

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 WEAVING IN THE CLOUDS: ACHIEVING SYNERGISTIC COLLABORATION AMONG LLM AGENTS VIA FEDER- ATED LEARNING

006 **Anonymous authors**

007 Paper under double-blind review

010 ABSTRACT

013 Multi-Agent Systems (MAS) powered by Large Language Models (LLMs) have
014 shown immense potential in solving complex, sequential tasks by simulating ex-
015 pert collaboration. However, their reliance on centralized data clashes with real-
016 world privacy constraints and data silos. Conversely, existing privacy-preserving
017 paradigms like Federated Learning (FL) typically ignore the inherent sequential
018 dependencies present in collaborative workflows, leading to suboptimal perfor-
019 mance. To bridge this critical gap, we introduce **FedWave**, a novel framework
020 for federated multi-agent collaboration. FedWave empowers LLM-based agents
021 to collaboratively solve complex sequential tasks under strict privacy constraints
022 by employing three core mechanisms: (1) a collaborative Value Chain Layer to
023 model sequential dependencies, enabling efficient local fine-tuning through Fed-
024 erated Learning with LoRA adapters; (2) an intelligent Mixture of Experts (MoE)
025 router at the server level for dynamic, task-aware aggregation of expert knowl-
026 edge, moving beyond simple averaging; and (3) a final Direct Preference Opti-
027 mization (DPO) stage to align the model’s collaborative outputs with human pref-
028 erences. Extensive experiments demonstrate that FedWave significantly outper-
029 forms both traditional federated learning and centralized multi-agent baselines, ef-
030 fectively achieving synergistic collaboration without compromising data privacy.
031 The codes are available at <https://anonymous.4open.science/r/FedWave-111A>.

1 INTRODUCTION

034 Large Language Models (LLMs) have demonstrated remarkable capabilities across a diverse range
035 of tasks, enabling numerous innovative applications (Webb et al., 2023; Ouyang et al., 2022; Achiam
036 et al., 2023; Team, 2024). Among these, frameworks based on Multi-Agent Systems (MAS) have
037 become particularly prominent for solving complex real-world problems by simulating collabora-
038 tion among experts from different domains (Zhao et al., 2024; Qian et al., 2025; Li et al., 2024).
039 In scenarios such as business planning (Zhao et al., 2025), financial analysis (Yang et al., 2023),
040 and medical diagnostics (Tang et al., 2024a), MAS effectively model and handle sequential depen-
041 dencies across different roles, showing immense potential. However, the success of these systems
042 heavily relies on massive, centralized datasets (Wu et al., 2024). This data demand faces two ma-
043 jor challenges: first, high-quality public datasets are projected to be exhausted by 2026 (Villalobos
044 et al., 2022); second, vast high-quality data is distributed among different parties, forming “data
045 silos” due to data sovereignty and privacy concerns (Ye et al., 2024; Fan et al., 2023). Therefore,
046 although MAS excel at handling process dependencies, their centralized data assumption overlooks
047 the critical need for privacy preservation in the real world.

048 Federated Learning (FL) (McMahan et al., 2023) offers a viable solution to this issue. It establishes a
049 privacy-preserving, distributed collaborative framework that allows multiple parties to train a model
050 jointly without sharing their local data, thereby greatly facilitating the secure integration of data
051 value across institutions. Nevertheless, mainstream FL frameworks (Kuang et al., 2024; Yao et al.,
052 2024) often overlook the potential sequential dependencies or causal chains that may exist between
053 different data sources in their design. For instance, in a typical automotive industry chain, design
data precedes production data, supply chain data is tightly coupled with the production phase, and
quality inspection data serves as the downstream validation for the entire process. Existing federated

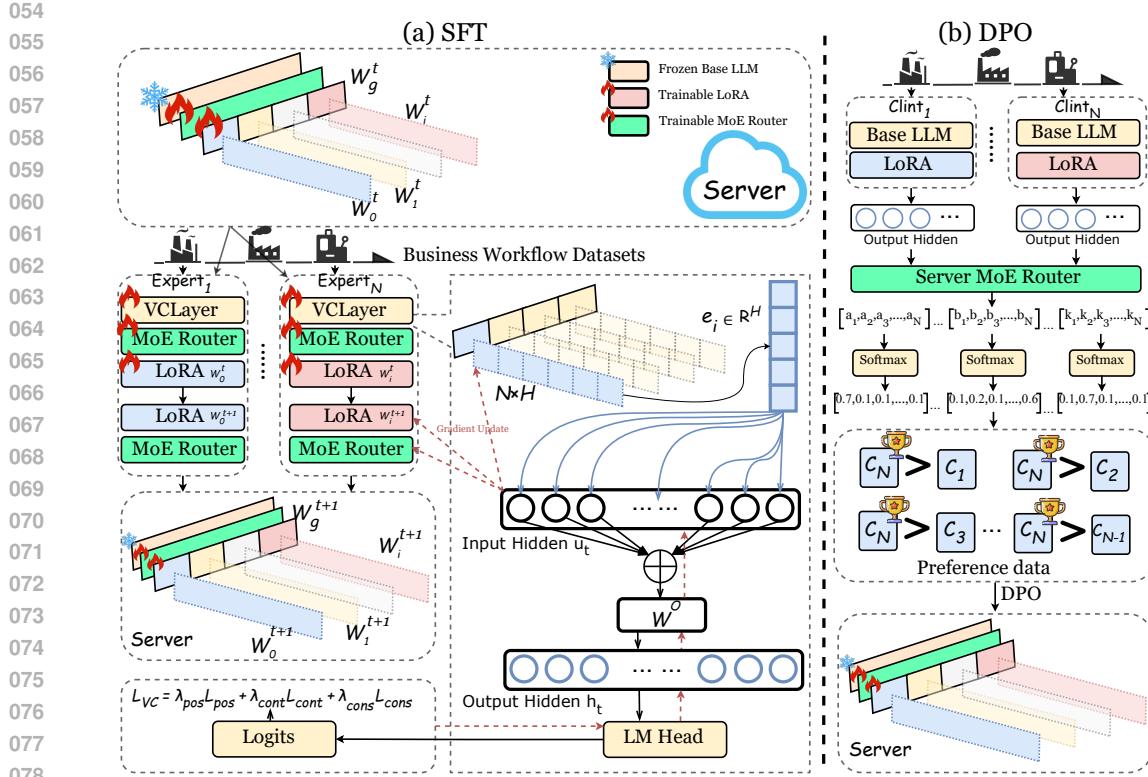


Figure 1: Overview of the FedWave framework. The framework consists of two main phases: (a) The Federated SFT phase, where individual expert agents collaboratively train trainable LoRA adapters, the Value Chain Layer (VCLayer), and a MoE Router on private data. The VCLayer models the workflow structure through a specialized loss function, while the MoE Router learns to dynamically coordinate experts and updates its gradients via backpropagation. Locally updated model parameters are sent to the server for aggregation. (b) The DPO alignment phase, which utilizes the trained MoE Router to automatically generate a preference dataset composed of responses from the most relevant (winning) and less relevant (losing) experts. The final global model is fine-tuned on this dataset using DPO to align its output with high-quality collaborative preferences.

aggregation strategies, typically based on the Independent and Identically Distributed (IID) data assumption (Li et al., 2020; Karimireddy et al., 2021; Hsu et al., 2019), fail to effectively model and leverage this process-aware knowledge embedded in the data, thus limiting the performance of the federated model in complex collaborative tasks with inherent logical sequences.

To address these challenges, we propose FedWave, a novel framework for federated multi-agent collaboration. FedWave empowers LLM-based agents for efficient collaboration while preserving data privacy, integrating distributed knowledge to solve complex sequential tasks. Its core mechanisms are as follows: First, at the client level, we combine FL with the Parameter-Efficient Fine-Tuning (PEFT) technique, LoRA. Each expert agent performs local LoRA-only fine-tuning on the base LLM, minimizing model updates sent to the server. This design significantly enhances communication efficiency and mitigates privacy leakage risks by reducing information exchange. Second, on the server, we designed an innovative Mixture of Experts (MoE) aggregation router. This router discards traditional FL aggregation, acting as an intelligent coordinator that learns to dynamically and selectively weave together knowledge from different expert agents based on the task. Finally, we introduce a Direct Preference Optimization (DPO) (Rafailov et al., 2023) stage to align the aggregated global model using preference data on collaborative outcomes. This ensures the model generates high-quality, collaborative outputs that align with human expectations. Extensive experiments demonstrate that our framework outperforms existing federated learning and multi-agent baselines.

