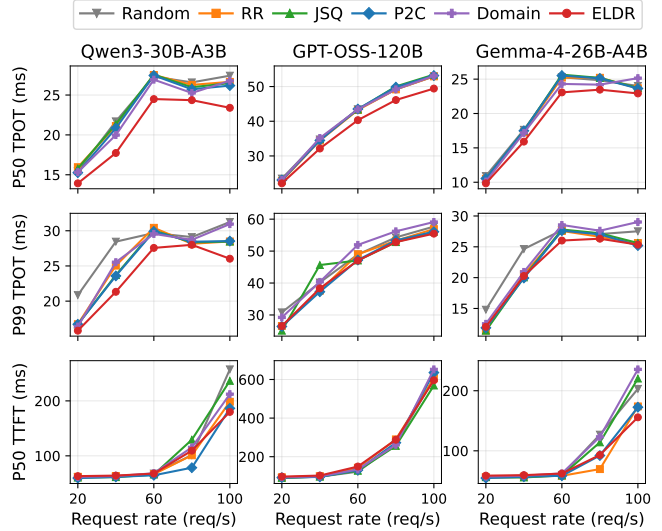


Figure 12. TPOT (median, p99) and median TTFT vs request rate on the **language** workload at 8P16D.

to its domain, load-balanced among them by JSQ. Within each cell, all six routers share the same prefill policy and decoder pool, so a cell isolates the decoder routing decision.

ELDR configuration. All experiments use the design from §4: count-idf signatures restricted to a greedy ρ -selected layer subset, Hungarian-balanced K -means with $K=D$ (one cluster per decoder), and $\tau=0.1$ locality-band routing. The 235B run additionally enables per-cluster per-layer expert-rank permutation for EP load balancing. Signature capture takes 4–15 minutes per (model, dataset) for the 1,000 calibration prompts; the offline fit (greedy mask selection plus balanced K -means) completes in under 10 seconds.

6.1 Main Results

ELDR reduces median TPOT across all three models and both workloads, with the size of the win tracking how separable the workload’s expert activations are. The two baseline families isolate ELDR’s two ingredients: the load balancers add load-awareness without locality, Domain adds static locality without live load, and ELDR combines finer locality with live load. We report the mean across the five request rates; the improvement holds across the whole sweep rather than at a single operating point.

On the task workload (Fig. 11), where prompts fall into well-defined domains (code, math, medical, legal), ELDR reduces median TPOT by 7.0–13.9% and tail TPOT by 3.4–6.0% over the best load balancer, and sits below every one at every rate. Domain is a stronger baseline here—its labels align with the model’s expert clusters, so a static per-domain partition concentrates each block’s expert working set and itself reduces median TPOT by 6.8–9.7% over the load balancers—yet ELDR still beats it, by 1.4–6.9% on median and 1.6–4.5%

Table 1. ELDR runtime overhead (Qwen3-30B-A3B, task, 8P16D, 60 req/s; median TTFT 69 ms).

Component	Locus	P50	% TTFT
record() hook	prefill host	0.02 ms	< 0.1
reduce() scatter	prefill GPU	0.48 ms	0.7
stage_sigs() D2H	prefill GPU	0.21 ms	0.3
pop_sig fetch	scheduler	7.0 μ s	< 0.1
Route (τ -JSQ)	proxy	0.15 ms	0.2
Total		0.86 ms	1.2

on tail TPOT, because its $K=16$ signature clusters resolve locality more finely than four labels and its τ -band spills load across cluster boundaries that Domain’s hard partition forbids. ELDR’s median TTFT tracks the baselines and is lower near saturation, where its faster decoders relieve prefill back-pressure.

On the language workload (Fig. 12), where domain boundaries are softer, ELDR reduces median TPOT by 5.9–10.0% over the best load balancer; mean tail TPOT reduces by 6.2% on Qwen3-30B-A3B and regresses by 1.5% on GPT-OSS-120B and 0.2% on Gemma-4-26B-A4B, though at the per-cell peak all three reduce (9.6%, 1.0%, 5.6%). Domain collapses here: a language label is a coarse proxy for expert activation—Fig. 9 shows several of ELDR’s centroids falling within a single language, so one static block per language mixes unrelated signature clusters—and its skewed mix over-subscribes one block, so Domain only matches the load balancers on median TPOT (within 3% of RR) and regresses tail TPOT by up to 6.1% as the hot block saturates. ELDR beats Domain by 5.7–9.1% on median and 7.0–9.5% on tail TPOT: its finer signature clusters capture the intra-language sub-structure that the coarse labels miss (Fig. 9), and its τ -band routing rebalances live load that the static partition cannot. Median TTFT again tracks the baselines.

