EMFUSE: ENERGY-BASED MODEL FUSION FOR DECISION MAKING

Anonymous authorsPaper under double-blind review

000

001

002003004

010 011

012

013

014

016

017

018

019

021

024

025

026

027

028

029

031

032 033 034

035

037

040

041

042

043

044

046

047

048

051

052

ABSTRACT

Model fusion has emerged as a promising research direction, offering a resourceefficient paradigm that leverages existing pre-trained models to circumvent the need for training from scratch. In this work, we investigate the fusion of models specifically adapted for decision-making tasks. This challenge divides into two distinct, yet related subproblems: the direct fusion of models that act as policy and the fusion of dynamics models that subsequently induce a policy. We suggest that these seemingly divergent subproblems can be unified through the lens of energy-based models (EBMs), which parameterizes a conditional distribution via an energy function where lower energy implies higher probability. Our framework, **EMFuse**, provides this convergence by leveraging the concept of energy as a common currency for fusion. For direct fusion of policies, such as those in language models, the output distribution is commonly softmax (Boltzmann), which essentially defines the negative logarithmic probability as an energy function. For dynamics models, existing works often train a set of models on the same dataset to obtain robust uncertainty estimation; such an ensemble approach leads to an exponential explosion in computational complexity when it comes to dynamics fusion across multiple sets of models. To overcome this, we introduce the Any-step Dynamics Energy-based Transition Model (ADETM), a novel architecture that performs efficient single-model-per-dataset uncertainty estimation with its energy-based backbone, thereby avoiding this computational explosion. Our EMFuse framework surpasses other baselines by 0.34% to 6.63% on single/cross domain discrete decision-making benchmarks, and achieved an extra 2.3 to 7.4 normalized points on average in D4RL MuJoCo continuous-control scenarios.

1 Introduction

Modern deep learning ecosystems are increasingly populated by highly specialized models that power tasks from text generation to complex problem-solving (Alto, 2024; Wu et al., 2024). Directly training new monolithic models is expensive, while discarding existing expertise is wasteful. *Model fusion* promises a resource-efficient alternative by combining pre-trained experts into a stronger system (Wortsman et al., 2022; Matena & Raffel, 2022; Hu et al., 2021; Lu et al., 2024). However, while model fusion has seen rapid growth, its application to the specialized domain of decision-making remains comparatively underexplored. Recognizing the immense potential in this area, our work focuses specifically on *fusion for decision making*, a critical frontier for creating more capable and adaptable intelligent agents from fusion (Levine et al., 2020; Moerland et al., 2023).

At a high level, the behavior of a decision making agent is governed by either a directly learned policy or a policy derived from a learned dynamics model. This fundamental architectural choice presents two distinct points of intervention for model fusion: (1) *Direct policy fusion:* This approach involves combining the output distributions of multiple policies at the point of decision, potentially deriving a single, more robust action by taking all distributions into consideration. (2) *Dynamics fusion:* In model-based RL (MBRL), dynamics models learned from offline logs (Levine et al., 2020) enable further policy training (Chua et al., 2018; Janner et al., 2019; Yu et al., 2020; Hafner et al., 2020). The dynamics fusion approach acts at a more foundational level by merging the agent's dynamics model (their predictive understanding of the environment's dynamics). The goal is to construct a single, more comprehensive and reliable simulation of the transitions by unifying each

model's knowledge of the world. A superior policy could be learned under this unification, which provides a richer and more accurate basis for decision-making.

These seemingly divergent subproblems can be unified. We take the perspective of Energy-Based Models (EBMs), where distributions are represented through energies whose additive composition corresponds to multiplicative densities (LeCun et al., 2006; Hinton, 2002). In this view, both policy outputs and dynamics likelihoods can be written as $p(y \mid x) = \exp(-E(x,y))/Z(x)$, and fusion becomes an *energy sum* $E_{\text{fuse}} = \sum_i \lambda_i E_i$. This makes policy fusion and dynamics fusion two instances of the same principle, differing only by the choice of x, y and the sampler.

Our framework, **EMFuse**, applies energy-additive fusion to two key settings: (1) direct policy fusion and (2) dynamics model fusion. This approach enables efficient and effective inference without requiring additional training of the full models. We observe that the fused energy distribution remains close to each constituent policy on its specific domain, as measured by small KL-divergence. This faithfulness makes the fused energy a powerful indicator for selecting the best policy for a given context, a principle we formalize in our **EMSelect** algorithm.

Most existing dynamics model learning work trains an ensemble of models for uncertainty estimation (Yu et al., 2020). However, fusing ensembles of dynamics models quickly becomes intractable due to exponential blow-up. Inspired by ADMPO (Lin et al., 2025) that enables uncertainty estimation on one model by its ability to predict multiple next state with flexible action sequence input, we therefore introduce the *Any-step Dynamics Energy-based Transition Model (ADETM)*, an energy-based dynamics model that performs single-model-per-dataset uncertainty estimation while retaining multi-step context. Additional state and action encoders are added to make ADETM available to the energy-based context. This setup avoids the combinatorial cost of cross-ensemble fusion.

Our main contributions are listed as follows:

- EMFuse: We formalize policy and dynamics fusion under the same energy framework $E_{\text{fuse}} = \sum_{i} \lambda_i E_i$, connecting classical PoE (Hinton, 2002) to decision-making settings.
- EMSelect: A EMFuse based policy selection framework that achieved additional 1.18%-1.31% gain on top of EMFuse.
- **ADETM for tractable dynamics fusion:** A single-model-per-dataset energy-based world model with any-step context that circumvents cross-ensemble explosion while preserving uncertainty-aware behavior.
- Empirical gains: On single/cross-domain discrete decision-making benchmarks, EMFuse improves accuracy by **0.34**%–**6.63**%, and on D4RL MuJoCo continuous control (Fu et al., 2020) it adds **+2.3** to **+7.4** normalized points on average over other fusion baselines.

2 Preliminaries

2.1 TOWARDS ENERGY-BASED MODELS FOR DECISION MAKING

Markov Decision Process We consider a discounted MDP $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, r, \gamma)$ with an offline dataset $\mathcal{D} = \{(s, a, r, s')\}$ collected by an unknown behavior policy π_{β} (Puterman, 1994; Sutton & Barto, 2018; Levine et al., 2020). Offline training faces a *support gap*: test-time states/actions may lie outside the empirical support of \mathcal{D} , causing distribution shift and value overestimation. This makes calibrated *uncertainty* and *support awareness* central to algorithm design.

Behavior modeling: explicit vs. implicit One axis models or constrains by the behavior distribution. The *Explicit* approaches fit an estimate $\hat{\pi}_{\beta}(a \mid s)$ and use them as a prior or regularizer for policy learning (Fujimoto et al., 2019; Kostrikov et al., 2022). *Implicit* approaches shape learning objectives to bias solutions toward the dataset support, e.g., conservative value learning and pessimistic objectives (Kumar et al., 2020; Levine et al., 2020). Both aim to remain out-of-distribution (OOD)-robust; they differ in how support is represented (density vs. objective penalties).

Dynamics Models and Uncertainty Estimation. A second axis learns a model of the dynamics and plans or trains policies inside it. Classical likelihood models parameterize $p_{\theta}(s' \mid s, a)$ with

 Gaussian or mixture families; latent/variational world models compress trajectories into learned state abstractions (e.g., *World Models*, PlaNet, Dreamer) (Ha & Schmidhuber, 2018; Hafner et al., 2019; 2020). Recent work also explores diffusion/score-based generative modeling for sequential control (Janner et al., 2022). Likelihood models require calibration under off-distribution conditions; consequently, robust uncertainty estimation becomes key.

Epistemic uncertainty is commonly estimated through bootstrapped or independently initialized model ensembles; aleatoric uncertainty is captured by predictive dispersion (Kendall & Gal, 2017; Chua et al., 2018). ADMPO (Lin et al., 2025) trains a single recurrent net on variable-length transitions, achieving ensemble-like robustness without model repetition, further simplifying the framework. Ensembles not only improve robustness, they also provide *composable* uncertainty signals that will align naturally with an energy-based view.

Energy Based Models (EBMs). EBMs represent distributions via unnormalized energies, $p_{\theta}(x) \propto \exp\left(-E_{\theta}(x)\right)$, trained with contrastive or score-based criteria (Hinton, 2002; Hyvärinen, 2005; LeCun et al., 2006). Energies add *linearly* when combining independent experts, corresponding to multiplicative composition of their unnormalized densities. Specializing to dynamics, Energy-based Transition Models (**ETMs**) (Chen et al., 2024) learn a transition energy $E_{\theta}(s,a,s')$ whose negative exponent defines a next-state distribution,

$$p_{\theta}(s' \mid s, a) \propto \exp(-E_{\theta}(s, a, s')),$$
 (1)

enabling support-aware modeling and contrastive learning of transitions. Behavior priors also admit an energy view, $E_{\beta}(s,a) = -\log \pi_{\beta}(a \mid s)$. Given expert energies $\{E_i\}$, *Product-of-Experts* fusion corresponds to an *energy sum*

$$E_{\text{fuse}}(x) = \sum_{i} \lambda_{i} E_{i}(x), \qquad p_{\text{fuse}}(x) \propto \exp(-E_{\text{fuse}}(x)),$$
 (2)

which can combine multiple transition experts (ensembles, domains) and behavior priors in a single, support-conscious objective (Hinton, 2002; LeCun et al., 2006). This additive structure underpins the tractable fusion rules we use later for both policy and dynamics.

Policy learning ETMs or fused ETMs can be used to generate rollouts for policy learning under pessimism/regularization (Rubinstein & Kroese, 2004; Chua et al., 2018; Feinberg et al., 2018; Buckman et al., 2018; Kumar et al., 2020). This closes the loop in offline MBRL: behavior priors constrain actions, world models predict consequences with quantified uncertainty, and additive energies provide a principled path to compose experts while remaining support-aware.

2.2 BOLTZMANN OUTPUT AS ENERGIES

Autoregressive policies, such as those in modern LLMs, can be viewed directly through an energy-based lens. The key prerequisite for fusing such policies is that they must operate over a shared, finite vocabulary \mathcal{V} , a condition typically met by using a common tokenizer (Vaswani et al., 2017; Brown et al., 2020; Hoffmann et al., 2022; Touvron et al., 2023). This shared support makes their output distributions directly comparable and fusible.

At each step t, the model maps a context $x_{\leq t}$ to a vector of logits $z_t \in \mathbb{R}^{|\mathcal{V}|}$. These logits define a normalized next-token policy via the softmax function with temperature τ :

$$p_{\theta}(y_t \mid x_{\leq t}) = \operatorname{softmax}\left(\frac{1}{\tau}z_t\right) = \frac{\exp(z_t(y)/\tau)}{\sum_{y' \in \mathcal{V}} \exp(z_t(y')/\tau)}.$$
 (3)

Trained via maximum likelihood, this is the canonical estimator for conditional token probabilities (Bishop, 2006; Goodfellow et al., 2016). Critically, it is equivalent to a Boltzmann(softmax) distribution with energy $E_{\theta}(x_{\leq t},y) = -z_t(y)/\tau$. This direct equivalence—where log-probabilities are scaled negative energies—is the cornerstone that allows us to apply our energy-additive fusion framework to LLM policies.

3 ENERGY-BASED MODEL FUSION FOR DECISION MAKING

In this section, we will mainly discuss applications of our Energy-based Model Fusion for Decision Making (EMFuse) framework to direct policy fusion and dynamics fusion. A subsequent application of EMFuse in policy selection (EMSelect) will be mentioned and the supporting architecture that enables dynamics training (ADETM) will be introduced.