108 The main contributions of this paper can be summarized as follows:
 109

- 110 • **Privacy-Preserving Collaboration for Workflow Tasks:** We introduce FedWave, the first
 111 framework enabling LLM agents to solve complex, sequential workflows across data silos.
 112 Its novel Value Chain Layer, within a Federated Learning paradigm, models inter-agent
 113 dependencies to facilitate structured, privacy-preserving collaboration.
- 114 • **Dynamic Knowledge Aggregation:** We designed a aggregation mechanism centered on
 115 MoE router. This shifts from static federated averaging to dynamic, task-aware knowledge
 116 fusion, allowing the global model to selectively leverage the most relevant expertise.
- 117 • **Extensive Experimental Validation:** Comprehensive experiments on business workflow
 118 datasets validate FedWave’s superiority. It significantly outperforms federated learning
 119 baselines and achieves competitive results against centralized multi-agent systems.

121 2 RELATED WORK

123 2.1 FEDERATED LEARNING

125 Federated Learning (FL) is a privacy-preserving paradigm for collaborative training on decentralized
 126 data (Kairouz et al., 2021). While the foundational FedAvg algorithm (McMahan et al., 2023) is
 127 effective, its performance often degrades with Non-IID data, prompting extensive research into solu-
 128 tions. These improvements typically mitigate data heterogeneity through client-side optimizations
 129 like FedProx (Li et al., 2020) and SCAFFOLD (Huang et al., 2024), or server-side aggregation re-
 130 finements such as FedNova (Wang et al., 2020) and FedAvgM (Hsu et al., 2019). However, existing
 131 works primarily address statistical heterogeneity, overlooking the critical issue of Sequential Depen-
 132 dency in process-based tasks. Their static aggregation methods treat clients as independent contrib-
 133 utors, failing to dynamically fuse knowledge based on the task’s inherent structure. In contrast, our
 134 FedWave framework fundamentally changes this paradigm by enabling the model to understand and
 135 leverage these sequential relationships among clients.

136 2.2 MULTI-AGENT COLLABORATION

138 The rise of Large Language Models (LLMs) has significantly advanced the development of Multi-
 139 Agent Systems, establishing them as a powerful paradigm for solving complex problems (Akata
 140 et al., 2023; Guo et al., 2024; Hao et al., 2023). These systems accomplish tasks by simulating
 141 collaboration among multiple agents, each with distinct roles or capabilities. Existing works have
 142 explored diverse collaboration models: some frameworks leverage discussion and debate (Du et al.,
 143 2024; Chen et al., 2024; Xiong et al., 2023) to refine and enhance reasoning abilities, while oth-
 144 ers construct hierarchical or sequential pipeline structures (Zhang et al., 2024; Zhao et al., 2024;
 145 2025). In these structures, agents process information progressively according to a predefined work-
 146 flow (such as a business workflow), effectively addressing complex sequential tasks in domains like
 147 long-text processing and medical diagnostics (Tang et al., 2024b; Sun et al., 2023). However, these
 148 frameworks’ reliance on centralized data access fundamentally conflicts with real-world, privacy-
 149 sensitive applications where data is siloed. This reveals a critical research gap: enabling complex,
 150 sequentially dependent agent collaboration in a decentralized, privacy-preserving setting. Our work,
 151 FedWave, directly addresses this challenge by merging the collaborative power of multi-agent sys-
 152 tems with the privacy guarantees of federated learning.

153 2.3 MIXTURE OF EXPERTS AND DYNAMIC AGGREGATION

155 The Mixture of Experts (MoE) architecture enhances model scalability through conditional compu-
 156 tation, where a router dynamically selects specialized ‘expert’ subnetworks for each input (Jacobs
 157 et al., 1991; Fedus et al., 2022). After its revival in modern deep learning (Shazeer et al., 2017), this
 158 paradigm has been instrumental in scaling Large Language Models (LLMs) to trillion-parameter
 159 scales while maintaining computational efficiency (Lepikhin et al., 2020; Jiang et al., 2024). This
 160 principle inspires Dynamic Aggregation in Federated Learning (FL), moving beyond the static, one-
 161 size-fits-all approach of methods like FedAvg (McMahan et al., 2017). While Personalized FL (PFL)
 adapts models for clients (Fallah et al., 2020; Arivazhagan et al., 2019; T Dinh et al., 2020; Li et al.,

2021), it primarily addresses statistical heterogeneity rather than collaborative, sequential relationships. However, a critical gap remains, as existing work has not integrated MoE’s dynamic routing into FL aggregation to specifically solve for Sequential Dependency among clients. Our FedWave framework is the first to introduce an MoE router as the core of federated aggregation, transforming it into an intelligent coordinator that can orchestrate expert knowledge for complex workflows.

3 METHODS

3.1 FEDERATED FINE-TUNING WITH THE VALUE CHAIN LAYER

The first phase of our framework fine-tunes a base LLM in a federated setting where each client, an expert agent (e.g., design, manufacturing), holds a private dataset for their role. To model the sequential relationships between these experts, we introduce the Value Chain Layer (VCLayer). As shown in Figure 1 (a), each expert agent i has a frozen base LLM, a trainable LoRA adapter (W_i), and our VCLayer. This lightweight, pluggable module processes the hidden states from the LoRA-adapted model, making them aware of the agent’s role and position in the workflow.

Given the final hidden states $H \in \mathbb{R}^{L \times d}$ from the base LLM (where L is sequence length, d is hidden dimension), the VCLayer applies a stage-specific transformation. For an agent at stage i , this is a specialized multi-head attention mechanism:

$$H'_i = \text{Attention}(HW_{q,i}, HW_{k,i}, HW_{v,i})W_{o,i} \quad (1)$$

where $W_{q,i}, W_{k,i}, W_{v,i} \in \mathbb{R}^{d \times d_{vc}}$ and $W_{o,i} \in \mathbb{R}^{d_{vc} \times d}$ are trainable, stage-specific projection matrices. This allows each expert to focus on task-relevant aspects of the input. The resulting state H'_i is then passed to the LM head to produce logits. A key innovation is the collaborative loss, \mathcal{L}_{VC} , which guides the VCLayer to learn the workflow structure. During local training for agent i , the total loss combines the standard Supervised Fine-Tuning (SFT) loss with our collaborative loss:

$$\mathcal{L}_{\text{total}}^{(i)} = \mathcal{L}_{\text{SFT}}^{(i)} + \gamma \mathcal{L}_{VC}^{(i)} \quad (2)$$

Here, γ is a balancing hyperparameter. \mathcal{L}_{SFT} is the conventional cross-entropy loss for next-token prediction. \mathcal{L}_{VC} comprises three terms that enforce the value chain’s relational structure:

$$\mathcal{L}_{VC} = \lambda_{\text{pos}} \mathcal{L}_{\text{pos}} + \lambda_{\text{cont}} \mathcal{L}_{\text{cont}} + \lambda_{\text{cons}} \mathcal{L}_{\text{cons}} \quad (3)$$

where λ_{pos} , λ_{cont} , and λ_{cons} are weighting coefficients.

Positional Loss (\mathcal{L}_{pos}): This loss enforces a geometric arrangement of experts in an embedding space, reflecting their sequential order. It is based on the similarity between learnable embeddings $\{e_0, \dots, e_{N-1}\}$ for each stage:

$$\mathcal{L}_{\text{pos}} = \sum_{i=0}^{N-2} \left(\cos(e_i, e_{i+1}) - T_{\text{pos}}^{(i)} \right)^2 \quad (4)$$

where $T_{\text{pos}}^{(i)}$ is a target cosine similarity, encouraging adjacent experts to be closer.

Continuity Loss ($\mathcal{L}_{\text{cont}}$): This loss promotes a smooth transition of knowledge by ensuring adjacent experts learn similar functions. It operates on the VCLayer’s projection matrices:

$$\mathcal{L}_{\text{cont}} = \sum_{i=0}^{N-2} \left(\cos(\text{vec}(W_{q,i}), \text{vec}(W_{q,i+1})) - T_{\text{cont}} \right)^2 \quad (5)$$

where $\text{vec}(\cdot)$ is the vectorization operator and T_{cont} is a high target similarity value, encouraging the attention mechanisms of consecutive stages to be functionally alike.

Consistency Loss ($\mathcal{L}_{\text{cons}}$): This loss ensures a coherent solution progression by aligning an expert’s output representation with its predecessor’s for the same input. For an agent at stage $i > 0$:

$$\mathcal{L}_{\text{cons}}^{(i)} = \left(\cos(f_i(\bar{H}), f_{i-1}(\bar{H})) - T_{\text{cons}} \right)^2 \quad (6)$$

where $f_i(\bar{H})$ is the mean-pooled output hidden state from the VCLayer of agent i for input H .