6.2 Overhead Analysis

ELDR’s overhead has a one-time offline part and a per-request serving part. Offline, the router is built once per deployment: a signature-capture pass over the 1,000 calibration prompts, then the fit—greedy mask plus balanced K -means—which completes in under 10 s on CPU and is cheap to re-run on configuration changes. At serving time (Table 1), ELDR adds 0.86 ms per request, 1.2% of the 69 ms median TTFT, dominated by per-forward signature capture on the prefiller; the scheduler fetch and routing decision are sub-percent. The signature cache occupies 0.24% of HBM and each request carries a 12 KiB signature, negligible against the multi-megabyte KV transfer. This is why median TTFT stays within the across-baseline spread in the main results.




Figure 13. Mean active experts per decode step on Qwen3-30B-A3B in the task domain on an 8P16D cluster.




Figure 14. TPOT P50/P99 % Δ vs. RR at $r=60$ req/s (8P16D, six (model, dataset) cells) for two signature transforms: the IDF-reweighted top- k count (**count-idf**) vs. the continuous softmax gate (**gate-prob**). Rest of the recipe fixed (greedy mask, balanced K -means $K=16$, $\tau=0.1$).

6.3 Design Validation

Active-Expert Reduction. We validate that ELDR’s TPOT improvements come from a reduction in the number of distinct experts activated per decode step. Fig. 13 measures this count on Qwen3-30B-A3B serving the task domain at 8P16D, aggregated over $\sim 50K$ decode steps. ELDR reduces the per-step active-expert count by 22.0% on average across decode batch sizes compared to RR.

Expert Signature. We validate ELDR’s signature choice: the IDF-reweighted top- k count (count-idf). Fig. 14 compares it against the strongest continuous alternative, the softmax gate-probability, to validate the discrete signature design (§4.2). count-idf reduces TPOT P50 by an additional 3 pp on average and up to 14 pp across the six cells, confirming that ρ faithfully ranks signature quality and that discrete prefill counts are the right primitive.

Offline Cluster Balance. Fig. 15 validates ELDR’s choice of Hungarian-balanced K -means over vanilla K -means at $\tau=0.1$. Vanilla K -means reduces median TPOT by up to 9.8% but regresses tail TPOT by as much as 17.4%: the per-decoder locality win is real, but the load imbalance pushes the tail well past round-robin. Hungarian-balanced K -means recovers both metrics, reducing P50 by up to 12.6% and P99 by up

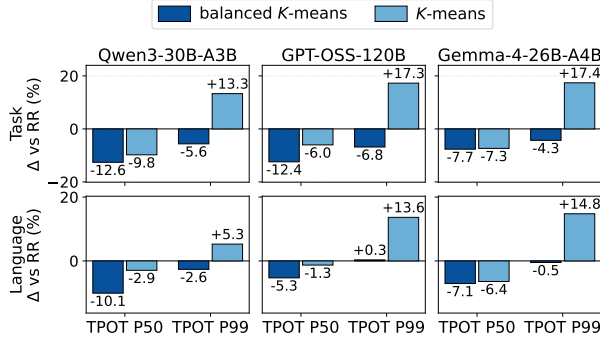


Figure 15. Mean % Δ vs. RR over five request rates (20–100 qps) at 8P16D with $\tau=0.1$.




Figure 16. Mean % Δ vs. RR (five rates, 20–100 qps; 8P16D) for six cells (three models \times two datasets) at four τ values.

to 6.8%—uniform decoder utilization keeps the tail in check without sacrificing the median win.

Locality Band Width. Fig. 16 validates ELDR’s choice of $\tau=0.1$ across $\tau \in \{0, 0.1, 0.2, 0.3\}$. Pure top-1 routing ($\tau=0$) regresses tail TPOT on four of six workloads—by as much as 7.9% on Gemma language and 6.7% on Gemma task—because a transient burst of similar requests lands on one decoder. A small band of $\tau=0.1$ removes the regression on every workload while reducing median TPOT by 5.2–12.7% relative to RR. Beyond that the tail reduction saturates and median TPOT erodes as more requests spill outside their locality band, so ELDR settles on $\tau=0.1$.