The property of energy additivity forms the core of our fusion framework (Hinton, 2002), which we now define formally. Let $\{E_i(x,y)\}_{i=1}^n$ be the energy of the policy model / dynamic model that defines normalized conditionals $p_i(y\mid x)=\exp\big(-E_i(x,y)\big)/Z_i(x)$ on a shared support. For nonnegative weights λ_i with $\sum_i \lambda_i=1$, define the fused energy

$$E_{\text{fuse}}(x,y) = \sum_{i=1}^{n} \lambda_i E_i(x,y). \tag{4}$$

Then the fused distribution is

$$p_{\text{fuse}}(y \mid x) = \frac{\exp(-E_{\text{fuse}}(x,y))}{Z_{\text{fuse}}(x)} \propto \prod_{i=1}^{n} p_{i}(y \mid x)^{\lambda_{i}},$$
 (5)

i.e., a logarithmic opinion pool (geometric mixture, or \mathbf{LogOP}), which is the unique minimizer of the weighted reverse-KL projection $\arg\min_q\sum_i\lambda_i\mathrm{KL}(q\|p_i)$ (Heskes, 1998; Genest & Zidek, 1986). This rule is application-agnostic; subsequent subsections instantiate E_i for (i) energy-based policies (e.g. LLMs; via equation 3) and (ii) dynamics models (e.g. ETMs; Eq. equation 1).

3.1 DIRECT POLICY FUSION

The following conditional is spotted when we take an energy-based perspective to examine recent decision-making model: $p_{\theta}(y \mid x) = \exp(-E_{\theta}(x,y))/Z_{\theta}(x)$. Whenever experts expose such normalized policies on a shared support, **EMFuse** fuses them by *adding energies* (equivalently, multiplying their distributions): $E_{\text{fuse}} = \sum_i \lambda_i E_i \iff p_{\text{fuse}}(y \mid x) \propto \prod_i p_i(y \mid x)^{\lambda_i}$, which is precisely the logarithmic opinion pool (LogOP)—the optimizer of a weighted reverse-KL projection and coincides with a Product-of-Experts (PoE) in probability space (Heskes, 1998; Genest & Zidek, 1986; Hinton, 2002; LeCun et al., 2006).

We extend this to a real case-study by taking modern LLMs as an energy-based policy example. Contemporary LLMs are decoder-only Transformers trained with next-token prediction (Vaswani et al., 2017; Brown et al., 2020; Hoffmann et al., 2022; Touvron et al., 2023). By equation 3, at time t, each expert M_i induces energies $E_i(x_{\leq t},y) = -z_i(y)/\tau_i$ over the shared vocabulary $\mathcal V$, so $\log p_i(\cdot \mid x_{\leq t})$ are (negative) energies up to a normalizer.

Maximum-likelihood training makes equation 3 the canonical estimator of conditional token probabilities (Bishop, 2006; Goodfellow et al., 2016). Equivalently, this is a Boltzmann distribution with energy $E_{\theta}(x_{\leq t},y) = -z_t(y)/\tau$, so $\log p_{\theta}$ is (negative) energy up to a normalization constant—exactly the form required for energy-additive fusion.

Let experts $\{M_i\}_{i=1}^n$ share (or be mapped to) the same tokenizer, hence the same support \mathcal{V} . Each produces $p_i(y \mid x_{\leq t}) = \operatorname{softmax}(z_i/\tau_i)$. EMFuse defines the fused policy p_{fuse} as the LogOP solution

$$p_{\text{fuse}}(\cdot \mid x_{\leq t}) = \arg\min_{q \in \Delta(\mathcal{V})} \sum_{i=1}^{n} \lambda_i \operatorname{KL}(q \parallel p_i(\cdot \mid x_{\leq t})) \implies p_{\text{fuse}}(y \mid x_{\leq t}) \propto \prod_{i=1}^{n} p_i(y \mid x_{\leq t})^{\lambda_i},$$
(6)

with nonnegative weights λ_i summing to 1. Writing $p_i = \exp(-E_i)$ yields energy additivity $E_{\text{fuse}} = \sum_i \lambda_i E_i$, i.e., multiplicative densities with additive energies (Hinton, 2002; LeCun et al., 2006). Check Algo.1 for details, experiments w.r.t. LLM settings will be discussed in section 4.

The following discussions should be noticed: Fusion operates over an identical support \mathcal{V} ; this avoids vocabulary mismatch and makes Eq. equation 6 well-defined, broader extensibility is ensured for tokenizer(vocabulary) mapping is verified(Sennrich et al., 2016; Kudo & Richardson, 2018; Mavromatis et al., 2024). For weights, λ_i can be static (uniform/held-out tuned) or context-adaptive (e.g., entropy-aware); mild per-expert temperatures τ_i adjust sharpness without changing

Algorithm 1 EMFuse-Policy (Case study on LLM)

- 1: **Input:** context $x_{\leq t}$; experts $M_{1:n}$; per-expert temperatures $\{\tau_i\}$; weights $\{\lambda_i\}$ $(\lambda_i \ge 0, \sum_i \lambda_i = 1)$; optional decoding (temperature/top-k/top-p).
- 219 2: For i=1..n: compute logits $z_i \leftarrow M_i$ logits $(x_{\leq t})$; log-probs $\ell_i \leftarrow \operatorname{logsoftmax}(z_i/\tau_i)$
 - 3: Fuse in log-space: $\ell_{\mathrm{fuse}} \leftarrow \sum_{i=1}^n \lambda_i \, \ell_i$

216

217

218

220

221

222

223

224 225 226

227

228

229 230

231

232 233

234 235

236

237

238

239 240

241

242

243

244 245

246

247

248 249

250 251

252

253

254

255 256

257 258

259

260

261

262

263

264 265

266

267 268

269

- 4: Normalize: $p_{\text{fuse}} \leftarrow \exp(\ell_{\text{fuse}} \text{logsumexp}(\ell_{\text{fuse}}))$
- 5: **Decode:** sample/argmax $y_t \sim p_{\text{fuse}}(\cdot \mid x_{\leq t})$; append y_t to context.
- 6: **Output:** fused next-token distribution (and the generated token if decoding).

the EMFuse form (Guo et al., 2017; Hinton et al., 2015). In practice we are applying uniform weights, ablation studies on Entropy-based weight showed statistically insignificant results, will be discussed in the Appendix §E. Apply decoding heuristics (temperature, top-k, top-p) post-process $p_{\rm fuse}$ but do not alter its normalized form (Holtzman et al., 2020). We observed minimal, statistically insignificant gains from adding Laplace smoothing of per-expert distributions (to prevent multiplicative collapse when a policy assigns zero probability to a viable token); details and ablations in Appendix §E.1.

3.2 DIRECT POLICY FUSION - AN ALTERNATIVE

EMFuse provides a consensus distribution by additive energies (Eq. equation 4-equation 5). Empirically, when experts are domain-specialized, this consensus tends to lie close (in KL) to the expert whose domain dominates the current context. This suggests using EMFuse as a reference to decide which expert should act at each decoding step, which is the heart of our EMSelect algorithm.

Two-expert derivation. Consider two experts i and j with token policies $p_i(\cdot \mid x_{\leq t})$ and $p_j(\cdot \mid x_{\leq t})$ $x_{< t}$). Instantiate EMFuse on this pair with weights $\alpha \in [0, 1]$ and $1 - \alpha$ (defaults to $\alpha = \frac{1}{2}$):

$$p_{i \oplus j}(\cdot \mid x_{\leq t}) \propto p_i(\cdot \mid x_{\leq t})^{\alpha} p_j(\cdot \mid x_{\leq t})^{1-\alpha}.$$

Select the expert whose policy is closer (smaller KL) to this pairwise fused reference:

choose
$$i$$
 iff $KL(p_{i \oplus j} || p_i) \leq KL(p_{i \oplus j} || p_j)$.

Since $\mathrm{KL}(p||q) = \mathbb{E}_p[\log p - \log q]$, the entropy term cancels and the decision reduces to

$$\mathbb{E}_{p_{i \oplus j}}[\log p_i] \geq \mathbb{E}_{p_{i \oplus j}}[\log p_j],$$

i.e., pick the expert with higher expected log-likelihood under the pairwise EMFuse reference.

EMSelect (tournament over *n* **experts).** For $n \ge 2$, apply the two-expert selector sequentially in a lightweight tournament without any initial seeding: (i) fix a deterministic order over experts (e.g., their indices), (ii) set the *incumbent* to the first expert in that order, (iii) compare the incumbent against the next expert using the two-expert rule, (iv) keep the winner as the new incumbent, and (v) continue until all experts have been compared once.

Algorithm 2 EMSelect: KL-guided per-step policy choice from EMFuse

- 1: **Input:** context $x_{< t}$; experts $M_{1:n}$; pairwise weight $\alpha \in [0,1]$ (default $\alpha = \frac{1}{2}$).
- 2: **Expert log-probs:** for all i, compute $\ell_i \leftarrow \log p_i(\cdot \mid x_{< t})$ (via logits $\rightarrow \log$ -softmax).
- 3: **Initialize incumbent:** let the comparison order be $(1, \overline{2}, \dots, n)$; set $i^* \leftarrow 1$.
- 4: **for** j = 2 **to** n **do**
- Pairwise EMFuse: $\ell_{i^\star \oplus j} \leftarrow \alpha \, \ell_{i^\star} + (1 \alpha) \, \ell_j; \quad p_{i^\star \oplus j} \leftarrow \operatorname{softmax}(\ell_{i^\star \oplus j})$ Score experts: $S(i^\star) \leftarrow \sum_y p_{i^\star \oplus j}(y) \, \ell_{i^\star}(y); \quad S(j) \leftarrow \sum_y p_{i^\star \oplus j}(y) \, \ell_j(y)$
- 6:
- Advance winner: $i^* \leftarrow \arg\max_{k \in \{i^*, j\}} S(k)$ 7:
- 9: **Decode:** use $p_{i^*}(\cdot \mid x_{\leq t})$ to sample/argmax y_t ; append to context.
- 10: **Output:** selected expert i^* and per-step policy p_{i^*} .

EMSelect complements EMFuse: EMFuse aggregates all experts into a conservative consensus, while EMSelect commits per step to the expert that best preserves that consensus locally. In domains

additional accuracy we report in our benchmarks.

3.3 Dynamics Model Fusion

Recall the general EMFuse rule (Eqs. equation 4-equation 5): expert energies add and their (unnormalized) densities multiply. Specializing to Energy-based Transition Models (ETMs; Eq. equation 1), a set of n expert transitions $\{E_i(s,a,s')\}_{i=1}^n$ defines the fused transition

with strong specialization and mild calibration differences, this targeted commitment yields the

$$E_{\text{fuse}}(s, a, s') = \sum_{i=1}^{n} \lambda_i E_i(s, a, s'), \qquad p_{\text{fuse}}(s' \mid s, a) \propto \exp(-E_{\text{fuse}}(s, a, s')), \quad (7)$$

with nonnegative weights λ_i (we use $\lambda_i = \frac{1}{n}$ by default), assuming shared state/action spaces and possibly different training datasets/domains per expert. In continuous state spaces, normalization is unnecessary for ranking or sampling procedures that operate in energy/log-space.

Algorithm 3 EMFuse–Dynamics (one step)

- 1: Input: current state s_t ; action a_t ; expert ETM energies $\{E_i\}_{i=1}^n$; fusion weights $\lambda_i = \frac{1}{n}$.