216 3.2 DYNAMIC AGGREGATION WITH MOE ROUTING
217

218 To overcome static aggregation limitations (e.g., FedAvg), we introduce a trainable Mixture of Ex-
219 perts (MoE) Router. This task-aware coordinator dynamically determines each expert’s contribution
220 based on the input prompt. As a shared global component, it is co-trained with the experts’ LoRA
221 adapters and VCLayers. The MoE Router, $G(\cdot)$, is a lightweight MLP. For any input, it computes a
222 representation by mean-pooling the base LLM’s hidden states, $\bar{H} = \frac{1}{L} \sum_{l=1}^L H_l$. It then processes
223 \bar{H} to produce logits $z \in \mathbb{R}^N$ over the N experts, which are converted to a probability distribution
224 via softmax:

$$225 \quad \alpha = \text{softmax}(G(\bar{H})) \quad (7)$$

226 where α_i is the routing weight for expert i . The MoE Router is trained implicitly via a loss weighting
227 mechanism during local SFT. We modulate each expert’s SFT loss for a sample by its assigned
228 routing weight. For expert i , the loss $\mathcal{L}_{\text{SFT}}^{(i)}$ is multiplied by α_i , strengthening the gradient signal
229 when the router correctly assigns a high weight to the appropriate expert.

230 We add two auxiliary losses to regularize the router. A **load balancing loss**, $\mathcal{L}_{\text{balance}}$, encourages even
231 expert utilization over a batch to prevent specialization collapse, and is formulated as the variance of
232 expert utilization. An **entropy-based confidence loss**, $\mathcal{L}_{\text{entropy}}$, penalizes uncertain routing decisions
233 to encourage sparse, confident weights. The total local loss for agent i is thus:

$$234 \quad \mathcal{L}_{\text{total}}^{(i)} = \alpha_i \mathcal{L}_{\text{SFT}}^{(i)} + \gamma \mathcal{L}_{\text{VC}}^{(i)} + \delta_1 \mathcal{L}_{\text{balance}} + \delta_2 \mathcal{L}_{\text{entropy}} \quad (8)$$

235 where δ_1 and δ_2 are hyperparameters. This composite loss enables end-to-end training of the LoRA
236 adapters, VCLayers, and the MoE Router. After each local training round, the router’s updated
237 weights are aggregated on the server, similar to FedAvg:

$$238 \quad W_{\text{router}}^{t+1} = \sum_{i \in S_t} \frac{n_i}{n} W_{\text{router},i}^{t+1} \quad (9)$$

242 3.3 PREFERENCE ALIGNMENT WITH DIRECT PREFERENCE OPTIMIZATION
243

244 After the federated SFT phase, the aggregated model has expert knowledge and a foundational
245 workflow understanding. The final phase (Figure 1 (b)) refines the model’s output by aligning it with
246 human preferences for quality and coherence using Direct Preference Optimization (DPO) (Rafailov
247 et al., 2023). This phase begins by using the final SFT model, π_{SFT} , to automatically generate a
248 preference dataset, $\mathcal{D}_{\text{pref}} = \{(x, y_w, y_l)\}$. The co-trained MoE Router’s intelligence is leveraged to
249 create these preference pairs. For each prompt x , the following steps are taken:

- 250 1. The MoE router within π_{SFT} computes routing weights α and ranks the N experts based on
251 their relevance to the prompt.
- 253 2. The top-ranked expert, $c_w = \arg \max_i \alpha_i$, is selected to generate the winning (chosen)
254 response, y_w .
- 255 3. A lower-ranked expert, c_l , is selected to generate the losing (rejected) response, y_l .
- 256 4. The tuple (x, y_w, y_l) is added to the preference dataset $\mathcal{D}_{\text{pref}}$.

257 This automated process uses the router’s expertise to create a large-scale dataset favoring outputs
258 from the most contextually appropriate expert. The aggregated model π_{SFT} is then fine-tuned on $\mathcal{D}_{\text{pref}}$
259 using DPO. DPO directly optimizes the model for preferences without a separate reward model. The
260 policy model, π_θ , is initialized from π_{SFT} , and a frozen copy of π_{SFT} serves as the reference model,
261 π_{ref} . DPO’s objective is to maximize the likelihood of preferred responses y_w and minimize that of
262 rejected responses y_l , constrained by a penalty term preventing large deviations from the reference
263 model. The loss function is:

$$264 \quad \mathcal{L}_{\text{DPO}}(\pi_\theta; \pi_{\text{ref}}) = -\log \sigma \left(\beta \log \frac{\pi_\theta(y_w|x)}{\pi_\theta(y_l|x)} - \beta \log \frac{\pi_{\text{ref}}(y_w|x)}{\pi_{\text{ref}}(y_l|x)} \right) \quad (10)$$

265 where σ is the logistic function and β controls the preference strength. Minimizing this loss over the
266 preference dataset aligns the policy model π_θ with the collaborative logic from the SFT phase. This
267 final step yields the fully optimized FedWave model, adept at specialized tasks and high-quality,
268 human-preferred collaboration.

270
271 Table 1: Performance comparison of FedWave and baselines across three BizWorkflow datasets and
272 three different backbone models. The best scores for each metric are highlighted in **bold**.
273
274

273 274 275 276 277 278 279 280 281 282 283 284 285 286 287 288 289 290 291 292 293 294 295 296 297 298 299 300 301 302 303 304 305 306 307 308 309 310 311 312 313 314 315 316 317 318 319 320 321 322 323	Automotive			E-commerce			Pharmaceutical		
	BS-F	Meteor	Rouge-L	BS-F	Meteor	Rouge-L	BS-F	Meteor	Rouge-L
<i>Qwen2-7B</i>									
FedAvg	71.11	23.35	22.18	79.17	43.29	38.68	77.12	26.76	28.54
FedAvgM	70.66	22.02	21.74	79.51	43.09	39.23	76.87	25.57	27.51
FedProx	71.36	23.58	22.53	79.23	43.53	38.84	77.18	27.08	28.65
FedAdam	71.39	23.74	21.74	78.78	41.01	37.07	76.41	26.07	27.29
FedYogi	71.49	24.18	22.48	78.54	40.31	36.81	76.00	24.59	26.11
Scaffold	71.26	23.53	22.49	79.04	42.32	38.50	77.06	27.09	28.68
FedWave	71.90	40.35	23.46	80.17	48.28	39.63	78.34	29.82	31.89
<i>Llama2-7B</i>									
FedAvg	69.43	17.41	20.98	72.96	16.45	22.66	70.61	10.10	14.98
FedAvgM	65.62	14.32	16.65	73.11	18.13	23.73	70.95	11.24	15.15
FedProx	69.26	17.08	20.74	72.82	16.30	22.52	70.92	10.27	15.06
FedAdam	68.90	17.03	19.55	70.83	13.49	19.88	68.79	8.75	13.34
FedYogi	69.12	17.20	19.99	70.42	13.15	19.40	68.36	8.31	12.94
Scaffold	69.25	16.78	20.64	72.95	16.62	22.72	70.97	10.39	15.25
FedWave	70.08	18.55	22.12	74.95	20.80	26.30	74.78	13.38	19.73
<i>Llama3-8B</i>									
FedAvg	58.62	10.73	9.55	68.81	26.25	25.29	61.34	14.17	14.49
FedAvgM	56.03	7.33	6.76	64.61	24.05	21.19	61.84	14.24	14.38
FedProx	58.44	10.50	9.27	68.52	27.16	25.66	61.79	14.50	15.06
FedAdam	51.47	2.67	2.25	57.22	18.28	14.65	58.08	12.16	11.49
FedYogi	52.78	4.17	3.66	57.37	18.16	14.77	58.20	11.68	11.05
Scaffold	59.15	11.31	10.21	68.12	25.65	24.60	61.62	14.85	14.97
FedWave	70.27	24.83	17.47	79.66	44.72	38.33	77.95	24.02	28.40

4 EXPERIMENTS

304 Our experiments evaluate FedWave to address three questions: (1) How does its collaborative per-
305 formance on sequential tasks compare to standard federated learning? (2) How does the privacy-
306 preserving FedWave perform against a centralized multi-agent system with full data access? (3)
307 Which design elements are most critical to its success?