Prefix Cache Composition. Fig. 17 validates that ELDR composes with the prefix cache. Turning the cache on collapses TTFT for both routers without changing TPOT, and ELDR’s TPOT advantage over RR ($\approx 13\%$ on P50, 5% on P99) is preserved across both cache states. The locality benefit and the prefix-cache benefit are additive. A secondary back-pressure effect shows up on TTFT: ELDR’s faster decoders drain requests sooner, keeping prefill queues shorter and reducing TTFT when the cache is off.

6.4 Generalization

Decoder Pool Size. To check that ELDR generalizes across decoder counts, we sweep three prefill–decode topologies on Qwen3-30B-A3B (language workload, $TP=1$), holding




Figure 17. Composition with prefix caching on GPT-OSS-120B in the task domain on an 8P16D cluster at $r=100$ req/s, with requests sampled cyclically from 2000 task prompts.

Table 2. Topology generalization on Qwen3-30B-A3B, language workload: mean % Δ vs. RR across 20–100 qps.

	8P8D	8P16D	8P24D
TPOT P50	-8.0%	-9.8%	-10.2%
TPOT P99	-2.5%	-0.8%	-1.1%




Figure 18. TPOT (median, p99) for Qwen3-235B-A22B on language workload at 2P8D with $TP=4$ and $EP=4$ per instance (40 GPUs across 5 nodes).

the prefiller pool at 8 and growing the decoder pool: 8P8D (16 GPUs), 8P16D (24), and 8P24D (32). Over five rates (20–100 qps), the mean median-TPOT reduction vs. round-robin grows monotonically with the pool—8.0%, 9.8%, 10.2% (Table 2)—while tail TPOT stays within noise of round-robin throughput. This follows the mechanism: more decoders split each workload into finer clusters with narrower expert coverage, so the per-decoder expert union shrinks and the latency reduction scales with it. ELDR thus generalizes across topologies, improving as the pool expands.

Large MoE with Expert Parallelism. We scale ELDR to Qwen3-235B-A22B at 2P8D, $TP=4$, $EP=4$ per instance (40 GPUs, 5 nodes), where each decoder shards attention and experts across four GPUs. Clustering alone is insufficient here—a cluster’s hot experts can concentrate on one EP rank and bottleneck the decoder—so each ELDR decoder pairs clustering with the per-decoder expert placement of $\$5$, spreading each cluster’s expected expert work evenly across the four GPUs. On WildChat (Fig. 18, 24–56 req/s), ELDR reduces median TPOT by 2.7–4.3% and tail TPOT by 0.6–2.0% at every rate, confirming that locality routing generalizes to large-MoE deployments with expert parallelism.

7 Related Work

PD disaggregated serving. Prefill/decode (PD) disaggregation runs compute-bound prefill and memory-bandwidth-bound decode on separate workers, removing interference and meeting phase SLOs [29, 30, 39]; we adopt its *xPyD* topology. Cache-aware routers exploit prefix/KV locality [31, 38] but treat decode workers as interchangeable for expert computation. ELDR complements them by adding *expert-activation* locality as the decode-routing signal, leaving outputs unchanged.

MoE Expert locality. Expert parallelism (EP) shards experts across devices, where skewed routing causes load imbalance. One line rebalances *within* a deployment, replicating hot experts and rerouting tokens across EP ranks [6, 8, 24, 36]—*intra*-worker balancing that evens per-GPU expert load but not which worker serves a request. Closest to us, systems exploiting expert-activation locality—clustering requests by prefill activations [2] or colocating experts with their tokens [18]—target inter-node all-to-all *communication*, not decode bandwidth. A third line hits the same bandwidth bottleneck but *approximately*, altering decode expert selection: dropping low-importance experts [10], piggybacking on already-loaded ones [27], or sharing experts across a batch [33], trading accuracy for speed.

ELDR differs on both axes. It performs *inter*-worker balancing, routing whole requests *across* decode workers by expert locality to shrink each worker’s active experts. It is also *lossless*: it changes only which worker serves a request, never a token’s expert selection, so outputs match standard top-*k* gating. It is thus exact where batch-aware methods approximate, and orthogonal to *intra*-worker balancing (each decode worker can still run EPLB internally), composing with these methods rather than replacing them.