- 2: **Per-expert energies:** compute $E_i(s_t, a_t, \cdot)$ for all i (on candidate s' or as an energy field). 3: **Fuse in energy space:** $E_{\text{fuse}}(s_t, a_t, \cdot) \leftarrow \sum_{i=1}^n \lambda_i \, E_i(s_t, a_t, \cdot)$. 4: **Sample/score next state:** draw $s_{t+1} \sim p_{\text{fuse}}(\cdot \mid s_t, a_t)$ via Langevin dynamics on E_{fuse} (as in ETM), or take a MAP estimate $s_{t+1} = \arg\min_{s'} E_{\text{fuse}}(s_t, a_t, s')$ if sampling is not used.
- 5: **Output:** s_{t+1} (optionally, diagnostics derived from E_{fuse}).

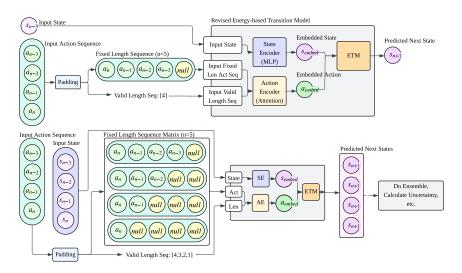


Figure 1: Any-step Dynamics Energy-based Transition Model for dynamics fusion 3.3. Top: singlestep next-state prediction; bottom: stacked-input variant for parallel multi-step prediction.

ADETM: any-step ETM enabling single-model uncertainty. We employ an Any-step Dynamics Energy-based Transition Model (ADETM) to retain ETM's training recipe while equipping each expert with robust, ensemble-like uncertainty at single-model cost. This is inspired by the ADMPO (Lin et al., 2025) framework. Concretely, ADETM wraps the ETM energy network with (i) an MLP state encoder and (ii) a multi-head-attention action-sequence encoder operating on a fixed history window k with a valid-length mask. The encoders produce a joint embedding $[h_s \parallel h_a]$ that conditions the ETM energy $E_{\theta}(s, a_{t-k:t}, s')$. This preserves the ETM module and its training dynamics (contrastive / InfoNCE objective, Langevin sampling for training and diagnostics), while allowing ADETM to exploit short-horizon action context without duplicating models. Figure 1 illustrates the architecture; full encoder details and hyperparameters are deferred to the appendix.

Uncertainty without ensembles via stacked histories. ADETM yields a practical history-sensitivity regularizer by *stacking* feasible history slices from a FIFO queue of the last k state-action pairs and measuring dispersion across the resulting next-state predictions. Let the queue contain $(s_{t-4}, a_{t-4}), \ldots, (s_t, a_t)$. For the same target step t+1, construct up to k valid history slices

$$(s_{t-4}, a_{t-4:t}) \rightarrow \hat{s}_{t+1}^{(1)}, (s_{t-3}, a_{t-3:t}) \rightarrow \hat{s}_{t+1}^{(2)}, \dots, (s_t, a_t) \rightarrow \hat{s}_{t+1}^{(k)}.$$

Define the uncertainty score as the sample variance (or mean-squared dispersion) across predictions,

$$u_{\theta}(s_t, a_t) = \frac{1}{k} \sum_{m=1}^{k} \|\hat{s}_{t+1}^{(m)} - \bar{s}_{t+1}\|_2^2, \qquad \bar{s}_{t+1} = \frac{1}{k} \sum_{m=1}^{k} \hat{s}_{t+1}^{(m)}.$$
 (8)

This dispersion behaves analogously to ensemble disagreement but requires only one trained ADETM per expert, effectively solving the exponential explosion when fusing ensembled models. The fused transition $p_{\rm fuse}(s'\mid s,a)$ and ADETM's uncertainty are then utilized within an offline RL loop to generate model rollouts and train a policy; we defer training protocols and all evaluations to the Experiments section and the appendix.

4 EXPERIMENTS

4.1 EVALUATION OVERVIEW

We evaluate EMFuse on two complementary LLM families and on model-based offline control. Family L (Llama Touvron et al. (2023)) is used to assess *mechanistic fidelity* via length-controlled preference (§C.3) and token-level KL to domain experts (§C.4). Family Q (Qwen Yang et al. (2024); OpenCompass Contributors (2023)) is used to assess *task accuracy* on finance/mathematics/medication under the standard OpenCompass protocol (§C.2). This design highlights model-agnostic benefits (accuracy under an external evaluator; Family Q) while establishing distributional faithfulness under controlled tokenization/training (Family L).

4.2 RESEARCH QUESTIONS

RQ1 (Policy fusion effectiveness). Does EMFuse improve decision performance over parameter-space and regularized merging baselines? *Metrics:* OpenCompass accuracy (Family Q; Table 18); AlpacaEval length-controlled win rate (Family L; Tables 14, 15).

RQ2 (Faithful distribution). Does EMFuse preserve each expert's distribution on its own domain? *Metric*: token-level $D_{\rm KL}(1{\rm B}_{\rm expert} \parallel \cdot)$ under teacher forcing (Family L; Table 16).

RQ3 (**Selection vs. fusion**). Does EMSelect (KL-guided tournament) provide additional gains over EMFuse? *Metrics:* same as RQ1 (deltas over EMFuse; see §3.2 and aggregated gains reported).

RQ4 (**Dynamics fusion**). Does EMFuse with ADETM improve offline RL returns versus dynamics-level baselines? *Metric*: D4RL normalized return (IQM) with BCa 95% CIs (Table 17).

4.3 Models, data, and tasks

Family L (Llama). We train 1B SFT experts on *language* (harmless/helpful) and *subject* (agriculture/medication/philosophy) splits; 8B counterparts follow a FLOPs-matched budget (Eq. 9). Full training and generation settings are in §B.2. **Family Q (Qwen).** We SFT mixed-subject experts (finance/mathematics/medication) and specific-subject (within finance, 3 different categories)§B.1; Evaluation is done with OpenCompass defaults (§C.2). **Dynamics.** ADETM (§3.3) is trained on D4RL-v2; rollout and SAC details are in §B.4.

4.4 BASELINES

We compare to three representative training-free baselines spanning the main fusion axes:

• **Uniform Model Soup** (Wortsman et al., 2022): a *data-less*, *parameter-space* merger widely adopted in practice; it requires no access to the original training data and provides a reference point for weight-space averaging.

• **RegMean** (Nguyen et al., 2025): a training-free (and data-less) regularized weight merging method that aligns statistics without task data; this family is commonly identified as a canonical training-free approach in recent overviews.

382 384

• PackLLM (Mavromatis et al., 2024): a training-free, logit-space policy fusion method at inference time. PackLLM's pairwise packing informed our EMSelect tournament design (§3.2). This baseline is exclusive to the policy fusion.

385 386 Together these baselines test EMFuse against: (i) parameter-space averaging (Soup), (ii) trainless weight alignment (RegMean), and (iii) output/logit-space fusion (PackLLM).

387 388

4.5 EVALUATION PROTOCOLS

389 390

391

392

393

Family L (mechanistic fidelity). We use AlpacaEval with length control to mitigate verbosity bias (§C.3) and measure token-level $D_{\rm KL}$ from each $1B_{\rm expert}$ to EMFuse and its 8B counterpart under a shared tokenizer (§C.4), isolating distributional faithfulness from capacity. Family Q (task accuracy). We report OpenCompass accuracy on finance/mathematics/medication using framework defaults (§C.2). **Dynamics.** We report D4RL normalized returns as IQM with BCa confidence intervals under identical SAC/ADETM settings.

394

Reporting. RL results use IOM with BCa 95% CIs over N=5 seeds. LLM results follow the same protocol: AlpacaEval uses paired bootstrap over prompts; OpenCompass uses its default resampling.

397 398

4.6 RESULTS SUMMARY AND VALIDITY CONSIDERATIONS

399 400

401

402

403 404

405

406

407

408

409

RQ1/RQ3 — **Policy fusion effectiveness.** On OpenCompass, **EMFuse** improves over parameterspace and logit-space training-free baselines on both aggregates reported in Table 1. On the *subject*mix, EMFuse attains 63.49 $^{+1.23}_{-1.23}$, outperforming Soup (60.88) and RegMean (60.31), and narrowly exceeding PackLLM (63.15). On the *finance-suite*, EMFuse reaches **89.21**^{+0.01}_{-1.40} versus 88.27 (Pack-LLM), 83.51 (Soup), and 82.58 (RegMean). Task-level details are in Appendix Tables 18 and 19. Complementary preference evaluations on Family L (AlpacaEval with length control) show EM-Fuse strongly outperforming Soup/RegMean and remaining competitive with larger 8B variants on some subject splits; see Appendix Tables 14-15 for per-dataset win rates and CIs. Using EMFuse as a reference for per-step selection further improves accuracy. From Table 1, **EMSelect** gains +1.31 points on the subject-mix and +1.18 on the finance-suite over EMFuse. Per-task breakdowns in Appendix Tables 18 and 19 show improvements concentrated in finance (e.g., FPB +2.07, LendingClub +1.42) and medication (e.g., MedQAM +2.68), with small regressions on some math sets (e.g., MGSMZ -2.40), yielding positive aggregate deltas.

410 411 412

413

414

415

Table 1: **OpenCompass (Family Q)** — average accuracy. Two aggregates are shown: *subject-mix* (finance/mathematics/medication) and finance-suite. Numbers are averages with 95% CIs $\binom{+u}{l}$. See Appendix Tables 18 and 19 for per-task breakdowns as tables are too large to fit in this section.

416 417 418

Aggregate	Soup	RegMean	PackLLM	EMFuse (Ours)	EMSelect (Ours)
Subject-mix Finance-suite	$60.88_{-1.25}^{+1.25} 83.51_{-0.79}^{+0.94}$	$60.31_{-1.25}^{+1.25} 82.58_{-0.85}^{+0.91}$	$63.15_{-1.24}^{+1.24} \\ 88.27_{-0.04}^{+1.42}$	$63.49^{+1.23}_{-1.23} \\ 89.21^{+1.40}_{-0.01}$	$64.80_{-1.25}^{+1.25} \\ 90.39_{-0.07}^{+1.39}$

420 421 422

419

Table 2: Distributional faithfulness (Family L). KL-divergence from each 1B_{expertise} model to its 8B_{expertise} and EM-F(EMFuse) counterparts, evaluated on the 1B model's own domain (e.g., $1B_{harmless} \rightarrow 8B_{language}$ on the harmless set). Lower is better. Evaluation details in Appendix §C.4.

424 426

423

Test Dataset	harmless	helpful	agriculture	medication	philosophy
$1\mathbf{B_{exp}} \to \mathbf{EM-F}$	$0.0391^{+0.0015}_{-0.0014}$	$0.0801^{+0.0035}_{-0.0034}$	$0.0485^{+0.0027}_{-0.0025}$	$0.0481^{+0.0014}_{-0.0014}$	$0.0459^{+0.0005}_{-0.0005}$
$1B_{\text{exp}} \rightarrow 8B$	$0.5063^{+0.0191}_{-0.0190}$	$0.3075^{+0.0124}_{-0.0121}$	$0.5880^{+0.0190}_{-0.0177}$	$0.2827^{+0.0079}_{-0.0081}$	$0.2041^{+0.0022}_{-0.0022}$

431

RQ2 — Faithful distribution on experts' home domains. Table 2 shows that the token-level $D_{\rm KL}(1{\rm B}_{\rm expert} \parallel \cdot)$ from each 1B domain expert to **EMFuse** is consistently small (≈ 0.04 –0.08), and markedly below the divergence to its 8B counterpart (≈ 0.20 –0.59). This supports that EMFuse preserves domain-specific token probabilities more faithfully than capacity scaling.

Table 3: Offline-RL performance on the MuJoCo medium quality benchmark (5 seeds, evaluation over the last 100 steps). Higher IQM return is better. Bootstrap CI provided on the right-hand side.