4.1 SETTINGS

310 **Datasets and evaluation metrics.** We evaluate our framework on the MSCoRe benchmark (Lei
311 et al., 2025), which is specifically designed for multi-stage collaborative reasoning. It provides
312 three challenging datasets with complex, sequential tasks representing distinct business workflows:
313 **E-commerce**, **Pharmaceutical**, and **Automotive**. To comprehensively assess the quality of the
314 generated outputs, we employ a diverse suite of metrics beyond simple lexical overlap. This includes
315 ROUGE-1/2/L (Lin, 2004) for lexical content, BLEU-4 (Papineni et al., 2002) and GLEU (Wu et al.,
316 2016) for fluency, and crucially, METEOR (Banerjee & Lavie, 2005) and BERTScore (Zhang et al.,
317 2019) to evaluate deeper semantic fidelity and contextual relevance.

318 **Baselines.** To validate our framework’s effectiveness, we compare it against two baseline cat-
319 egories. For federated learning, we adapt widely-recognized algorithms: the foundational Fed-
320 Avg (McMahan et al., 2017); FedProx (Li et al., 2020) with its proximal term to miti-
321 gate heterogeneity; FedAvgM (Hsu et al., 2019), which adds server-side momentum; SCAF-
322 FOLD (Karimireddy et al., 2021) for client-drift correction; and the adaptive optimizers FedAdam
323 and FedYogi (Reddi et al., 2020). For multi-agent systems, we benchmark against centralized meth-

324
 325 Table 2: Performance comparison of FedWave against centralized multi-agent baselines on the
 326 Qwen2-7B backbone. FedWave operates in a decentralized, privacy-preserving setting, while the
 327 baselines have access to the full, centralized dataset. The best scores for each metric are highlighted
 328 in **bold**.

329 Dataset	330 Baselines	331 BS-F	332 GLEU	333 BLEU-4	334 ROUGE-1	335 ROUGE-2	336 ROUGE-L
337 Automotive	PMC (Zhang et al., 2025)	65.81	13.25	5.24	26.44	4.51	18.37
	MedAgents (Tang et al., 2024b)	64.94	12.47	4.81	23.68	4.07	18.34
	Debate(long) (Du et al., 2024)	65.56	12.92	6.49	25.49	5.91	18.03
	Debate(short) (Du et al., 2024)	65.32	12.66	6.30	25.39	6.02	17.81
	CoA (Zhang et al., 2024)	70.77	22.30	<u>14.41</u>	<u>33.99</u>	<u>11.15</u>	24.01
	FedWave (Ours)	71.47	<u>20.42</u>	15.44	35.11	12.19	<u>22.47</u>
338 E-commerce	PMC (Zhang et al., 2025)	70.88	24.01	19.73	33.47	11.45	29.91
	MedAgents (Tang et al., 2024b)	69.57	17.48	12.60	33.62	10.77	28.25
	Debate(long) (Du et al., 2024)	72.89	20.81	15.46	36.43	13.12	28.04
	Debate(short) (Du et al., 2024)	72.93	20.85	15.56	37.11	13.51	28.23
	CoA (Zhang et al., 2024)	79.20	<u>38.08</u>	<u>34.23</u>	<u>50.02</u>	<u>25.35</u>	<u>39.02</u>
	FedWave (Ours)	80.17	42.60	39.56	51.04	26.22	39.63
339 Pharmaceutical	PMC (Zhang et al., 2025)	68.55	17.34	11.67	28.16	7.52	25.93
	MedAgents (Tang et al., 2024b)	65.96	10.60	4.46	25.91	6.17	21.13
	Debate(long) (Du et al., 2024)	61.14	9.58	5.33	16.18	4.69	12.23
	Debate(short) (Du et al., 2024)	70.57	15.78	9.11	31.76	9.60	23.03
	CoA (Zhang et al., 2024)	77.26	29.92	<u>18.18</u>	<u>42.38</u>	<u>19.93</u>	33.70
	FedWave (Ours)	78.34	<u>26.27</u>	19.22	43.24	20.00	<u>31.89</u>