8 Conclusion

Decode routing in PD-disaggregated MoE serving optimizes only decode-worker load. ELDR adds a second axis, *expert locality*: a prefill-derived signature drives load-tolerant routing that stays coherent with the prefix cache, losslessly lowering median and tail TPOT across models, workloads, and topologies.

References

- [1] Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, and Charles Sutton. 2021. Program Synthesis with Large Language Models. arXiv:2108.07732 [cs.PL] <https://arxiv.org/abs/2108.07732>
- [2] Abhimanyu Bambhaniya, Geonhwa Jeong, Jason Park, Jiecao Yu, Jaewon Lee, Pengchao Wang, Changkyu Kim, Chunqiang Tang, and Tushar Krishna. 2026. Scaling Multi-Node Mixture-of-Experts Inference Using Expert Activation Patterns. arXiv:2604.23150 [cs.LG] <https://arxiv.org/abs/2604.23150>
- [3] Ilias Chalkidis, Abhik Jana, Dirk Hartung, Michael Bommarito, Ion Androutsopoulos, Daniel Katz, and Nikolaos Aletras. 2022. LexGLUE: A

Benchmark Dataset for Legal Language Understanding in English. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Smaranda Muresan, Preslav Nakov, and Aline Villavicencio (Eds.). Association for Computational Linguistics, Dublin, Ireland, 4310–4330. doi:10.18653/v1/2022.acl-long.297

- [4] Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. 2021. Evaluating Large Language Models Trained on Code. (2021). arXiv:2107.03374 [cs.LG]
- [5] Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. Training Verifiers to Solve Math Word Problems. arXiv:2110.14168 [cs.LG] <https://arxiv.org/abs/2110.14168>
- [6] DeepSeek-AI. 2025. EPLB: Expert Parallelism Load Balancer. <https://github.com/deepseek-ai/EPLB>.
- [7] DeepSeek-AI, Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Daya Guo, Dejian Yang, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Haowei Zhang, Honghui Ding, Huajian Xin, Huazuo Gao, Hui Li, Hui Qu, J. L. Cai, Jian Liang, Jianzhong Guo, Jiaqi Ni, Jiashi Li, Jiawei Wang, Jin Chen, Jingchang Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, Junxiao Song, Kai Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Lei Xu, Leyi Xia, Liang Zhao, Litong Wang, Liyue Zhang, Meng Li, Miaojun Wang, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Mingming Li, Ning Tian, Panpan Huang, Peiyi Wang, Peng Zhang, Qiancheng Wang, Qihao Zhu, Qinyu Chen, Qiushi Du, R. J. Chen, R. L. Jin, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, Runxin Xu, Ruoyu Zhang, Ruyi Chen, S. S. Li, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shaoqing Wu, Shengfeng Ye, Shengfeng Ye, Shirong Ma, Shiyu Wang, Shuang Zhou, Shuiping Yu, Shunfeng Zhou, Shuting Pan, T. Wang, Tao Yun, Tian Pei, Tianyu Sun, W. L. Xiao, Wangding Zeng, Wanbiao Zhao, Wei An, Wen Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, X. Q. Li, Xiangyue Jin, Xianzu Wang, Xiao Bi, Xiaodong Liu, Xiaohan Wang, Xiaojin Shen, Xiaokang Chen, Xiaokang Zhang, Xiaosha Chen, Xiaotao Nie, Xiaowen Sun, Xiaoxiang Wang, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xingkai Yu, Xinnan Song, Xinxia Shan, Xinyi Zhou, Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, Y. K. Li, Y. Q. Wang, Y. X. Wei, Y. X. Zhu, Yang Zhang, Yanhong Xu, Yanhong Xu, Yanping Huang, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Li, Yaohui Wang, Yi Yu, Yi Zheng, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Ying Tang, Yishi Piao, Yisong Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yu Wu, Yuan Ou, Yuchen Zhu, Yudian Wang, Yue Gong, Yuheng Zou, Yujia He, Yukun Zha, Yunfan Xiong, Yunxian Ma, Yuting Yan, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Z. F. Wu, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhen Huang, Zhen Zhang, Zhenda Xie, Zhengyan Zhang, Zhewen Hao, Zhibin Gou, Zhicheng Ma, Zhigang Yan, Zhihong Shao, Zhipeng Xu, Zhiyu Wu, Zhongyu Zhang, Zhuoshu Li, Zihui Gu, Zijia