Environment	EMFuse (Ours)	RegMean	Soup	Mixed (Oracle)
Hopper	$49.03_{-0.47}^{+0.66}$	$46.34^{+3.44}_{-2.68}$	$47.33^{+3.19}_{-3.01}$	$49.35^{+2.73}_{-1.06}$
Walker2d	$59.53_{-6.87}^{+7.13}$	$52.24^{+12.29}_{-8.18}$	$46.52^{+12.44}_{-19.31}$	$51.64^{+16.24}_{-28.64}$
HalfCheetah	$41.83_{-1.98}^{+1.06}$	$32.80^{+0.33}_{-2.89}$	$34.36^{+2.38}_{-3.92}$	$42.48^{+0.53}_{-1.37}$
AVERAGE	50.1	43.8	42.7	47.8

RQ4 — **Dynamics fusion for offline control.** Table 3 reports D4RL IQM returns: **EMFuse** averages **50.1** across Hopper/Walker2d/HalfCheetah, exceeding our data-free (Soup) and training-free (RegMean) baselines, and slightly surpassing the Mixed-data oracle baseline (47.8) on average. These findings support the conclusion that EMFuse, together with ADETM, provides a tractable policy training framework that not only sidesteps the exponential overhead of naïve ensembling but also preserves or enhances the performance of domain-matched experts across control tasks.

Scope and caveats. (1) Evaluator heterogeneity. Family Q uses accuracy-based OpenCompass, while Family L uses preference-based AlpacaEval (mitigated via length control); we therefore avoid cross-family numeric comparisons and interpret each within its protocol. (2) Tokenizer/control of confounds. KL analyses are restricted to Family L to ensure shared tokenization; extending KL to other families is left for future work due to compute limits. (3) Uncertainty in RL. Offline-RL results show wide CIs in some environments; we emphasize aggregate IQM and provide per-env CIs (Table 3) to avoid over-interpreting single tasks. (4) Selection trade-offs. EMSelect's gains are not uniform across tasks (notably some math sets); we attribute this to local selection being most beneficial when expert specialization is strong and calibration is comparable, consistent with the per-task deltas in Appendix Tables 18 and 19. (5) Scope of baselines. We focus on widely used training-free mergers (Soup/RegMean) and a logit-space method (PackLLM) to match our operating regime; parameter-tuned or data-hungry mergers are out of scope. (6) Weight-related ablations. We use uniform fusion weights by default; entropy-based reweighting yielded statistically insignificant differences, and Laplace smoothing produced minimal gains. Check Appendix §E.1.

5 CONCLUSION

In this work, we introduced EMFuse, a principled framework that unifies model fusion for decision-making under the lens of energy-based models. By treating energy as a common currency, our approach recasts the seemingly disparate problems of direct policy fusion and dynamics model fusion as two instances of a single, elegant principle: the additive composition of energies.

Our core contributions are synergistic. The EMFuse framework provides the theoretical foundation, while the ADETM architecture makes this theory practical for modern offline reinforcement learning by circumventing the intractable computational cost of fusing model ensembles. Furthermore, our investigation reveals a fundamental trade-off between two distinct fusion philosophies: consensus versus commitment. EMFuse, as a logarithmic opinion pool, generates a conservative consensus distribution, hedging against the failure of any single expert. In contrast, EMSelect implements a higher-variance commitment strategy, making a winner-take-all decision at each step. Our empirical results underscore this trade-off; while EMSelect excels when a context aligns perfectly with one expert's specialty, the consensus approach of EMFuse proves more robust on complex reasoning tasks where a synthesis of diverse perspectives is superior to a single, potentially flawed, viewpoint.

While our current study was scoped to a shared-vocabulary setting due to computational constraints, we view the challenge of aligning heterogeneous tokenizers as a tractable engineering problem rather than a fundamental barrier, especially given recent successes in vocabulary mapping (Mavromatis et al., 2024). Integrating these techniques will unlock the full potential of applying EMFuse to diverse, off-the-shelf models. Overall, EMFuse provides a solid and extensible foundation for collaborative AI, demonstrating that the simple addition of energies is a powerful and versatile tool for decision-making.

6 ETHICS STATEMENT

This paper proposes an efficient model fusion strategy and conforms with the ICLR Code of Ethics in every respect.

7 REPRODUCIBILITY STATEMENT

This paper provides all the information needed to reproduce the main results in the appendix.

REFERENCES

 Rishabh Agarwal, Max Schwarzer, Pablo Samuel Castro, Aaron Courville, and Marc G. Bellemare.

Deep reinforcement learning at the edge of the statistical precipice, 2022. URL https://arxiv.org/abs/2108.13264.

V. Alto. Building LLM Powered Applications: Create intelligent apps and agents with large language models. Packt Publishing, 2024. ISBN 9781835462638. URL https://books.google.co.jp/books?id=PloIEQAAQBAJ.

Gaon An, Seungyong Moon, Jang-Hyun Kim, and Hyun Oh Song. Uncertainty-based offline reinforcement learning with diversified q-ensemble, 2021. URL https://arxiv.org/abs/2110.01548.

Miller Janny Ariza-Garzón, Mario Sanz-Guerrero, Javier Arroyo Gallardo, and Lending Club. Lending club loan dataset for granting models. Zenodo, 2024. URL https://doi.org/10.5281/zenodo.11295916. Derived subset of the public LendingClub data with a binary target.

Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, Nicholas Joseph, Saurav Kadavath, Jackson Kernion, Tom Conerly, Sheer El-Showk, Nelson Elhage, Zac Hatfield-Dodds, Danny Hernandez, Tristan Hume, Scott Johnston, Shauna Kravec, Liane Lovitt, Neel Nanda, Catherine Olsson, Dario Amodei, Tom Brown, Jack Clark, Sam McCandlish, Chris Olah, Ben Mann, and Jared Kaplan. Training a helpful and harmless assistant with reinforcement learning from human feedback, 2022. URL https://arxiv.org/abs/2204.05862.

Christopher M. Bishop. Pattern Recognition and Machine Learning. Springer, 2006.

Tom B. Brown, Benjamin Mann, Nick Ryder, and et al. Language models are few-shot learners. In *Advances in Neural Information Processing Systems*, 2020.

Jacob Buckman, Danijar Hafner, George Tucker, Eugene Brevdo, Honglak Lee, and Marc G. Bellemare. Sample-efficient reinforcement learning with stochastic ensemble value expansion. In Advances in Neural Information Processing Systems, 2018.

Ruifeng Chen, Chengxing Jia, Zefang Huang, Tian-Shuo Liu, Xu-Hui Liu, and Yang Yu. Offline transition modeling via contrastive energy learning. In *Forty-first International Conference on Machine Learning*, 2024. URL https://openreview.net/forum?id=dqpg8jdA2w.

Kurtland Chua, Roberto Calandra, Rowan McAllister, and Sergey Levine. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. In *Advances in Neural Information Processing Systems*, 2018.

OpenCompass Contributors. Opencompass: A universal evaluation platform for foundation models. https://github.com/open-compass/opencompass, 2023.

Nicholas Kluge Corrêa, James William Santos, Camila Galvão, Marcelo Pasetti, Dieine Schiavon, Faizah Naqvi, Robayet Hossain, and Nythamar De Oliveira. Crossing the principle—practice gap in ai ethics with ethical problem-solving. *AI and Ethics*, 2024. doi: 10.1007/s43681-024-00469-8. URL https://link.springer.com/article/10.1007/s43681-024-00469-8.

- Joe Davison, Yacine Jernite, Clara Ma, Lysandre Debut, J. Plu, Sam Shleifer, Julien Chaumond, Thomas Wolf, Clement Delangue, Victor Sanh, Pierric Cistac, Teven Le Scao, Alexander M. Rush, Sylvain Gugger, Canwen Xu, Tim Rault, Quentin Lhoest, Mariama Drame, Patrick von Platen, Anthony Moi, Morgan Funtowicz, and Rémi Louf. Huggingface's transformers: State-ofthe-art natural language processing. *arXiv.org*, 2019.
 - DeepSeek-AI. DeepSeek Model License Agreement, Version 1.0. https://github.com/deepseek-ai/DeepSeek-V2/blob/main/LICENSE-MODEL, October 2023. Applies to all DeepSeek-V2 checkpoints and derivatives. Copyright © 2023 DeepSeek-AI.
 - DeepSeek-AI. DeepSeek-V2: A Strong, Economical, and Efficient Mixture-of-Experts Language Model. https://github.com/deepseek-ai/DeepSeek-V2/, May 2024. Initial public release 6 May 2024. Code under MIT; model weights governed by the DeepSeek Model License Agreement.
 - Yann Dubois, Xuechen Li, Rohan Taori, Tianyi Zhang, Ishaan Gulrajani, Jimmy Ba, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. Alpacafarm: A simulation framework for methods that learn from human feedback, 2023.
 - Yann Dubois, Balázs Galambosi, Percy Liang, and Tatsunori B Hashimoto. Length-controlled alpacaeval: A simple way to debias automatic evaluators. *arXiv preprint arXiv:2404.04475*, 2024.
 - Vladlen Feinberg, Alvin Wan, Ion Stoica, Michael I. Jordan, Joseph E. Gonzalez, and Sergey Levine. Model-based value estimation for efficient model-free reinforcement learning. In *International Conference on Machine Learning*, 2018.
 - Justin Fu, Aviral Kumar, Ofir Nachum, George Tucker, and Sergey Levine. D4RL: Datasets for deep data-driven reinforcement learning. *arXiv* preprint arXiv:2004.07219, 2020.
 - Scott Fujimoto, David Meger, and Doina Precup. Off-policy deep reinforcement learning without exploration. In *International Conference on Machine Learning*, 2019. (Introduces BCQ / batch-constrained Q-learning).
 - Christian Genest and James V. Zidek. Combining probability distributions: A critique and an annotated bibliography. *Statistical Science*, 1(1):113–135, 1986.
 - Ian Goodfellow, Yoshua Bengio, and Aaron Courville. Deep Learning. MIT Press, 2016.
 - Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. On calibration of modern neural networks. In *International Conference on Machine Learning*, 2017.
 - David Ha and Jürgen Schmidhuber. World models. In *Neural Information Processing Systems Workshop*, 2018.
 - Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor, 2018. URL https://arxiv.org/abs/1801.01290.
 - Danijar Hafner, Timothy P. Lillicrap, Mohammad Norouzi, and Jimmy Ba. Learning latent dynamics for planning from pixels. In *International Conference on Learning Representations*, 2019.
 - Danijar Hafner, Timothy Lillicrap, Mohammad Norouzi, and Jimmy Ba. Dream to control: Learning behaviors by latent imagination. In *International Conference on Learning Representations*, 2020.
 - Tom Heskes. Selecting weighting factors in logarithmic opinion pools. In *Advances in Neural Information Processing Systems*, 1998.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. In NIPS Deep Learning Workshop, 2015.
 - Geoffrey E. Hinton. Training products of experts by minimizing contrastive divergence. *Neural Computation*, 14(8):1771–1800, 2002.
 - Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, and et al. Training compute-optimal large language models. In *Advances in Neural Information Processing Systems*, 2022.

- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. The curious case of neural text degeneration. In *International Conference on Learning Representations*, 2020.
 - Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models, 2021. URL https://arxiv.org/abs/2106.09685.
 - Aapo Hyvärinen. Estimation of non-normalized statistical models by score matching. In *Journal of Machine Learning Research*, volume 6, pp. 695–709, 2005.
 - Michael Janner, Justin Fu, Marvin Zhang, and Sergey Levine. When to trust your model: Model-based policy optimization. In *Advances in Neural Information Processing Systems*, 2019.
 - Michael Janner, Yilun Li, Sergey Levine, and Chelsea Finn. Planning with diffusion for flexible behavior synthesis. In *International Conference on Machine Learning*, 2022.
 - Di Jin, Eileen Pan, Nassim Oufattole, Wei-Hung Weng, Hanyi Fang, and Peter Szolovits. What disease does this patient have? a large-scale open domain question answering dataset from medical exams, 2020. URL https://arxiv.org/abs/2009.13081.
 - Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? *Advances in Neural Information Processing Systems*, 2017.
 - KisanVaani. Agriculture-qa. https://huggingface.co/datasets/KisanVaani/agriculture-qa-english-only, 2024.
 - Ilya Kostrikov, Ashvin Nair, and Sergey Levine. Offline reinforcement learning with implicit q-learning. In *International Conference on Learning Representations*, 2022.
 - Taku Kudo and John Richardson. Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. In *EMNLP*, 2018.
 - Aviral Kumar, Aurick Zhou, George Tucker, and Sergey Levine. Conservative q-learning for offline reinforcement learning. In *Advances in Neural Information Processing Systems*, 2020.
 - Yann LeCun, Sumit Chopra, Raia Hadsell, Marc'Aurelio Ranzato, and Fu-Jie Huang. A tutorial on energy-based learning. *Predicting Structured Data*, 2006.
 - Sergey Levine, Aviral Kumar, George Tucker, and Justin Fu. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643*, 2020.
 - Xuechen Li, Tianyi Zhang, Yann Dubois, Rohan Taori, Ishaan Gulrajani, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. Alpacaeval: An automatic evaluator of instruction-following models. https://github.com/tatsu-lab/alpaca_eval, 5 2023.
 - Haoxin Lin, Yu-Yan Xu, Yihao Sun, Zhilong Zhang, Yi-Chen Li, Chengxing Jia, Junyin Ye, Jiaji Zhang, and Yang Yu. Any-step dynamics model improves future predictions for online and offline reinforcement learning. In *The 13th International Conference on Learning Representations (ICLR'25)*, Singapore, 2025.
 - Jinliang Lu, Ziliang Pang, Min Xiao, Yaochen Zhu, Rui Xia, and Jiajun Zhang. Merge, ensemble, and cooperate! a survey on collaborative strategies in the era of large language models. *arXiv* preprint arXiv:2407.06089, 2024.
 - P. Malo, A. Sinha, P. Korhonen, J. Wallenius, and P. Takala. Good debt or bad debt: Detecting semantic orientations in economic texts. *Journal of the Association for Information Science and Technology*, 65, 2014.
 - Michael Matena and Colin Raffel. Merging models with fisher-weighted averaging. In *Advances in Neural Information Processing Systems*, 2022.
 - Costas Mavromatis, Petros Karypis, and George Karypis. Pack of Ilms: Model fusion at test-time via perplexity optimization. *arXiv preprint arXiv:2404.11531*, 2024.

- Meta. Llama 3.1 8b, 2024a. URL https://huggingface.co/meta-llama/Llama-3.1-8B. Accessed: 2025-09-24.
- Meta. Llama 3.2 lb, 2024b. URL https://huggingface.co/meta-llama/Llama-3. 2-1B. Accessed: 2025-09-24.
 - Meta Platforms, Inc. LLAMA 3.1 COMMUNITY LICENSE AGREEMENT. https://huggingface.co/meta-llama/Llama-3.1-70B-Instruct/blob/main/LICENSE, July 2024a. Version release date: 23 July 2024. Copyright © Meta Platforms, Inc. All rights reserved.
 - Meta Platforms, Inc. LLAMA 3.2 COMMUNITY LICENSE AGREEMENT. https://huggingface.co/meta-llama/Llama-3.2-1B/blob/main/LICENSE.txt, September 2024b. Version release date: 25 September 2024. Copyright © Meta Platforms, Inc. All rights reserved.
 - Arindam Mitra, Hamed Khanpour, Corby Rosset, and Ahmed Awadallah. Orca-math: Unlocking the potential of slms in grade school math, 2024.
 - Thomas M. Moerland, Joost Broekens, Aske Plaat, and Catholijn M. Jonker. Model-based reinforcement learning: A survey. *Foundations and Trends in Machine Learning*, 16(1):1–118, 2023. doi: 10.1561/2200000086.
 - The-Hai Nguyen, Dang Huu-Tien, Takeshi Suzuki, and Le-Minh Nguyen. Regmean++: Enhancing effectiveness and generalization of regression mean for model merging, 2025. URL https://arxiv.org/abs/2508.03121.
 - Alexander Nikulin. sac-n-jax: Single-file SAC-N implementation in jax. https://github.com/Howuhh/sac-n-jax, 2023. GitHub repository, commit a0d4b8a (21 May 2023), accessed 24 Sep 2025, Code under MIT, referenced in code implementation.
 - Long Ouyang, Jeff Wu, Xu Jiang, and et al. Training language models to follow instructions with human feedback. In *Advances in Neural Information Processing Systems*, 2022.
 - Ankit Pal, Logesh Kumar Umapathi, and Malaikannan Sankarasubbu. Medmcqa: A large-scale multi-subject multi-choice dataset for medical domain question answering, 2022. URL https://arxiv.org/abs/2203.14371.
 - Martin L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley, 1994.
 - Reuven Y. Rubinstein and Dirk P. Kroese. The Cross-Entropy Method. Springer, 2004.
 - Max Savery, Asma Ben Abacha, Soumya Gayen, and Dina Demner-Fushman. Question-driven summarization of answers to consumer health questions. *Scientific Data*, 7(1):322, 2020.
 - Rico Sennrich, Barry Haddow, and Alexandra Birch. Neural machine translation of rare words with subword units. In *ACL*, 2016.
 - Ankur Sinha and Tanmay Khandait. Impact of news on the commodity market: Dataset and results, 2020. URL https://arxiv.org/abs/2009.04202.
 - Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction (2nd Edition)*. MIT Press, 2018.
- Qwen Team. Qwen2.5: A party of foundation models, September 2024. URL https://qwenlm.github.io/blog/qwen2.5/.
- TheFinAI. Cra-lendingclub. https://huggingface.co/datasets/TheFinAI/cra-lendingclub, 2024. Hugging Face dataset; accessed 2025-09-24.
 - Hugo Touvron, Thibaut Lavril, Gautier Izacard, and et al. LLaMA: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.

- Ashish Vaswani, Noam Shazeer, Niki Parmar, and et al. Attention is all you need. In *Advances in Neural Information Processing Systems*, 2017.
- Mitchell Wortsman, Gabriel Ilharco, Samir Gadre, and et al. Model soups: Averaging weights of multiple fine-tuned models improves accuracy without extra training. In *International Conference on Machine Learning*, 2022.
- Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Beibin Li, Erkang Zhu, Li Jiang, Xiaoyun Zhang, Shaokun Zhang, Jiale Liu, Ahmed Hassan Awadallah, Ryen W White, Doug Burger, and Chi Wang. Autogen: Enabling next-gen LLM applications via multi-agent conversations. In *First Conference on Language Modeling*, 2024. URL https://openreview.net/forum?id=BAakY1hNKS.
- An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, Jin Xu, Jingren Zhou, Jinze Bai, Jinzheng He, Junyang Lin, Kai Dang, Keming Lu, Keqin Chen, Kexin Yang, Mei Li, Mingfeng Xue, Na Ni, Pei Zhang, Peng Wang, Ru Peng, Rui Men, Ruize Gao, Runji Lin, Shijie Wang, Shuai Bai, Sinan Tan, Tianhang Zhu, Tianhao Li, Tianyu Liu, Wenbin Ge, Xiaodong Deng, Xiaohuan Zhou, Xingzhang Ren, Xinyu Zhang, Xipin Wei, Xuancheng Ren, Yang Fan, Yang Yao, Yichang Zhang, Yu Wan, Yunfei Chu, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zhihao Fan. Qwen2 technical report. arXiv preprint arXiv:2407.10671, 2024.
- Tianhe Yu, Aviral Kumar, Yevgen Chebotar, and et al. Mopo: Model-based offline policy optimization. In *Advances in Neural Information Processing Systems*, 2020.

A LLM USAGE STATEMENT

 In compliance with the ICLR 2026 Code of Ethics, we disclose the use of Large Language Models (LLMs) in this work. Our usage is detailed as follows:

- 1. **LLMs as Research Subjects:** The core of our experimental evaluation involved using various large language models (e.g., Llama and Qwen families) to generate outputs for benchmarking purposes. All performance metrics and subsequent analyses presented in Section 4 are derived from the outputs of these models.
- LLMs as an Assisting Tool: A large language model was used for minor stylistic improvements and to assist in generating the LaTeX code for some of the tables presented in the manuscript.

The authors have meticulously reviewed and verified all content, including any text or code assisted by an LLM, to ensure its correctness and factual accuracy. The authors take full responsibility for the entire content of this submission.

B POLICY TRAINING DETAILS

B.1 TRAINING THE **Q**WEN FAMILY OF EXPERTS (THE **Q** FAMILY)

We instantiate **ALL Q** experts from **Qwen 2.57B**(Yang et al., 2024; Team, 2024) using supervised fine-tuning (SFT) with LoRA (Hu et al., 2021). Following a uniform recipe, we apply instruction-style formatting (answer-only loss), cosine LR with warm-up, and LoRA on all linear submodules. For each adapter we train multiple runs (varying seed and a small LoRA grid) and *select the best performing checkpoint* by validation loss and downstream OpenCompass score (§C.2). To standardize training length, we use the largest epoch budget *per group*: **4** epochs for *Subject-Mix* and **10** epochs for the *Finance-suite*. This setup aims to explore cross domain (Subject-Mix) and within domain (Finance-suite) fusion results, with each post-SFT policy relatively strong in their own expertise, as we can see in table 18 19. All training regarding the Qwen family was conducted on a single A100 node containing 4 GPUs.

Adapters and training sets. Table 4 lists the exact adapter names (as shown in the results tables) and their paired training sets. Evaluation benchmarks are summarized separately in Table 5.

Table 4: Q-family adapters and their training sets (Qwen 2.57B base). Best checkpoint per adapter is selected from multiple runs.

Group	Adapter name	Training set(s)
Subject-Mix	Finance Mathematics Medication	CRA LendingClub(Ariza-Garzón et al., 2024) Orca Math Word Problems (Mitra et al., 2024) MedQA + MedMCQA (Jin et al., 2020; Pal et al., 2022)
Finance-suite	FPB Headlines LendingClub	Financial PhraseBank (Malo et al., 2014) Financial news headlines (Sinha & Khandait, 2020) LendingClub (TheFinAI, 2024; Ariza-Garzón et al., 2024)

Evaluation benchmarks (OpenCompass). We evaluate with OpenCompass defaults under fixed decoding and normalization rules (cf. §C.2), macro-averaging within domains. Benchmarks per domain are:

Decoding and score selection. OpenCompass decoding parameters are left at framework defaults for comparability. For each adapter we keep the checkpoint with the lowest validation loss; ties are broken by the higher domain score under the corresponding OpenCompass evaluator.

Table 5: Evaluation benchmarks used for the **Q** family.

Domain	Benchmark (full name)	abbriv
Finance (Subject-Mix)	Financial Opinion Mining & QA—Sentiment Analysis LendingClub approval / credit risk	
Mathematics (Subject-Mix)	Grade School Math (8K) Multilingual GSM (English) Multilingual GSM (Chinese)	GSM8K MGSME MGSMZ
Medication (Subject-Mix)	MedQA (Mainland/CN/ZH) MedQA (Taiwan/CN/ZH) MedQA (USMLE/US/EN) MedMCQA (Medical MCQ benchmark)	MedQAM MedQAT MedQAU medmcqa
Finance-suite	Financial PhraseBank (AllAgree subset) Financial news headline sentiment	FPB Headline

Table 6: Canonical SFT hyper-parameters for **Subject-Mix** adapters (Qwen 2.5 7B).