344
 345
 346
 347
 348
 349
 350
 351
 352
 353
 354
 355
 356
 357
 358
 359
 360
 361
 362
 363
 364
 365
 366
 367
 368
 369
 370
 371
 372
 373
 374
 375
 376
 377
 378
 379
 380
 381
 382
 383
 384
 385
 386
 387
 388
 389
 390
 391
 392
 393
 394
 395
 396
 397
 398
 399
 400
 401
 402
 403
 404
 405
 406
 407
 408
 409
 410
 411
 412
 413
 414
 415
 416
 417
 418
 419
 420
 421
 422
 423
 424
 425
 426
 427
 428
 429
 430
 431
 432
 433
 434
 435
 436
 437
 438
 439
 440
 441
 442
 443
 444
 445
 446
 447
 448
 449
 450
 451
 452
 453
 454
 455
 456
 457
 458
 459
 460
 461
 462
 463
 464
 465
 466
 467
 468
 469
 470
 471
 472
 473
 474
 475
 476
 477
 478
 479
 480
 481
 482
 483
 484
 485
 486
 487
 488
 489
 490
 491
 492
 493
 494
 495
 496
 497
 498
 499
 500
 501
 502
 503
 504
 505
 506
 507
 508
 509
 510
 511
 512
 513
 514
 515
 516
 517
 518
 519
 520
 521
 522
 523
 524
 525
 526
 527
 528
 529
 530
 531
 532
 533
 534
 535
 536
 537
 538
 539
 540
 541
 542
 543
 544
 545
 546
 547
 548
 549
 550
 551
 552
 553
 554
 555
 556
 557
 558
 559
 560
 561
 562
 563
 564
 565
 566
 567
 568
 569
 570
 571
 572
 573
 574
 575
 576
 577
 578
 579
 580
 581
 582
 583
 584
 585
 586
 587
 588
 589
 590
 591
 592
 593
 594
 595
 596
 597
 598
 599
 600
 601
 602
 603
 604
 605
 606
 607
 608
 609
 610
 611
 612
 613
 614
 615
 616
 617
 618
 619
 620
 621
 622
 623
 624
 625
 626
 627
 628
 629
 630
 631
 632
 633
 634
 635
 636
 637
 638
 639
 640
 641
 642
 643
 644
 645
 646
 647
 648
 649
 650
 651
 652
 653
 654
 655
 656
 657
 658
 659
 660
 661
 662
 663
 664
 665
 666
 667
 668
 669
 670
 671
 672
 673
 674
 675
 676
 677
 678
 679
 680
 681
 682
 683
 684
 685
 686
 687
 688
 689
 690
 691
 692
 693
 694
 695
 696
 697
 698
 699
 700
 701
 702
 703
 704
 705
 706
 707
 708
 709
 710
 711
 712
 713
 714
 715
 716
 717
 718
 719
 720
 721
 722
 723
 724
 725
 726
 727
 728
 729
 730
 731
 732
 733
 734
 735
 736
 737
 738
 739
 740
 741
 742
 743
 744
 745
 746
 747
 748
 749
 750
 751
 752
 753
 754
 755
 756
 757
 758
 759
 760
 761
 762
 763
 764
 765
 766
 767
 768
 769
 770
 771
 772
 773
 774
 775
 776
 777
 778
 779
 780
 781
 782
 783
 784
 785
 786
 787
 788
 789
 790
 791
 792
 793
 794
 795
 796
 797
 798
 799
 800
 801
 802
 803
 804
 805
 806
 807
 808
 809
 810
 811
 812
 813
 814
 815
 816
 817
 818
 819
 820
 821
 822
 823
 824
 825
 826
 827
 828
 829
 830
 831
 832
 833
 834
 835
 836
 837
 838
 839
 840
 841
 842
 843
 844
 845
 846
 847
 848
 849
 850
 851
 852
 853
 854
 855
 856
 857
 858
 859
 860
 861
 862
 863
 864
 865
 866
 867
 868
 869
 870
 871
 872
 873
 874
 875
 876
 877
 878
 879
 880
 881
 882
 883
 884
 885
 886
 887
 888
 889
 890
 891
 892
 893
 894
 895
 896
 897
 898
 899
 900
 901
 902
 903
 904
 905
 906
 907
 908
 909
 910
 911
 912
 913
 914
 915
 916
 917
 918
 919
 920
 921
 922
 923
 924
 925
 926
 927
 928
 929
 930
 931
 932
 933
 934
 935
 936
 937
 938
 939
 940
 941
 942
 943
 944
 945
 946
 947
 948
 949
 950
 951
 952
 953
 954
 955
 956
 957
 958
 959
 960
 961
 962
 963
 964
 965
 966
 967
 968
 969
 970
 971
 972
 973
 974
 975
 976
 977
 978
 979
 980
 981
 982
 983
 984
 985
 986
 987
 988
 989
 990
 991
 992
 993
 994
 995
 996
 997
 998
 999
 1000
 1001
 1002
 1003
 1004
 1005
 1006
 1007
 1008
 1009
 1010
 1011
 1012
 1013
 1014
 1015
 1016
 1017
 1018
 1019
 1020
 1021
 1022
 1023
 1024
 1025
 1026
 1027
 1028
 1029
 1030
 1031
 1032
 1033
 1034
 1035
 1036
 1037
 1038
 1039
 1040
 1041
 1042
 1043
 1044
 1045
 1046
 1047
 1048
 1049
 1050
 1051
 1052
 1053
 1054
 1055
 1056
 1057
 1058
 1059
 1060
 1061
 1062
 1063
 1064
 1065
 1066
 1067
 1068
 1069
 1070
 1071
 1072
 1073
 1074
 1075
 1076
 1077
 1078
 1079
 1080
 1081
 1082
 1083
 1084
 1085
 1086
 1087
 1088
 1089
 1090
 1091
 1092
 1093
 1094
 1095
 1096
 1097
 1098
 1099
 1100
 1101
 1102
 1103
 1104
 1105
 1106
 1107
 1108
 1109
 1110
 1111
 1112
 1113
 1114
 1115
 1116
 1117
 1118
 1119
 1120
 1121
 1122
 1123
 1124
 1125
 1126
 1127
 1128
 1129
 1130
 1131
 1132
 1133
 1134
 1135
 1136
 1137
 1138
 1139
 1140
 1141
 1142
 1143
 1144
 1145
 1146
 1147
 1148
 1149
 1150
 1151
 1152
 1153
 1154
 1155
 1156
 1157
 1158
 1159
 1160
 1161
 1162
 1163
 1164
 1165
 1166
 1167
 1168
 1169
 1170
 1171
 1172
 1173
 1174
 1175
 1176
 1177
 1178
 1179
 1180
 1181
 1182
 1183
 1184
 1185
 1186
 1187
 1188
 1189
 1190
 1191
 1192
 1193
 1194
 1195
 1196
 1197
 1198
 1199
 1200
 1201
 1202
 1203
 1204
 1205
 1206
 1207
 1208
 1209
 1210
 1211
 1212
 1213
 1214
 1215
 1216
 1217
 1218
 1219
 1220
 1221
 1222
 1223
 1224
 1225
 1226
 1227
 1228
 1229
 1230
 1231
 1232
 1233
 1234
 1235
 1236
 1237
 1238
 1239
 1240
 1241
 1242
 1243
 1244
 1245
 1246
 1247
 1248
 1249
 1250
 1251
 1252
 1253
 1254
 1255
 1256
 1257
 1258
 1259
 1260
 1261
 1262
 1263
 1264
 1265
 1266
 1267
 1268
 1269
 1270
 1271
 1272
 1273
 1274
 1275
 1276
 1277
 1278
 1279
 1280
 1281
 1282
 1283
 1284
 1285
 1286
 1287
 1288
 1289
 1290
 1291
 1292
 1293
 1294
 1295
 1296
 1297
 1298
 1299
 1300
 1301
 1302
 1303
 1304
 1305
 1306
 1307
 1308
 1309
 1310
 1311
 1312
 1313
 1314
 1315
 1316
 1317
 1318
 1319
 1320
 1321
 1322
 1323
 1324
 1325
 1326
 1327
 1328
 1329
 1330
 1331
 1332
 1333
 1334
 1335
 1336
 1337
 1338
 1339
 1340
 1341
 1342
 1343
 1344
 1345
 1346
 1347
 1348
 1349
 1350
 1351
 1352
 1353
 1354
 1355
 1356
 1357
 1358
 1359
 1360
 1361
 1362
 1363
 1364
 1365
 1366
 1367
 1368
 1369
 1370
 1371
 1372
 1373
 1374
 1375
 1376
 1377
 1378
 1379
 1380
 1381
 1382
 1383
 1384
 1385
 1386
 1387
 1388
 1389
 1390
 1391
 1392
 1393
 1394
 1395
 1396
 1397
 1398
 1399
 1400
 1401
 1402
 1403
 1404
 1405
 1406
 1407
 1408
 1409
 1410
 1411
 1412
 1413
 1414
 1415
 1416
 1417
 1418
 1419
 1420
 1421
 1422
 1423
 1424
 1425
 1426
 1427
 1428
 1429
 1430
 1431
 1432
 1433
 1434
 1435
 1436
 1437
 1438
 1439
 1440
 1441
 1442
 1443
 1444
 1445
 1446
 1447
 1448
 1449
 1450
 1451
 1452
 1453
 1454
 1455
 1456
 1457
 1458
 1459
 1460
 1461
 1462
 1463
 1464
 1465
 1466
 1467
 1468
 1469
 1470
 1471
 1472
 1473
 1474
 1475
 1476
 1477
 1478
 1479
 1480
 1481
 1482
 1483
 1484
 1485
 1486
 1487
 1488
 1489
 1490
 1491
 1492
 1493
 1494
 1495
 1496
 1497
 1498
 1499
 1500
 1501
 1502
 1503
 1504
 1505
 1506
 1507
 1508
 1509
 1510
 1511
 1512
 1513
 1514
 1515
 1516
 1517
 1518
 1519
 1520
 1521
 1522
 1523
 1524
 1525
 1526
 1527
 1528
 1529
 1530
 1531
 1532
 1533
 1534
 1535
 1536
 1537
 1538
 1539
 1540
 1

378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431

Table 3: Ablation study of the key components of FedWave on the Automotive dataset using the Llama3-8B backbone.

Method	BS-F	Meteor	Rouge-L
FedWave (Full Model)	70.27	24.83	17.47
<i>- Ablation Variants -</i>			
w/o DPO	66.55	18.40	17.14
w/o MoE Router	69.00	23.65	16.99
w/o VCLayer	69.07	23.96	16.21
<i>- Baseline -</i>			
SFT Only (FedAvg)	58.62	10.73	9.55

Table 4: Performance of the FedWave framework when integrated with different federated optimization algorithms. Experiments are conducted with the Llama-8B backbone.

Aggregation Algorithm	BS-F	Meteor	Rouge-L
FedWave (Default)	66.55	18.40	17.14
+ FedAvgM	70.59	23.43	21.73
+ FedProx	70.86	<u>23.71</u>	22.20
+ FedAdam	70.68	23.68	21.89
+ FedAdagrad	70.71	23.79	<u>21.97</u>
+ FedYogi	70.35	23.42	21.46

4.3 COLLABORATIVE PERFORMANCE OF FEDWAVE COMPARED TO MULTI-AGENT

We further benchmark FedWave against several centralized Multi-Agent baselines, as shown in Table 2. Crucially, while baselines operate with full data access, FedWave performs in a decentralized, privacy-preserving setting. Despite this challenging condition, our framework demonstrates highly competitive or even superior performance. On the E-commerce dataset, FedWave surpasses all centralized methods across every reported metric. On the Automotive and Pharmaceutical datasets, it achieves state-of-the-art results on key metrics such as BERTScore-F and ROUGE, proving its ability to generate high-quality, semantically rich outputs.

4.4 KEY DESIGN FACTORS AND HYPERPARAMETER INFLUENCE IN FEDWAVE

Ablation Study. The ablation study results in Table 3 reveal the powerful synergistic effect of our framework’s components. While removing any single component—the DPO stage, MoE Router, or VCLayer—causes a noticeable performance degradation, the decline is not catastrophic as the remaining parts partially compensate to maintain a degree of collaborative intelligence. For instance, without the MoE Router, the VCLayer still ensures the agents learn their sequential roles. However, when all three components are removed (reverting to the ‘SFT Only (FedAvg)’ baseline), the collaborative intelligence system collapses, causing a drastic drop across all metrics (e.g., Meteor falls from 24.83 to 10.73). This substantial gap demonstrates that our components have a multiplicative, not merely additive, effect. It is the integrated combination of explicit workflow modeling (VCLayer), dynamic knowledge aggregation (MoE Router), and preference alignment (DPO) that collectively enables FedWave’s superior performance.

Hyper-parameter Sensitivity Analysis. We conduct a comprehensive analysis to evaluate the sensitivity of FedWave to its key hyperparameters, demonstrating the framework’s robustness.