- Zhu, Zijun Liu, Zilin Li, Ziwei Xie, Ziyang Song, Ziyi Gao, and Zizheng Pan. 2025. DeepSeek-V3 Technical Report. arXiv:2412.19437 [cs.CL] <https://arxiv.org/abs/2412.19437>
- [8] Seokjin Go and Divya Mahajan. 2025. MoETuner: Optimized Mixture of Expert Serving with Balanced Expert Placement and Token Routing. arXiv:2502.06643 [cs.LG] <https://arxiv.org/abs/2502.06643>
- [9] Google DeepMind. 2025. Gemma 4 26B-A4B. <https://huggingface.co/google/gemma-4-26b-a4b>.
- [10] Vima Gupta, Jae Hyung Ju, Kartik Sinha, Ada Gavrilovska, and Anand Padmanabha Iyer. 2026. Lynx: Enabling Efficient MoE Inference through Dynamic Batch-Aware Expert Selection. arXiv:2411.08982 [cs.LG] <https://arxiv.org/abs/2411.08982>
- [11] Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Leng Thai, Junhao Shen, Jinyi Hu, Xu Han, Yujie Huang, Yuxiang Zhang, Jie Liu, Lei Qi, Zhiyuan Liu, and Maosong Sun. 2024. OlympiadBench: A Challenging Benchmark for Promoting AGI with Olympiad-Level Bilingual Multimodal Scientific Problems. arXiv:2402.14008 [cs.CL] <https://arxiv.org/abs/2402.14008>
- [12] Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. Measuring Massive Multitask Language Understanding. arXiv:2009.03300 [cs.CY] <https://arxiv.org/abs/2009.03300>
- [13] Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. Measuring Mathematical Problem Solving With the MATH Dataset. arXiv:2103.03874 [cs.LG] <https://arxiv.org/abs/2103.03874>
- [14] Di Jin, Eileen Pan, Nassim Oufattole, Wei-Hung Weng, Hanyi Fang, and Peter Szolovits. 2021. What Disease Does This Patient Have? A Large-Scale Open Domain Question Answering Dataset from Medical Exams. *Applied Sciences* 11, 14 (2021). doi:10.3390/app11146421
- [15] Qiao Jin, Bhuwan Dhingra, Zhengping Liu, William Cohen, and Xinghua Lu. 2019. PubMedQA: A Dataset for Biomedical Research Question Answering. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, Kentaro Inui, Jing Jiang, Vincent Ng, and Xiaojun Wan (Eds.). Association for Computational Linguistics, Hong Kong, China, 2567–2577. doi:10.18653/v1/D19-1259
- [16] Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient Memory Management for Large Language Model Serving with PagedAttention. arXiv:2309.06180 [cs.LG] <https://arxiv.org/abs/2309.06180>
- [17] Yuhang Lai, Chengxi Li, Yiming Wang, Tianyi Zhang, Ruiqi Zhong, Luke Zettlemoyer, Wen-tau Yih, Daniel Fried, Sida Wang, and Tao Yu. 2023. DS-1000: a natural and reliable benchmark for data science code generation. In *Proceedings of the 40th International Conference on Machine Learning (Honolulu, Hawaii, USA) (ICML '23)*. JMLR.org, Article 756, 27 pages.
- [18] Yan Li, Zhenyu Zhang, Zhengang Wang, Pengfei Chen, and Pengfei Zheng. 2026. Semantic Parallelism: Redefining Efficient MoE Inference via Model-Data Co-Scheduling. arXiv:2503.04398 [cs.LG] <https://arxiv.org/abs/2503.04398>
- [19] Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. 2023. Let’s Verify Step by Step. arXiv:2305.20050 [cs.LG] <https://arxiv.org/abs/2305.20050>
- [20] Wang Ling, Dani Yogatama, Chris Dyer, and Phil Blunsom. 2017. Program Induction by Rationale Generation : Learning to Solve and Explain Algebraic Word Problems. arXiv:1705.04146 [cs.AI] <https://arxiv.org/abs/1705.04146>
- [21] S. Lloyd. 1982. Least squares quantization in PCM. *IEEE Transactions on Information Theory* 28, 2 (1982), 129–137. doi:10.1109/TIT.1982.1056489
- [22] Mikko I. Malinen and Pasi Fränti. 2014. Balanced K-Means for Clustering. In *Structural, Syntactic, and Statistical Pattern Recognition*, Pasi Fränti, Gavin Brown, Marco Loog, Francisco Escolano, and Marcello Pelillo (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 32–41.