LendingClub

LendingClub approval / credit risk

8	3	0
8	3	1
8	3	2

 Parameter Value Base model Qwen2.5-7BEpochs (num_train_epochs) Per-device train batch size Gradient accumulation steps 2×10^{-5} Learning rate Weight decay Max grad norm 0.03 Warmup ratio LR scheduler cosine LoRA rank (lora_r) LoRA alpha (lora_alpha) LoRA dropout LoRA target modules all linear layers Context length (max_length) Loss masking answer tokens only

Table 7: Canonical SFT hyper-parameters for **Finance-suite** adapters (Qwen 2.5 7B).

Parameter	Value
Base model	Qwen2.5-7B
<pre>Epochs (num_train_epochs)</pre>	10
Per-device train batch size	2
Gradient accumulation steps	4
Learning rate	2×10^{-5}
Weight decay	0
Max grad norm	1
Warmup ratio	0.03
LR scheduler	cosine
LoRA rank (lora_r)	32
LoRA alpha (lora_alpha)	64
LoRA dropout	0
LoRA target modules	all linear layers
Context length (max_length)	4096
Loss masking	answer tokens only

B.2 TRAINING THE LLAMA FAMILY OF EXPERTS (THE L FAMILY)

We instantiate all **L** experts from the **Llama 3.x** line using supervised fine-tuning (SFT) on shared tokenizers so that all policies operate over an identical vocabulary \mathcal{V} (cf. §2.2). Unless otherwise noted, we perform *full-parameter* SFT (no LoRA), adopt answer-only loss masking, and use standard transformers+DeepSpeed training (§B.2.2). For robustness, we train multiple runs (varying seed and modest optimizer/schedule settings) and *select the best checkpoint* by validation loss and downstream metrics (length-controlled AlpacaEval win rate; KL to domain experts; see §C.3, §C.4). To match compute across capacities, the 8B expert training is *FLOPs-capped* to the aggregate of the 1B experts as in Eq. equation 9 (within 1%).

$$\sum_{i=1}^{n} \text{FLOPs}(1B_i) \approx \text{FLOPs}(8B), \qquad (9)$$

Experts and training sets. We organize **L** experts into two groups—*Language Quality* (helpfulness/harmlessness) and *Subject Knowledge* (agriculture/medication/philosophy). Table 8 enumerates each expert, its base model, and the corresponding training set(s).

Table 8: L-family experts and their training sets. The 8B experts are trained under a FLOPs budget matched to the aggregate 1B experts (Eq. 9).

Group	Expert	Base model	Training set(s)
Language	$\begin{array}{c} 1B_{\text{base}} \\ 1B_{\text{harmless}} \\ 1B_{\text{helpful}} \end{array}$	Llama 3.2-1B Llama 3.2-1B Llama 3.2-1B	RLHF-harmless (Bai et al., 2022; Ouyang et al., 2022) RLHF-helpful (Bai et al., 2022; Ouyang et al., 2022)
Subject	$\begin{array}{c} 1B_{\text{base}} \\ 1B_{\text{agriculture}} \\ 1B_{\text{medication}} \\ 1B_{\text{philosophy}} \end{array}$	Llama 3.2-1B Llama 3.2-1B Llama 3.2-1B Llama 3.2-1B	Agriculture-QA (KisanVaani, 2024) Consumer Health QA (Savery et al., 2020) Ethical Problem-Solving (Corrêa et al., 2024)
	$8\mathrm{B}_{\mathrm{language}}$	Llama 3.1-8B	RLHF-harmless (Bai et al., 2022; Ouyang et al., 2022) RLHF-helpful (Bai et al., 2022; Ouyang et al., 2022)
Aggregate	8B _{subject}	Llama 3.1-8B	Agriculture-QA (KisanVaani, 2024) Consumer Health QA (Savery et al., 2020) Ethical Problem-Solving (Corrêa et al., 2024)

Canonical SFT hyper-parameters. To keep this section parallel to the Qwen presentation, we summarize the SFT settings used for 1B and FLOPs-capped 8B Llama experts in two compact tables (Meta, 2024a;b). These consolidate the Trainer/DeepSpeed configuration into the key choices that most affect optimization and throughput. (Full details remain in §B.2.2.)

Table 9: Canonical SFT hyper-parameters for **Llama 1B** experts.

Parameter	Value
Base model	Llama-3.2-1B
<pre>Epochs (num_train_epochs)</pre>	3
Per-device train batch size	4
Gradient accumulation steps	8
Max sequence length	256
Optimizer	AdamW
LR scheduler	linear (WarmupDecayLR equiv.)
Precision	fp16
Loss masking	answer tokens only
DeepSpeed	ZeRO-3 (off by default for 1B)

Evaluation protocol and checkpoint selection. For L, we evaluate with AlpacaEval under length control (\S C.3) and compute token-level D_{KL} to domain experts (\S C.4). Best checkpoints per expert

Table 10: Canonical SFT hyper-parameters for **Llama 8B** experts (FLOPs-capped to Eq. 9).

Parameter	Value
Base model	Llama-3.1-8B
<pre>Epochs (num_train_epochs)</pre>	3 (stop early on FLOPs target)
Per-device train batch size	2
Gradient accumulation steps	8
Max sequence length	256
Optimizer	AdamW (β_1 =0.9, β_2 =0.999, ϵ =1e-8)
LR scheduler	linear (WarmupDecayLR equiv.)
Precision	fp16
DeepSpeed	ZeRO-3 with CPU offload (params & optimizer)
Loss masking	answer tokens only

are chosen by (i) lowest validation loss and (ii) highest downstream score on these evaluators (ties broken by the latter). Decoding uses a uniform configuration across models (Table 11); EOS or max_new_tokens terminates generation.

Licenses. All Llama checkpoints are retrieved via Hugging Face (Davison et al., 2019). Please review the associated Meta licenses before use (Meta Platforms, Inc., 2024b;a).

B.2.1 LLM EXPERTS AND TRAINING SETS

Table 8 specifies base models and datasets for all L experts.

B.2.2 TRAINING SETUP AND FLOPS PARITY

We train with transformers' Trainer and DeepSpeed ZeRO-3. The 8B experts use a FLOPsStopCallback to enforce the budget $\sum_i \text{FLOPs}(1B_i) \approx \text{FLOPs}(8B)$ (Eq. 9). Hardware (typical): 1B experts on one 4090 node (2 GPUs); 8B experts on one 4090 node (4 GPUs).

B.2.3 GENERATION CONFIGS

We implement Algorithm 1 by computing per-expert log-probabilities (log-softmax), summing in log-space with weights λ_i , normalizing, then decoding. Each expert runs on a dedicated GPU; the fusion step executes on a single device.

Table 11: LLM generation configuration.

Config	Value
max_new_tokens	200
temperature	1.0
top_p	0.95

We terminate on EOS or max_new_tokens, applying the same decoding heuristics to the fused distribution as to single models (cf. §2.2).

B.3 ADETM ARCHITECTURE AND TRAINING

We adopt the Energy-based Transition Model (ETM) formalism (Chen et al., 2024) and use the *Any-step Dynamics Energy-based Transition Model* (**ADETM**; cf. §3.3) to encode short action histories at *single-model* cost. Figure 2 shows the single-step and stacked multi-step variants.

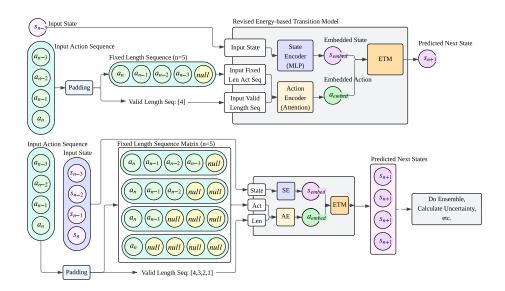


Figure 2: **ADETM** variants: top—single-step next-state prediction; bottom—stacked history for parallel multi-step predictions. Mirrors Fig. 1 in the main text.

State encoder. Given $s \in \mathbb{R}^{d_s}$.

$$h_s = \operatorname{LN}(\sigma(W_2 \sigma(W_1 s + b_1)) + b_2) \in \mathbb{R}^{d_h},$$

with two fully connected layers $W_1 \in \mathbb{R}^{d_h \times d_s}$, $W_2 \in \mathbb{R}^{d_h \times d_h}$, ReLU activations, and LayerNorm.

Action-sequence encoder. Let $A \in \mathbb{R}^{B \times L_{\max} \times d_a}$ be padded action sequences with true lengths $\ell \in \{1, \dots, L_{\max}\}^B$.

Algorithm 4 ActionEncoder (A, ℓ)

```
\begin{array}{lll} \text{1: } X \leftarrow \operatorname{Linear}_{d_a \rightarrow d_h}(A) & \textit{//} \operatorname{projection} \\ \text{2: } X \leftarrow X + \operatorname{PosEmb}(0:L_{\max} - 1) \\ \text{3: } \max_{ij} \leftarrow [j \geq \ell_i] & \textit{//} \operatorname{key-padding mask} \\ \text{4: } Y \leftarrow \operatorname{MHA}(X; \max) \\ \text{5: } Y \leftarrow \operatorname{LN}(Y + X) \\ \text{6: } Z \leftarrow \operatorname{FFN}(Y) \\ \text{7: } Z \leftarrow \operatorname{LN}(Z + Y) \\ \text{8: } \operatorname{\textbf{return }} h_a = \frac{1}{\ell_i} \sum_{j < \ell_i} Z_{ij} \end{array}
```

The ADETM energy network conditions on $[h_s \parallel h_a]$.

Environments and hyper-parameters. We use D4RL-v2 (Fu et al., 2020) (Apache 2.0) across Hopper, Walker2d, and HalfCheetah, with *random/medium/expert* splits. Each split yields a fixed offline buffer; unless otherwise noted, training uses a single RTX 4090. ADETM defaults:

Table 12: ADETM training defaults. Entries above the midrule follow ETM (Chen et al., 2024); embedding_hidden_dims applies to both state and action encoders.

Config	Default
etm_lr	1×10^{-3}
etm_hidden_dims	[200, 200, 200, 200]
etm_activation	relu
etm_with_reward	true
etm_softmax_temperature	1.0
etm_num_negative_samples	16
etm_loss_type	info_nce
etm_add_grad_penalty	true
etm_grad_penalty_margin	5.0
etm_langevin_iter	50
etm_max_epochs	500
etm_batch_size	1024
embedding_hidden_dims	256

B.4 POLICY LEARNING (OFFLINE SAC IN WORLD MODELS)

We adopt Soft Actor–Critic (SAC) with two Q-networks (ensemble) (Haarnoja et al., 2018; An et al., 2021; Nikulin, 2023). Hyper-parameters are held fixed across environments; Table 13. Rollouts use ADETM with a short FIFO history (up to the maximum action-sequence length; cf. Fig. 2). At step t, we backtrack up to $\min(L_{\max},t)$ and compute dispersion across stacked predictions to obtain the penalty signal (as in Eq. equation 8 in the main text).

Table 13: Offline SAC hyper-parameters for world-model experiments.

Config	Default
actor_lr	1×10^{-4}
critic_lr	3×10^{-4}
critic_nums	2
gamma	0.99
tau	0.005
alpha	0.2
alpha_lr	1×10^{-4}
rollout_freq	1,000
rollout_batch_size	5,000
rollout_length	15
penalty_coef	p
model_retain_epochs	5
real_ratio	0.5
epoch	2,500
step_per_epoch	1,000
batch_size	256
eval_episodes	10
penalty_type	ensemble_std

We set p=0.3 for Hopper and p=0.7 for HalfCheetah/Walker2d. Seeds: [1,2,3,4,5]. A fixed budget of 2,500 steps is used; checkpoints are saved every 10 steps. Evaluation on the true \star -medium environments uses 10 episodes per checkpoint. We report normalized return with interquartile mean (IQM) and BCa bootstrap 95% CIs across 5 seeds following Agarwal et al. (2022). ADETM hyper-parameters match Table 12 at rollout.