VCLayer Loss Components. As illustrated in Figure 2, we first vary the individual weights for the Positional Loss (λ_{pos}), Continuity Loss (λ_{cont}), and Consistency Loss (λ_{cons}). For both λ_{pos} and

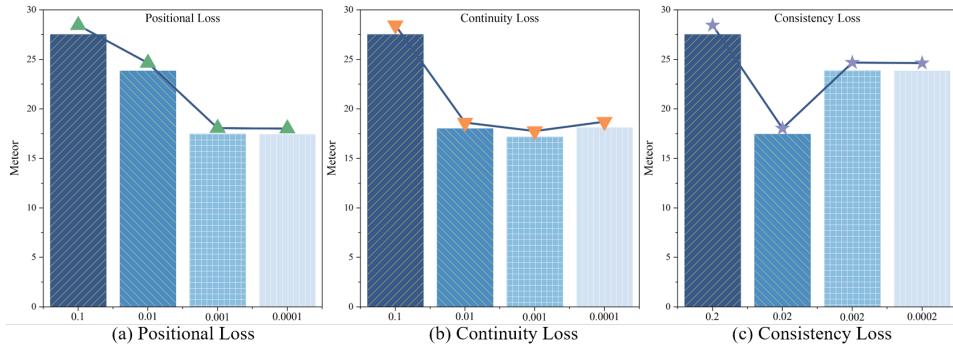


Figure 2: Sensitivity analysis for the weights of the collaborative loss components in the VCLayer.

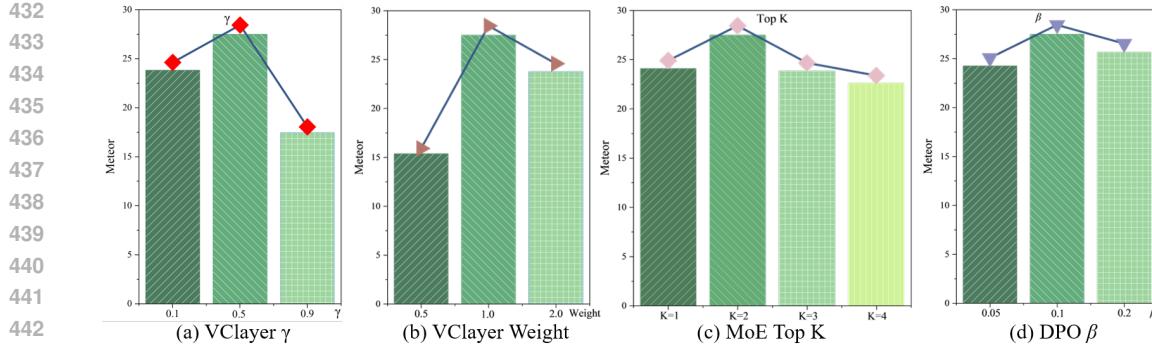


Figure 3: Sensitivity analysis for core hyperparameters of FedWave: (a) Balancing coefficient γ for SFT and VC losses; (b) Overall scaling factor for VCLayer weights; (c) MoE router’s top_k value; and (d) DPO’s beta value.

λ_{cont} , the performance, measured by Meteor, peaks at a weight of 0.1 and gracefully degrades as the weight decreases, confirming their significant contribution to modeling the workflow structure. The framework also shows robustness to changes in λ_{cons} , maintaining stable performance across a range of values.

Other Core Hyperparameters. We further analyze several critical hyperparameters in Figure 3. We study the balancing coefficient γ (Figure 3 (a)), which trades off between the SFT and VCLayer losses, finding that $\gamma = 0.5$ yields an optimal balance. Our analysis of an overall scaling factor for the VCLayer weights (Figure 3 (b)) shows performance peaks at the default scaling of 1.0, confirming our weights are well-calibrated. The study of MoE top_k (Figure 3 (c)) reveals that performance peaks at k=2, suggesting that activating a small, focused group of experts is most effective. Finally, the analysis of DPO β (Figure 3 (d)) confirms that a value of 0.1 provides the best alignment without significant deviation from the base model. Overall, these results demonstrate that while our proposed components are crucial, FedWave is not overly sensitive to their precise hyperparameter values, highlighting its stability and reliability.

4.5 ANALYSIS ON FEDERATED AGGREGATION ALGORITHMS

To assess the compatibility and modularity of our framework, we integrated FedWave with several advanced federated optimization algorithms, replacing the default FedAvg-based aggregation. The results, presented in Table 4, show that all tested optimizers yield a substantial performance improvement over the default configuration. Notably, algorithms like FedProx and those with adaptive optimization (FedAdagrad, FedAdam) achieve the higher scores. This suggests that the structured, non-IID environment created by our VCLayer and MoE router benefits significantly from optimizers designed to handle client drift and heterogeneity. FedProx, with its regularization term, likely prevents the specialized expert models from diverging too far from the global consensus, while adaptive methods better navigate the complex loss landscape.

5 CONCLUSION

In this paper, we introduced FedWave, a novel federated multi-agent collaboration framework designed to address the critical challenge of solving complex sequential tasks across decentralized, privacy-sensitive data silos. By integrating a collaborative VCLayer, a dynamic MoE router for intelligent aggregation, and a final DPO stage for preference alignment, our framework successfully bridges the gap between the collaborative capabilities of multi-agent systems and the privacy guarantees of federated learning. Our extensive experiments demonstrate that FedWave not only significantly outperforms standard federated learning baselines but also achieves performance competitive with, and often superior to, centralized multi-agent systems that have unrestricted data access.

486 REFERENCES
487

488 Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Ale-
489 man, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical
490 report. *arXiv preprint arXiv:2303.08774*, 2023.

491 Elif Akata, Lion Schulz, Julian Coda-Forno, Seong Joon Oh, Matthias Bethge, and Eric Schulz.
492 Playing repeated games with large language models. *arXiv preprint arXiv:2305.16867*, 2023.

493

494 Manoj Ghuhan Arivazhagan, Vinay Aggarwal, Aaditya Kumar Singh, and Sunav Choudhary. Fed-
495 erated learning with personalization layers, 2019. URL <https://arxiv.org/abs/1912.00818>.

496

497 Satanjeev Banerjee and Alon Lavie. Meteor: An automatic metric for mt evaluation with improved
498 correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic*
499 *evaluation measures for machine translation and/or summarization*, pp. 65–72, 2005.

500

501 Justin Chih-Yao Chen, Swarnadeep Saha, and Mohit Bansal. Reconcile: Round-table conference
502 improves reasoning via consensus among diverse llms, 2024. URL <https://arxiv.org/abs/2309.13007>.

503

504 Yilun Du, Shuang Li, Antonio Torralba, Joshua B. Tenenbaum, and Igor Mordatch. Improving
505 factuality and reasoning in language models through multiagent debate, 2024. URL <https://openreview.net/forum?id=QAwaalJNCK>.

506

507 Alireza Fallah, Aryan Mokhtari, and Asuman Ozdaglar. Personalized federated learning with the-
508 retorical guarantees: A model-agnostic meta-learning approach. *Advances in neural information*
509 *processing systems*, 33:3557–3568, 2020.

510

511 Tao Fan, Yan Kang, Guoqiang Ma, Weijing Chen, Wenbin Wei, Lixin Fan, and Qiang Yang. Fate-
512 llm: A industrial grade federated learning framework for large language models. *Symposium on*
513 *Advances and Open Problems in Large Language Models (LLM@IJCAI'23)*, 2023.

514

515 William Fedus, Barret Zoph, and Noam Shazeer. Switch transformers: Scaling to trillion parameter
516 models with simple and efficient sparsity. *Journal of Machine Learning Research*, 23(120):1–39,
517 2022.

518

519 Taicheng Guo, Xiuying Chen, Yaqi Wang, Ruidi Chang, Shichao Pei, Nitesh V Chawla, Olaf Wiest,
520 and Xiangliang Zhang. Large language model based multi-agents: A survey of progress and
521 challenges. *arXiv preprint arXiv:2402.01680*, 2024.

522

523 Rui Hao, Linmei Hu, Weijian Qi, Qingliu Wu, Yirui Zhang, and Liqiang Nie. Chatllm network:
524 More brains, more intelligence. *arXiv preprint arXiv:2304.12998*, 2023.

525

526 Tzu-Ming Harry Hsu, Hang Qi, and Matthew Brown. Measuring the effects of non-identical data dis-
527 tribution for federated visual classification, 2019. URL <https://arxiv.org/abs/1909.06335>.

528

529 Xinmeng Huang, Ping Li, and Xiaoyun Li. Stochastic controlled averaging for federated learning
530 with communication compression. In *The Twelfth International Conference on Learning Repre-
sentations*, 2024. URL <https://openreview.net/forum?id=jj5ZjZsWJe>.

531

532 Robert A Jacobs, Michael I Jordan, Steven J Nowlan, and Geoffrey E Hinton. Adaptive mixtures of
533 local experts. *Neural computation*, 3(1):79–87, 1991.