- [23] M. Mitzenmacher. 2001. The power of two choices in randomized load balancing. *IEEE Transactions on Parallel and Distributed Systems* 12, 10 (2001), 1094–1104. doi:10.1109/71.963420
- [24] Xuan-Phi Nguyen, Shrey Pandit, Austin Xu, Caiming Xiong, and Shafiq Joty. 2026. Least-Loaded Expert Parallelism: Load Balancing An Imbalanced Mixture-of-Experts. arXiv:2601.17111 [cs.LG] <https://arxiv.org/abs/2601.17111>
- [25] NVIDIA. 2026. NVIDIA Dynamo: A Datacenter-Scale Distributed Inference Serving Framework. <https://github.com/ai-dynamo/dynamo>.
- [26] NVIDIA Corporation. 2024. NIXL: NVIDIA Inference Xfer Library. <https://github.com/ai-dynamo/nixl>.
- [27] Costin-Andrei Oncescu, Qingyang Wu, Wai Tong Chung, Robert Wu, Bryan Gopal, Junxiong Wang, Tri Dao, and Ben Athiwaratkun. 2025. Opportunistic Expert Activation: Batch-Aware Expert Routing for Faster Decode Without Retraining. arXiv:2511.02237 [cs.LG] <https://arxiv.org/abs/2511.02237>
- [28] OpenAI, , Sandhini Agarwal, Lama Ahmad, Jason Ai, Sam Altman, Andy Applebaum, Edwin Arbus, Rahul K. Arora, Yu Bai, Bowen Baker, Haiming Bao, Boaz Barak, Ally Bennett, Tyler Bertao, Nivedita Brett, Eugene Brevdo, Greg Brockman, Sebastien Bubeck, Che Chang, Kai Chen, Mark Chen, Enoch Cheung, Aidan Clark, Dan Cook, Marat Dukhan, Casey Dvorak, Kevin Fives, Vlad Fomenko, Timur Garipov, Kristian Georgiev, Mia Glaese, Tarun Gogineni, Adam Goucher, Lukas Gross, Katia Gil Guzman, John Hallman, Jackie Hehir, Johannes Heidecke, Alec Helyar, Haitang Hu, Romain Huet, Jacob Huh, Saachi Jain, Zach Johnson, Chris Koch, Irina Kofman, Dominik Kundel, Jason Kwon, Volodymyr Kyrylov, Elaine Ya Le, Guillaume Leclerc, James Park Lennon, Scott Lessans, Mario Lezcano-Casado, Yuanzhi Li, Zhuohan Li, Ji Lin, Jordan Liss, Lily, Liu, Jiancheng Liu, Kevin Lu, Chris Lu, Zoran Martinovic, Lindsay McCallum, Josh McGrath, Scott McKinney, Aidan McLaughlin, Song Mei, Steve Mostovoy, Tong Mu, Gideon Myles, Alexander Neitz, Alex Nichol, Jakub Pachocki, Alex Paino, Dana Palmie, Ashley Pantuliano, Giambattista Parascandolo, Jongsoo Park, Leher Pathak, Carolina Paz, Ludovic Peran, Dmitry Pimenov, Michelle Pokrass, Elizabeth Proehl, Huida Qiu, Gaby Raila, Filippo Raso, Hongyu Ren, Kimmy Richardson, David Robinson, Bob Rotsted, Hadi Salman, Suvansh Sanjeev, Max Schwarzer, D. Sculley, Harshit Sikchi, Kendal Simon, Karan Singhal, Yang Song, Dane Stuckey, Zhiqing Sun, Philippe Tillet, Sam Toizer, Foivos Tsimpourlas, Nikhil Vyas, Eric Wallace, Xin Wang, Miles Wang, Olivia Watkins, Kevin Weil, Amy Wendling, Kevin Whinnery, Cedric Whitney, Hannah Wong, Lin Yang, Yu Yang, Michihiro Yasunaga, Kristen Ying, Wojciech Zaremba, Wenting Zhan, Cyril Zhang, Brian Zhang, Eddie Zhang, and Shengjia Zhao. 2025. gpt-oss-120b & gpt-oss-20b Model Card. arXiv:2508.10925 [cs.CL] <https://arxiv.org/abs/2508.10925>
- [29] Pratyush Patel, Esha Choukse, Chaojie Zhang, Aashaka Shah, Íñigo Goiri, Saeed Maleki, and Ricardo Bianchini. 2024. Splitwise: Efficient generative LLM inference using phase splitting. arXiv:2311.18677 [cs.AR] <https://arxiv.org/abs/2311.18677>
- [30] Ruoyu Qin, Zheming Li, Weiran He, Mingxing Zhang, Yongwei Wu, Weimin Zheng, and Xinran Xu. 2025. Mooncake: A KVCache-centric Disaggregated Architecture for LLM Serving. arXiv:2407.00079 [cs.DC] <https://arxiv.org/abs/2407.00079>
- [31] Vikranth Srivatsa, Zijian He, Reyna Abhyankar, Dongming Li, and Yiyang Zhang. 2024. Preble: Efficient Distributed Prompt Scheduling for LLM Serving. arXiv:2407.00023 [cs.DC] <https://arxiv.org/abs/2407.00023>
- [32] The llm-d Authors. 2026. llm-d: Kubernetes-Native Distributed Inference. <https://github.com/llm-d/llm-d>.