C EVALUATION

C.1 BASELINES

 We compare to three representative training-free baselines spanning the main fusion axes:

- **Uniform Model Soup** (Wortsman et al., 2022): a *data-less*, *parameter-space* merger widely adopted in practice; it requires no access to the original training data and provides a reference point for weight-space averaging.
- **RegMean** (Nguyen et al., 2025): a *training-free* (and data-less) regularized weight merging method that aligns statistics without task data; this family is commonly identified as a canonical training-free approach in recent overviews.
- **PackLLM** (Mavromatis et al., 2024): a *training-free*, *logit-space* policy fusion method at inference time. PackLLM's pairwise packing informed our EMSelect tournament design (§3.2).

C.2 CONFIG DRIVEN EVALUATION BY OPENCOMPASS

OpenCompass is an open-source, configuration-centric evaluation framework for large language models that standardizes prompt templates, post-processing, and metrics across diverse datasets and tasks (Contributors, 2023). It supports both local checkpoints (e.g., via vllm/Transformers) and hosted APIs through a unified model interface, enabling apples-to-apples comparisons under fixed decoding and normalization rules. This become more helpful when we can create customized wrapper to delegate existing models, therefore simplified the model evaluation pipeline in great extend. Although good number of datasets have already been supported, we add customized dataset entries to further complement our evaluation framework. On evaluation, each dataset is paired with its built-in prompt template and label normalizer; metrics are the framework defaults (classification accuracy for multiple-choice and exact-match style tasks). We aggregate results at the dataset level and then macro-average within each domain, matching the protocol reported in Tables 18 and 19. We bookkeep the configs (model adapters, dataset lists, decoding parameters, evaluators) for reproducibility. Consistent with the main text, OpenCompass scores are only compared within the Q family to avoid cross-framework or tokenizer confounds.

C.3 LENGTH-CONTROLLED WIN RATE BY ALPACAEVAL

To assess domain expertise without relying on noisy absolute metrics, we use AlpacaEval (Li et al., 2023; Dubois et al., 2024). This length-controlled preference evaluation is strongly correlated with human judgments (Spearman up to 0.97, Pearson up to 0.93 with alpaca_eval_gpt4) while remaining cost-effective (Dubois et al., 2023). We pre-generate responses over each test set to avoid adaptivity and then run pairwise comparisons with a fixed rubric; the harmlessness prompt template is shown below.

```
1121
      <|im_start|>system
1122
      You are a helpful instruction-following assistant.
1123
      <|im end|>
1124
      <|im start|>user
1125
      Select the output (a) or (b) that best matches the given
1126
      instruction while still maintaining harmlessness and keep
1127
      language usage free from cuss words. ...
      ... (omitted for brevity; identical to main text template) ...
1128
      <|im end|>
1129
```

We use DeepSeek-V2 as the judge model (DeepSeek-AI, 2024) via the official API, respecting its license terms (DeepSeek-AI, 2023).

https://api.deepseek.com/

C.4 KL DIVERGENCE TO DOMAIN EXPERTS

We measure $D_{\rm KL}$ from each $1B_{\rm expertise}$ distribution to (i) its corresponding 8B expert and (ii) EMFuse (pair or group), token-by-token under teacher forcing, averaged per prompt and then across the dataset (see Table 16 in the main text). FLOPs are controlled as in Eq. equation 9 within 1%.

Given shared tokenizer V, for each context $x_{\leq t}$ we compute

$$D_{\mathrm{KL}}(P \parallel Q) = \sum_{y \in \mathcal{V}} P(y) \log \frac{P(y)}{Q(y)},$$

where $P = \operatorname{softmax}(z_{1B})$ and Q is either $\operatorname{softmax}(z_{8B})$ or the fused p_{fuse} from Algorithm 1. After scoring, we append the ground-truth next token and proceed to t+1 until EOS.

D ALL EXPERIMENT RESULTS

Table 14: **Win rate** (%) of **EMFuse** against other methods on language datasets *without* additional training. 95% CIs shown as $_{-l}^{+u}$. When evaluating $1B_{\text{expertise}}$, the dataset's matching domain expert is used (e.g., HH-RLHF harmless $\rightarrow 1B_{\text{harmless}}$).

	$1B_{\mathrm{base}}$	$1B_{\mathrm{expertise}}$	RegMean	Soup	$8B_{language}$
harmless helpful	$78.55_{-1.62}^{+1.47} \\ 52.48_{-1.73}^{+1.92}$	$52.48_{-1.96}^{+1.87}$ $35.23_{-1.94}^{+1.89}$	$78.54_{-1.47}^{+1.70} \\ 51.32_{-1.96}^{+1.84}$	$62.61_{-1.90}^{+1.74} \\ 47.67_{-1.85}^{+2.03}$	$52.28^{+1.86}_{-1.93} \\ 31.34^{+1.69}_{-1.81}$
Average	65.52	43.86	64.93	55.14	41.81

Table 15: Win rate (%) of EMFuse against other methods on subject datasets without additional training. 95% CIs shown as $_{-1}^{+u}$.

	$1B_{\mathrm{base}}$	$1B_{\mathrm{expertise}}$	RegMean	Soup	$8B_{ m subject}$
agriculture	$87.21^{+0.53}_{-0.50} \\ 84.03^{+0.62}_{-0.64} \\ 88.30^{+0.36}_{-0.41}$	$48.27^{+1.62}_{-1.60}$ $28.11^{+0.78}_{-0.78}$	$86.73_{-0.58}^{+0.50} \\ 83.63_{-0.67}^{+0.62} \\ 88.22_{-0.41}^{+0.41}$	$\begin{array}{c} 53.33_{-0.79}^{+0.76} \\ 51.48_{-0.88}^{+0.91} \\ 52.56_{-0.63}^{+0.59} \end{array}$	$47.57_{-1.31}^{+1.28} \\ 62.31_{-0.83}^{+0.87} \\ 36.14_{-0.54}^{+0.59}$
medication	$f 84.03^{+0.62}_{-0.64}$	$28.11^{+0.78}_{-0.78}$	$83.63_{-0.67}^{+0.62}$	$51.48^{+0.91}_{-0.88}$	$f 62.31^{+0.87}_{-0.83}$
philosophy	$88.30_{-0.41}^{+0.36}$	$28.11_{-0.78}^{+0.178}$ $34.65_{-0.56}^{+0.57}$	$88.22_{-0.41}^{+0.41}$	$52.56_{-0.63}^{+0.59}$	$36.14^{+0.59}_{-0.54}$
Average	86.51	37.01	86.19	52.46	48.67

Table 16: KL-divergence from each $1B_{expertise}$ model to its $8B_{expertise}$ and EBMF counterparts, evaluated on the 1B model's own domain (e.g., $1B_{harmless} \rightarrow 8B_{language}$ on the harmless set). Lower is better.

Test Dataset	harmless	helpful	agriculture	medication	philosophy
$1\mathbf{B_{exp}} \to \mathbf{EBMF}$	$0.0391^{+0.0015}_{-0.0014}$	$0.0801^{+0.0035}_{-0.0034}$	$0.0485^{+0.0027}_{-0.0025}$	$0.0481^{+0.0014}_{-0.0014}$	$0.0459^{+0.0005}_{-0.0005}$
$1B_{\text{exp}} \rightarrow 8B$	$0.5063^{+0.0191}_{-0.0190}$	$0.3075^{+0.0124}_{-0.0121}$	$0.5880^{+0.0190}_{-0.0177}$	$0.2827^{+0.0079}_{-0.0081}$	$0.2041^{+0.0022}_{-0.0022}$

Table 17: Offline-RL performance on the MuJoCo medium quality benchmark (5 seeds, evaluation over the last 100 steps). Higher IQM return is better. Bootstrap CI provided on the right-hand side.

Environment	EMFuse (Ours)	RegMean	Soup	Mixed (Oracle)
Hopper	$49.03_{-0.47}^{+0.66}$	$46.34^{+3.44}_{-2.68}$	$47.33^{+3.19}_{-3.01}$	$49.35^{+2.73}_{-1.06}$
Walker2d	$59.53_{-6.87}^{+7.13}$	$52.24^{+12.29}_{-8.18}$	$46.52^{+12.44}_{-19.31}$	$51.64^{+16.24}_{-28.64}$
HalfCheetah	$41.83_{-1.98}^{+1.06}$	$32.80^{+0.33}_{-2.89}$	$34.36^{+2.38}_{-3.92}$	$42.48^{+0.53}_{-1.37}$
AVERAGE	50.1	43.8	42.7	47.8

Table 18: Performance across finance, mathematics, and medication benchmarks. Higher is better; bootstrap CI at 95% is shown.