534

535 Albert Q Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bam-
536 ford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, et al.
Mixtral of experts. *arXiv preprint arXiv:2401.04088*, 2024.

537

538 Peter Kairouz, H. Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin
539 Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, Rafael
G. L. D’Oliveira, Hubert Eichner, Salim El Rouayheb, David Evans, Josh Gardner, Zachary
Garrett, Adrià Gascón, Badih Ghazi, Phillip B. Gibbons, Marco Gruteser, Zaid Harchaoui,

540 Chaoyang He, Lie He, Zhouyuan Huo, Ben Hutchinson, Justin Hsu, Martin Jaggi, Tara Ja-
 541 vidi, Gauri Joshi, Mikhail Khodak, Jakub Konečný, Aleksandra Korolova, Farinaz Koushanfar,
 542 Sanmi Koyejo, Tancrède Lepoint, Yang Liu, Prateek Mittal, Mehryar Mohri, Richard Nock,
 543 Ayfer Özgür, Rasmus Pagh, Mariana Raykova, Hang Qi, Daniel Ramage, Ramesh Raskar,
 544 Dawn Song, Weikang Song, Sebastian U. Stich, Ziteng Sun, Ananda Theertha Suresh, Flo-
 545 rian Tramèr, Praneeth Vepakomma, Jianyu Wang, Li Xiong, Zheng Xu, Qiang Yang, Felix X.
 546 Yu, Han Yu, and Sen Zhao. Advances and open problems in federated learning, 2021. URL
 547 <https://arxiv.org/abs/1912.04977>.

548 Sai Praneeth Karimireddy, Satyen Kale, Mehryar Mohri, Sashank J. Reddi, Sebastian U. Stich, and
 549 Ananda Theertha Suresh. Scaffold: Stochastic controlled averaging for federated learning, 2021.
 550 URL <https://arxiv.org/abs/1910.06378>.

551 Weirui Kuang, Bingchen Qian, Zitao Li, Daoyuan Chen, Dawei Gao, Xuchen Pan, Yuexiang Xie,
 552 Yaliang Li, Bolin Ding, and Jingren Zhou. Federatedscope-llm: A comprehensive package for
 553 fine-tuning large language models in federated learning. In *Proceedings of the 30th ACM SIGKDD*
 554 *Conference on Knowledge Discovery and Data Mining*, pp. 5260–5271, 2024.

555 Yuzhen Lei, Hongbin Xie, Jiaxing Zhao, Shuangxue Liu, and Xuan Song. Mscore: A benchmark for
 556 multi-stage collaborative reasoning in llm agents, 2025. URL <https://arxiv.org/abs/2509.17628>.

557 Dmitry Lepikhin, HyoukJoong Lee, Yuanzhong Xu, Dehao Chen, Orhan Firat, Yanping Huang,
 558 Maxim Krikun, Noam Shazeer, and Zhifeng Chen. Gshard: Scaling giant models with conditional
 559 computation and automatic sharding. *arXiv preprint arXiv:2006.16668*, 2020.

560 Huao Li, Hossein Nourkhiz Mahjoub, Behdad Chalaki, Vaishnav Tadiparthi, Kwonjoon Lee,
 561 Ehsan Moradi-Pari, Charles Michael Lewis, and Katia P Sycara. Language grounded multi-
 562 agent reinforcement learning with human-interpretable communication, 2024. URL <https://arxiv.org/abs/2409.17348>.

563 Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith.
 564 Federated optimization in heterogeneous networks, 2020. URL <https://arxiv.org/abs/1812.06127>.

565 Tian Li, Shengyuan Hu, Ahmad Beirami, and Virginia Smith. Ditto: Fair and robust federated
 566 learning through personalization. In *International conference on machine learning*, pp. 6357–
 567 6368. PMLR, 2021.

568 Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization
 569 branches out*, pp. 74–81, 2004.

570 Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas.
 571 Communication-efficient learning of deep networks from decentralized data. In *Artificial intelli-
 572 gence and statistics*, pp. 1273–1282. PMLR, 2017.

573 H. Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas.
 574 Communication-efficient learning of deep networks from decentralized data, 2023. URL <https://arxiv.org/abs/1602.05629>.

575 Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong
 576 Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to fol-
 577 low instructions with human feedback. *Advances in neural information processing systems*, 35:
 578 27730–27744, 2022.

579 Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic
 580 evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association
 581 for Computational Linguistics*, pp. 311–318, 2002.

582 Yiyue Qian, Shinan Zhang, Yun Zhou, Haibo Ding, Diego Socolin-
 583 sky, and Yi Zhang. Enhancing llm-as-a-judge via multi-agent collabora-
 584 tion. 2025. URL <https://www.amazon.science/publications/enhancing-llm-as-a-judge-via-multi-agent-collaboration>.

594 Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea
 595 Finn. Direct preference optimization: Your language model is secretly a reward model. In
 596 *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=HPuSIXJaa9>.
 597

598 Sashank Reddi, Zachary Charles, Manzil Zaheer, Zachary Garrett, Keith Rush, Jakub Konečný,
 599 Sanjiv Kumar, and H Brendan McMahan. Adaptive federated optimization. *arXiv preprint*
 600 *arXiv:2003.00295*, 2020.
 601

602 Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc Le, Geoffrey Hinton,
 603 and Jeff Dean. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer.
 604 *arXiv preprint arXiv:1701.06538*, 2017.
 605

606 Simeng Sun, Yang Liu, Shuohang Wang, Chenguang Zhu, and Mohit Iyyer. Pearl: Prompt-
 607 ing large language models to plan and execute actions over long documents. *arXiv preprint*
 608 *arXiv:2305.14564*, 2023.
 609

610 Canh T Dinh, Nguyen Tran, and Josh Nguyen. Personalized federated learning with moreau
 611 envelopes. *Advances in neural information processing systems*, 33:21394–21405, 2020.
 612

613 Xiangru Tang, Anni Zou, Zhuosheng Zhang, Ziming Li, Yilun Zhao, Xingyao Zhang, Arman Cohan,
 614 and Mark Gerstein. MedAgents: Large language models as collaborators for zero-shot medical
 615 reasoning. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Findings of the Association
 616 for Computational Linguistics: ACL 2024*, pp. 599–621, Bangkok, Thailand, August
 617 2024a. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-acl.33. URL
 618 <https://aclanthology.org/2024.findings-acl.33/>.
 619

620 Xiangru Tang, Anni Zou, Zhuosheng Zhang, Ziming Li, Yilun Zhao, Xingyao Zhang, Arman Cohan,
 621 and Mark Gerstein. MedAgents: Large language models as collaborators for zero-shot medical
 622 reasoning. In *Findings of the Association for Computational Linguistics: ACL 2024*, pp. 599–621,
 623 Bangkok, Thailand, August 2024b. Association for Computational Linguistics. doi: 10.18653/
 624 v1/2024.findings-acl.33. URL [https://aclanthology.org/2024.findings-acl.
 33/](https://aclanthology.org/2024.findings-acl.33/).
 625

626 Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin,
 627 Percy Liang, and Tatsunori B Hashimoto. Alpaca: A strong, replicable instruction-
 628 following model. *Stanford Center for Research on Foundation Models*. <https://crfm.stanford.edu/2023/03/13/alpaca.html>, 3(6):7, 2023.
 629

630 Qwen Team. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*, 2024.
 631

632 Pablo Villalobos, Jaime Sevilla, Lennart Heim, Tamay Besiroglu, Marius Hobbahn, and Anson Ho.
 633 Will we run out of data? an analysis of the limits of scaling datasets in machine learning. *arXiv
 634 preprint arXiv:2211.04325*, 1:1, 2022.
 635

636 Jianyu Wang, Qinghua Liu, Hao Liang, Gauri Joshi, and H. Vincent Poor. Tackling the objective
 637 inconsistency problem in heterogeneous federated optimization, 2020. URL <https://arxiv.org/abs/2007.07481>.
 638

639 Taylor Webb, Keith J Holyoak, and Hongjing Lu. Emergent analogical reasoning in large language
 640 models. *Nature Human Behaviour*, 7(9):1526–1541, 2023.
 641

642 Feijie Wu, Zitao Li, Yaliang Li, Bolin Ding, and Jing Gao. Fedbiot: Llm local fine-tuning in
 643 federated learning without full model. In *Proceedings of the 30th ACM SIGKDD Conference on
 644 Knowledge Discovery and Data Mining*, pp. 3345–3355, 2024.
 645

646 Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey,
 647 Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. Google’s neural machine trans-
 648 lation system: Bridging the gap between human and machine translation. *arXiv preprint*
 649 *arXiv:1609.08144*, 2016.