- [33] Daniil Vankov, Nikita Ivkin, Kyle Ulrich, Xiang Song, Ashish Khetan, and George Karypis. 2026. XShare: Collaborative in-Batch Expert Sharing for Faster MoE Inference. arXiv:2602.07265 [cs.LG] <https://arxiv.org/abs/2602.07265>
- [34] Wayne Winston. 1977. Optimality of the shortest line discipline. *Journal of Applied Probability* 14, 1 (1977), 181–189. doi:10.2307/3213271
- [35] An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu. 2025. Qwen3 Technical Report. *arXiv preprint arXiv:2505.09388* (2025).
- [36] Yanpeng Yu, Haiyue Ma, Krish Agarwal, Nicolai Oswald, Qijing Huang, Hugo Linsenmaier, Chunhui Mei, Ritchie Zhao, Ritika Borkar, Bitu Darvish Rouhani, David Nellans, Ronny Krashinsky, and Anurag Khandelwal. 2025. Efficient MoE Serving in the Memory-Bound Regime: Balance Activated Experts, Not Tokens. arXiv:2512.09277 [cs.DC] <https://arxiv.org/abs/2512.09277>
- [37] Wenting Zhao, Xiang Ren, Jack Hessel, Claire Cardie, Yejin Choi, and Yuntian Deng. 2024. WildChat: 1M ChatGPT Interaction Logs in the Wild. arXiv:2405.01470 [cs.CL] <https://arxiv.org/abs/2405.01470>
- [38] Lianmin Zheng, Liangsheng Yin, Zhiqiang Xie, Chuyue Sun, Jeff Huang, Cody Hao Yu, Shiyi Cao, Christos Kozyrakis, Ion Stoica, Joseph E. Gonzalez, Clark Barrett, and Ying Sheng. 2024. SGLang: Efficient Execution of Structured Language Model Programs. arXiv:2312.07104 [cs.AI] <https://arxiv.org/abs/2312.07104>
- [39] Yinmin Zhong, Shengyu Liu, Junda Chen, Jianbo Hu, Yibo Zhu, Xuanzhe Liu, Xin Jin, and Hao Zhang. 2024. DistServe: Disaggregating Prefill and Decoding for Goodput-optimized Large Language Model Serving. In *18th USENIX Symposium on Operating Systems Design and Implementation (OSDI 24)*. USENIX Association, Santa Clara, CA, 193–210. <https://www.usenix.org/conference/osdi24/presentation/zhong-yinmin>
- [40] Terry Yue Zhuo, Minh Chien Vu, Jenny Chim, Han Hu, Wenhao Yu, Ratnadira Widayarsi, Imam Nur Bani Yusuf, Haolan Zhan, Junda He, Indraneil Paul, Simon Brunner, Chen Gong, Thong Hoang, Armel Randy Zebaze, Xiaoheng Hong, Wen-Ding Li, Jean Kaddour, Ming Xu, Zhihan Zhang, Prateek Yadav, Naman Jain, Alex Gu, Zhoujun Cheng, Jiawei Liu, Qian Liu, Zijian Wang, Binyuan Hui, Niklas Muennighoff, David Lo, Daniel Fried, Xiaoning Du, Harm de Vries, and Leandro Von Werra. 2025. BigCodeBench: Benchmarking Code Generation with Diverse Function Calls and Complex Instructions. arXiv:2406.15877 [cs.SE] <https://arxiv.org/abs/2406.15877>