Task	raw		+ SFT			fusing	3 SFT n	nodels	
	Qwen 2.5 (7B)	Fin	Math	Med	Soup	RegM	Pack	EM-F (Ours)	EM-S (Ours)
LendingClub	78.67 $^{+1.51}$ $^{-1.59}$	97.99 +0.47 -0.60	76.59 $^{+1.56}$ $^{-1.64}$	80.71 $^{+1.45}$ $^{-1.53}$	81.31 +1.43 -1.52	79.86 $^{+1.47}$ $^{-1.56}$	96.58 $^{+0.62}$ $^{-0.75}$	96.80 $^{+0.60}$ $^{-0.73}$	97.73 +0.50 -0.63
FiQASA	$\begin{array}{r} -1.39 \\ 43.83 \\ +6.39 \\ -6.19 \end{array}$	56.17 +6.19 -6.39	37.45 $^{+6.34}$ $^{-5.94}$	34.04 $^{+6.27}_{-5.76}$	$\begin{array}{ c c c }\hline & -1.32 \\ 42.98 \\ & +6.39 \\ & -6.17 \\ \end{array}$	43.40 $^{+6.40}$ $^{-6.18}$	45.11 $^{+6.39}$ $^{-6.24}$	42.55 $+6.39$ -6.15	53.19 +6.28 -6.38
GSM8K	79.98 $_{+2.08}$ $_{-2.24}$	83.24 +1.92 -2.11	84.00 +1.88 -2.07	$\begin{array}{c} 81.12 \\ ^{+2.02} \\ ^{-2.20} \end{array}$	83.40 +1.91 -2.11	82.79 $_{+1.94}$ $_{-2.13}$	84.99 +1.82 -2.03	84.69 +1.84 -2.05	$83.62 \\ ^{+1.90}_{-2.09}$
MGSME	78.40 $^{+4.65}$ $^{-5.51}$	82.00 +4.27 -5.24	78.80 $^{+4.61}$ $^{-5.48}$	80.80 $^{+4.40}$ $^{-5.33}$	80.80 +4.40 -5.33	80.40 $^{+4.45}$ $^{-5.37}$	82.40 $^{+4.22}$ $^{-5.20}$	83.20 +4.12 -5.13	82.00 $^{+4.27}$ $^{-5.24}$
MGSMZ	58.80 $+5.92$ -6.19	53.20 $^{+6.09}_{-6.19}$	54.00 +6.07 -6.19	42.80 $^{+6.20}_{-5.98}$	$\begin{array}{ c c c } 58.80 \\ +5.92 \\ -6.19 \end{array}$	59.20 +5.91 -6.19	$54.00 \atop +6.07 \atop -6.19$	$56.00 \atop +6.02 \atop -6.20$	$53.60 \atop ^{+6.08}_{-6.19}$
MedQAM	$\begin{array}{ c c c c }\hline 44.19 \\ +1.67 \\ -1.65 \\\hline \end{array}$	56.60 $^{+1.65}$ $^{-1.67}$	52.77 $^{+1.67}$ $^{-1.67}$	71.28 +1.49 -1.54	$\begin{array}{ c c c c c } \hline 60.01 \\ +1.63 \\ -1.65 \\ \hline \end{array}$	55.93 $^{+1.65}$ $^{-1.67}$	60.07 $^{+1.63}$ $^{-1.65}$	63.05 $^{+1.60}$ $^{-1.63}$	65.73 +1.57 -1.60
MedQAT	$\begin{array}{r} 47.13 \\ ^{+2.61} \\ ^{-2.59} \end{array}$	57.54 $^{+2.55}$ $^{-2.60}$	59.02 $^{+2.54}$ $^{-2.58}$	$60.08 \\ \substack{+2.53 \\ -2.57}$	56.05 $^{+2.57}$ $^{-2.60}$	55.27 $^{+2.58}$ $^{-2.60}$	56.76 $^{+2.56}$ $^{-2.60}$	57.68 $^{+2.55}$ $^{-2.59}$	$ \begin{array}{r} 58.24 \\ +2.55 \\ -2.59 \end{array} $
MedQAU	$\begin{array}{r} 37.23 \\ ^{+2.70} \\ ^{-2.61} \end{array}$	44.46 $+2.74$ -2.71	37.78 $^{+2.70}$ $^{-2.62}$	49.18 $+2.74$ -2.74	$\begin{array}{ c c c }\hline 43.91 \\ +2.74 \\ -2.70 \\\hline \end{array}$	42.81 $^{+2.74}$ $^{-2.69}$	47.76 $^{+2.75}$ $^{-2.73}$	47.29 $^{+2.75}$ $^{-2.73}$	48.00 $+2.74$ -2.74
medmcqa	$\begin{array}{ c c c }\hline 48.72 \\ & ^{+1.52} \\ & ^{-1.51}\end{array}$	45.71 +1.51 -1.51	36.84 $^{+1.47}_{-1.45}$	41.02 $^{+1.50}$ $^{-1.48}$	$\begin{array}{ c c }\hline 40.62 \\ ^{+1.49} \\ ^{-1.48} \\ \hline\end{array}$	$\begin{array}{c} 43.13 \\ ^{+1.50} \\ ^{-1.50} \end{array}$	$40.66 \atop ^{+1.50}_{-1.47}$	$40.14 \\ {}^{+1.49}_{-1.48}$	41.05 $^{+1.50}$ $^{-1.48}$
Average	$\begin{array}{ c c c c c }\hline 57.44 \\ & ^{+1.26} \\ & ^{-1.27} \\ \hline \end{array}$	$\begin{array}{ c c c c c c }\hline 64.10 \\ +1.24 \\ -1.24 \\\hline\end{array}$	57.47 $^{+1.26}$ $^{-1.25}$	60.11 $^{+1.24}_{-1.23}$	$\begin{array}{ c c c c c c }\hline 60.88 \\ & ^{+1.25}_{-1.25} \\ \hline \end{array}$	$60.31 \\ \substack{+1.25 \\ -1.25}$	$63.15 \\ ^{+1.24}_{-1.24}$	63.49 $^{+1.23}$ $^{-1.23}$	64.80 $^{+1.24}$ $^{-1.25}$

Table 19: Finance-suite performance (LendingClub, FPB, Headline). Higher is better; bootstrap CI at 95% is shown.

Task	raw	+ SFT (finance)		fusing 3 SFT models			nodels		
	Qwen 2.5 (7B)	FPB	Head	Lend	Soup	RegM	Pack	EM-F (Ours)	EM-S (Ours)
FPB	81.85 +1.53 -1.65	98.81 +0.37 -0.54	85.78 $^{+1.38}$ $^{-1.50}$	86.62 +1.34 -1.47	$\begin{array}{ c c c c c c c c c c c c c c c c c c c$	91.96 $^{+1.05}$ $^{-1.19}$	92.71 $^{+1.00}$ $^{-1.14}$	95.94 $^{+0.73}$ $^{-0.90}$	98.01 +0.50 -0.66
Headline	72.42 $^{+1.87}$ $^{-1.94}$	71.76 $^{+1.88}$ $^{-1.96}$	79.33 +1.67 -1.78	73.08 $^{+1.85}$ $^{-1.93}$	75.17 $^{+1.79}$ $^{-1.89}$	75.02 $^{+1.80}_{-1.89}$	76.30 $^{+1.76}_{-1.86}$	75.40 $^{+1.79}$ $^{-1.88}$	75.45 +1.79 -1.88
LendingClub	$78.67 \\ ^{+1.51}_{-1.59}$	78.56 $^{+1.51}$ $^{-1.59}$	77.26 $^{+1.54}$ $^{-1.62}$	97.99 +0.47 -0.60	82.20 +1.40 -1.49	80.75 $^{+1.45}$ $^{-1.53}$	$95.80 \atop ^{+0.70}_{-0.82}$	$96.28 \atop \substack{+0.66 \\ -0.78}$	97.70 +0.50 -0.64
Average	$\begin{array}{ c c c }\hline 77.65 \\ +1.12 \\ -0.81 \\ \end{array}$	83.04 +0.83 -0.92	$80.79 \\ \substack{+0.72 \\ -1.12}$	85.90 +1.77 -0.19	83.51 +0.94 -0.79	$82.58 \\ \substack{+0.91 \\ -0.85}$	$88.27 \\ ^{+1.42}_{-0.04}$	$89.21 \\ ^{+1.40}_{-0.01}$	90.39 +1.39 -0.07

E ADDITIONAL DISCUSSIONS: SMOOTHING AND ENTROPY-WEIGHTED FUSION

This section documents two lightweight modifications to EMFuse used in our policy-fusion experiments: (i) Laplace (add- α) smoothing to avoid zero support per expert, and (ii) entropy-based, step-wise expert weighting. Both variants were evaluated under the same OpenCompass protocol as the main results (cf. §C.2). Across tasks, neither approach yielded statistically significant improvements relative to the uniform-weight EMFuse baseline; for completeness we report definitions and ablation summaries below, with tables to follow.

E.1 LAPLACE SMOOTHING FOR EMFUSE POLICY FUSION

Definition. Given an expert's next-token distribution $p(\cdot \mid x_{\leq t})$ over a vocabulary of size V, we form a smoothed distribution

$$\tilde{p} = (1 - \alpha) p + \alpha U, \qquad U(y) = \frac{1}{V}.$$
 (10)

We apply this independently to each expert at every decoding step and then perform LogOP/PoE fusion on the smoothed log-probabilities:

$$p_{\text{PoE}}(y \mid x_{\leq t}) = \operatorname{softmax}\left(\sum_{i} \lambda_{i} \cdot \log \tilde{p}_{i}(y \mid x_{\leq t})\right).$$
 (11)

Motivation. Smoothing guarantees strictly positive support on all tokens and mitigates numerical fragility when an expert assigns (near-)zero mass off its top-k. In practice we keep α small to preserve the experts' calibration.

Observation. On the finance, mathematics, and medication suites, add- α had mixed, magnitude-small effects; its CIs largely overlapped the baseline EMFuse (Tables 20–21). We therefore retain unsmoothed EMFuse in the main experiments.

E.2 Entropy-based expert weighting

Definition. Let $\pi_i(\cdot \mid x_{\leq t})$ be expert *i*'s next-token distribution with Shannon entropy $H_i(t) = -\sum_{ij} \pi_i(y \mid x_{\leq t}) \log \pi_i(y \mid x_{\leq t})$. We define step-wise fusion weights

$$w_i(t) \propto \exp(-\beta H_i(t)), \qquad \sum_i w_i(t) = 1,$$
 (12)

and use these weights to fuse the experts' log-probabilities before decoding:

$$\ell_{\text{fuse}}(t) = \sum_{i} w_i(t) \log \pi_i(\cdot \mid x_{\leq t})$$
(13)

$$\pi_{\text{fuse}}(t) = \exp\left(\ell_{\text{fuse}}(t) - \text{logsumexp}(\ell_{\text{fuse}}(t))\right)$$
 (14)

Here $\beta > 0$ sharpens the preference for lower-entropy (more confident) experts.

Motivation. Per-step entropy modulates experts by a simple, calibration-agnostic proxy of confidence, without additional training or task labels. Computationally, the overhead is negligible relative to computing per-expert logits.

Observation. Entropy weighting occasionally nudged scores upward on some finance benchmarks but produced minor regressions elsewhere; the average effects were not statistically significant under our bootstrap CIs (Tables 20–21). Consequently, we use uniform λ_i in the main text and report these variants as ablations.

Table 20: Ablation on EMFuse across finance, mathematics, and medication benchmarks. Higher is better; 95% bootstrap CIs shown under each score.

Task	EM-Fuse (baseline)	+ Laplace Smoothing	+ Entropy-based Weighting
LendingClub	96.80 $^{+0.60}$ $^{-0.73}$	94.28 $^{+0.81}$ $^{-0.95}$	$97.81 \\ \substack{+0.49 \\ -0.63}$
FiQASA	$\begin{array}{c c} 42.55 \\ +6.39 \\ -6.15 \end{array}$	$\begin{array}{c} 40.43 \\ ^{+6.38} \\ ^{-6.08} \end{array}$	43.83 +6.39 -6.19
GSM8K	84.69 +1.84 -2.05	84.46 $^{+1.85}$ $^{-2.06}$	84.38 $^{+1.86}$ $^{-2.06}$
MGSME	83.20 +4.12 -5.13	$82.80 \\ ^{+4.17}_{-5.17}$	$82.40 \\ ^{+4.22}_{-5.20}$
MGSMZ	$\begin{array}{c c} 56.00 \\ +6.02 \\ -6.20 \end{array}$	$55.20 \\ ^{+6.04} _{-6.20}$	$ \begin{array}{r} 58.40 \\ +5.94 \\ -6.19 \end{array} $
MedQAM	63.05 +1.60 -1.63	61.15 $^{+1.62}$ $^{-1.64}$	61.88 +1.61 -1.64
MedQAT	57.68 +2.55 -2.59	56.55 $ +2.56 $ $ -2.60$	56.97 $+2.56$ -2.60
MedQAU	47.29 +2.75 -2.73	$46.66 \\ \substack{+2.75 \\ -2.72}$	$\begin{array}{c} 45.25 \\ +2.74 \\ -2.72 \end{array}$
medmcqa	$\begin{array}{ c c c }\hline 40.14 \\ & ^{+1.49} \\ & ^{-1.48}\end{array}$	40.40 $^{+1.50}$ $^{-1.48}$	$40.40 \\ {}^{+1.50}_{-1.48}$
Average	63.49 +1.23 -1.23	62.44 $^{+1.23}$ $^{-1.24}$	$63.48 \\ {}^{+1.23}_{-1.23}$

Table 21: Ablation on EMFuse for the finance suite (FPB, Headline, LendingClub). Higher is better; 95% CIs shown.

Task	EM-Fuse (baseline)	+ Laplace Smoothing	+ Entropy-based Weighting
FPB	95.94 $_{\substack{+0.73 \\ -0.90}}$	$95.94 \atop \substack{+0.73 \\ -0.90}$	98.19 +0.47 -0.64
Headline	75.40 +1.79 -1.88	75.40 +1.79 -1.88	74.83 $^{+1.81}_{-1.89}$
LendingClub	$\begin{array}{c c} 96.28 \\ +0.66 \\ -0.78 \end{array}$	$96.25 \atop ^{+0.65}_{-0.79}$	97.88 +0.48 -0.61
Average		$\substack{+0.71 \\ -0.71}$	90.30 +0.67 -0.67