648 Kai Xiong, Xiao Ding, Yixin Cao, Ting Liu, and Bing Qin. Examining inter-consistency of large
 649 language models collaboration: An in-depth analysis via debate. In *The 2023 Conference on*
 650 *Empirical Methods in Natural Language Processing*, 2023. URL <https://openreview.net/forum?id=XEwQ1fDbDN>.

652 Hongyang Yang, Xiao-Yang Liu, and Christina Dan Wang. Fingpt: Open-source financial large
 653 language models, 2023. URL <https://arxiv.org/abs/2306.06031>.

655 Yuhang Yao, Jianyi Zhang, Junda Wu, Chengkai Huang, Yu Xia, Tong Yu, Ruiyi Zhang, Sungchul
 656 Kim, Ryan Rossi, Ang Li, Lina Yao, Julian McAuley, Yiran Chen, and Carlee Joe-Wong. Federated
 657 large language models: Current progress and future directions, 2024. URL <https://arxiv.org/abs/2409.15723>.

659 Rui Ye, Wenhao Wang, Jingyi Chai, Dihan Li, Zexi Li, Yinda Xu, Yaxin Du, Yanfeng Wang, and
 660 Siheng Chen. Openfedllm: Training large language models on decentralized private data via fed-
 661 erated learning. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery*
 662 *and Data Mining*, pp. 6137–6147, 2024.

664 Cong Zhang, Xin Deik Goh, Dexun Li, Hao Zhang, and Yong Liu. Planning with multi-constraints
 665 via collaborative language agents. In *Proceedings of the 31st International Conference on Com-
 666 putational Linguistics*, pp. 10054–10082, 2025.

667 Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. Bertscore: Evaluat-
 668 ing text generation with bert. *arXiv preprint arXiv:1904.09675*, 2019.

669 Yusen Zhang, Ruoxi Sun, Yanfei Chen, Tomas Pfister, Rui Zhang, and Sercan Arik. Chain of agents:
 670 Large language models collaborating on long-context tasks. *Advances in Neural Information*
 671 *Processing Systems*, 37:132208–132237, 2024.

673 Jiaxing Zhao, Hongbin Xie, Yuzhen Lei, Xuan Song, Zhuoran Shi, Lianxin Li, Shuangxue Liu,
 674 and Haoran Zhang. Connecting the dots: A chain-of-collaboration prompting framework for llm
 675 agents. *arXiv preprint arXiv:2505.10936*, 2025.

676 Jun Zhao, Can Zu, Hao Xu, Yi Lu, Wei He, Yiwen Ding, Tao Gui, Qi Zhang, and Xuanjing Huang.
 677 Longagent: Scaling language models to 128k context through multi-agent collaboration, 2024.
 678 URL <https://arxiv.org/abs/2402.11550>.

680
 681
 682
 683
 684
 685
 686
 687
 688
 689
 690
 691
 692
 693
 694
 695
 696
 697
 698
 699
 700
 701

702 A QUALITATIVE ANALYSIS: A CASE STUDY

704 To provide a more intuitive understanding of the performance differences, we present a qualitative
 705 case study on a representative task from the Automotive workflow dataset. We prompted the models
 706 with a complex, two-part request that requires both marketing creativity (Expert 1’s domain) and
 707 supply chain knowledge (Expert 2’s domain). The results, shown in Figure 4, highlight the distinct
 708 collaborative capabilities of each approach.

709 The analysis of the outputs in Figure 4 reveals the clear superiority of our proposed framework. The
 710 **FedAvg**, lacking any mechanism for structured collaboration, produces a convoluted and generic
 711 response. It struggles to differentiate between the distinct tasks of marketing and supply chain
 712 planning, mixing concepts and failing to provide the specific, actionable details required by either
 713 role. This output exemplifies the shortcomings of simple model averaging, which dilutes expert
 714 knowledge into a generalized but ultimately ineffective model.

715 In contrast, the **Centralized MAS** baseline, with full data access, performs well. It correctly sepa-
 716 rates the two sub-tasks and provides relevant, technically sound details for each, such as focusing on
 717 the core chassis components. This demonstrates the effectiveness of the multi-agent paradigm when
 718 privacy is not a concern.

719 Most importantly, **FedWave** achieves a remarkable level of performance that is highly competi-
 720 tive with the centralized approach, despite operating under strict privacy constraints. The output
 721 demonstrates a clear and seamless transition between the two roles. The marketing announcement is
 722 creative, while the supply chain requirements are precise and technical. Crucially, the generated de-
 723 tails are not only plausible but also qualitatively different from the centralized model’s output—for
 724 instance, specifying innovative components like ‘BladeCell’ batteries and ‘WaveDrive’ motors, and
 725 focusing on sustainable interior textiles. This suggests that FedWave successfully synthesizes spe-
 726 cialized knowledge learned from distributed data sources, a feat not possible with simple averaging.
 727 This successful role-switching is a direct result of our **MoE router** dynamically identifying the
 728 relevant expert for each part of the prompt. Furthermore, the high quality and coherence of the com-
 729 bined output showcase the effectiveness of the **VCLayer** in maintaining a logical flow and the **DPO**
 730 stage in refining the final generation to align with high-quality, human-like collaborative solutions.
 731 This case study compellingly illustrates that FedWave successfully instills genuine collaborative
 732 intelligence in a decentralized environment.

733
 734
 735
 736
 737
 738
 739
 740
 741
 742
 743
 744
 745
 746
 747
 748
 749
 750
 751
 752
 753
 754
 755

756

757

758

759

760

761

762

763

764

765

766

767

768

769

770

771

772

773

774

775

776

777

778

779

780

781

782

783

784

785

786

787

788

789

790

791

792

793

794

795

796

797

798

799

800

801

802

803

804

805

806

807

808

809

Figure 4: Qualitative comparison of model outputs for a multi-stage task. The prompt requires the model to first act as a marketing expert and then as a supply chain analyst.

Input Prompt	
<p><i>"Based on the latest market analysis indicating a strong demand for sustainable family vehicles, draft a marketing announcement for our new electric SUV, the 'EcoVoyager'. Then, outline the key component requirements for the initial production run that the supply chain team needs to source."</i></p>	
Model	Generated Output
Centralized MAS (Full Data Access)	<p>Marketing Announcement:</p> <p><i>Introducing the all-new EcoVoyager! Experience the future of family adventure with zero emissions. Designed for the conscious modern family with an exceptional range and advanced safety features. Pre-orders open this fall. #EcoVoyager #ElectricSUV #SustainableDriving</i></p> <p>Supply Chain Requirements:</p> <p><i>For the initial production run of 10,000 units, the key components to source are:</i></p> <ol style="list-style-type: none"> <i>1. Battery Packs: High-capacity lithium-ion battery packs meeting our long-range specifications.</i> <i>2. Electric Motors: High-efficiency dual-axle permanent magnet motors for the all-wheel-drive system.</i> <i>3. Infotainment System: The latest-generation central touchscreen with our proprietary OS.</i> <i>4. Chassis Components: Lightweight, high-strength aluminum alloy for the vehicle's unibody frame.</i>
FedAvg	<p>Marketing and Supply Chain Plan:</p> <p><i>Unveil your family's next chapter with the EcoVoyager! Fusing sustainable luxury with thrilling performance, our all-electric SUV is here. Featuring a panoramic sunroof and an interior crafted from premium eucalyptus fiber, the EcoVoyager delivers an impressive range on a single charge. Drive the change. #EcoVoyager #FutureIsElectric #FamilySUV</i></p>
FedWave (Privacy-Preserving)	<p>Marketing Announcement:</p> <p><i>Unveil your family's next chapter with the EcoVoyager! Fusing sustainable luxury with thrilling performance, our all-electric SUV is here. Featuring a panoramic sunroof and an interior crafted from premium eco-friendly materials, the EcoVoyager delivers a 300+ mile range on a single charge. Drive the change. #EcoVoyager #FutureIsElectric #FamilySUV</i></p> <p>Key Component Requirements:</p> <p><i>Supply chain priority sourcing for the initial 10,000-unit production run of the EcoVoyager:</i></p> <ol style="list-style-type: none"> <i>1. Battery System: High-density 'BladeCell' battery modules with advanced integrated thermal management.</i> <i>2. Drivetrain: Next-generation 'WaveDrive' axial flux synchronous motors for superior power delivery (front and rear).</i> <i>3. Central Computing Unit: The 'VisionCore AI' automotive-grade processor for infotainment and driver-assist features.</i> <i>4. Sustainable Textiles: Sourcing of certified eucalyptus fiber textiles for all seating and interior trim..</i>