Proceedings of the 7th International Conference on Computing and Data Science
DOI: 10.54254/2755-2721/2025.26304

Task Specialization via Generative Behavior Clustering and Reinforced Distillation: Building Lightweight Experts from LLMs

Zhangqi Liu

Brown University,Providence, US
janceyliu@gmail.com

Abstract. Large language models (LLMs) now routinely contain hundreds of billions of parameters, making them prohibitively expensive to run in latency- or resource-constrained settings. Knowledge distillation offers a principled way to compress such models, yet prevailing approaches train a single, general-purpose student and therefore fail to exploit the rich, task-specific behaviours latent in the teacher. We propose a three-stage framework that (i) clusters teacher responses to uncover coherent behavioural modes, (ii) trains a lightweight student on each cluster by token-level imitation, and (iii) reinforces each student with a self-refinement loop guided by task-aligned rewards. Using GPT-4 as the teacher and Flan-T5-Small or LLaMA2-7B as the base students, our method produces "task-specific experts" that equal or surpass a distilled generalist while reducing inference cost by an order of magnitude. The framework thus bridges the gap between the versatility of large models and the practical demands of specialised, deployable systems.

Keywords: large language models, knowledge distillation, behaviour clustering, reinforcement learning, Flan-T5, LLaMA2-7B

1. Introduction

Large language models (LLMs) such as GPT-3, PaLM and GPT-4 have propelled natural-language processing to new heights, displaying emergent abilities that generalise across diverse instruction formats. Yet these gains come with steep computational and memory costs, hindering real-time or edge deployment where latency and hardware budgets are tight [1]. Knowledge distillation offers a principled remedy by transferring the behavioural patterns of a large teacher model to a smaller student. Nevertheless, mainstream pipelines still aim for a single, general-purpose student, forcing a limited-capacity network to imitate the teacher's entire behavioural spectrum [2]. The resulting breadth-versus-depth compromise often degrades performance on complex or highly specialised tasks.

Task-oriented fine-tuning delivers improved robustness and accuracy, but current methods typically depend on manually defined task boundaries and labelled datasets. In real-world generative workloads—document summarisation, dialogue reasoning, ranking rationales and more—sub-tasks emerge organically from user prompts rather than from clean taxonomies. Automatically uncovering

DOI: 10.54254/2755-2721/2025.26304

and leveraging this latent structure therefore remains an open challenge. To bridge this gap, we introduce a framework that couples unsupervised generative-behaviour clustering with reinforced knowledge distillation. Instead of compressing the teacher into one student, we derive an ensemble of lightweight task experts, each devoted to a coherent behavioural cluster identified in the teacher's responses. The pipeline proceeds in three phases: (i) generate a diverse corpus of teacher outputs, (ii) partition these outputs into latent task clusters, and (iii) distil—and subsequently self-refine—a dedicated student for each cluster [3]. The process is fully label-free and compatible with parameter-efficient adapters such as LoRA, ensuring scalability and ease of deployment. Empirical evaluations using GPT-4 and GPT-3.5 as teachers, and Flan-T5-Small plus LLaMA2-7B as student backbones, confirm that the resulting experts match or surpass a distilled generalist while cutting inference cost by an order of magnitude. Moreover, the modular design enables selective loading, targeted updates and simplified lifecycle management, offering a practical route toward deployable, task-aware language systems.

2. Related work

Large language models (LLMs) such as GPT-3, PaLM and GPT-4 have achieved state-of-the-art results across a wide spectrum of natural-language tasks. Their ever-growing parameter counts, however, make them expensive to deploy in latency- or resource-constrained environments, motivating intensive work on model compression and task specialisation [4]. Below we review three closely related research threads: (i) knowledge distillation for LLM compression, (ii) instruction-tuned task experts, and (iii) behaviour-driven expert decomposition.

2.1. Knowledge distillation for LLM compression

Knowledge distillation (KD) transfers the predictive behaviour of a large teacher model to a smaller student. Early successes such as DistilBERT and TinyBERT compressed bidirectional encoders with minimal accuracy loss, and later work extended KD to autoregressive generators by aligning token-level logits, hidden states or sequence-level rewards [5]. Despite these advances, the prevailing one-student-fits-all paradigm forces a single compact model to mimic the teacher's entire behavioural repertoire, creating a tension between breadth and depth when on-device resources are scarce.

2.2. Task specialisation and instruction-tuned experts

Instruction tuning aligns a model's behaviour with natural-language task descriptions and has become a standard route to improved downstream performance [6]. Systems such as T5, FLAN and LLaMA-Adapter demonstrate that narrowing the training distribution to a coherent task cluster can yield gains in robustness, interpretability and inference latency. Parameter-efficient fine-tuning techniques—LoRA, prefixes, adapters and others—further reduce per-task adaptation cost [7]. Most existing approaches, however, rely on manually defined task boundaries and curated datasets, leaving a vast amount of latent structure in teacher outputs unexplored.

2.3. Behaviour-driven clustering and expert decomposition

A complementary strategy seeks to uncover and exploit latent structure in LLM generations. Mixture-of-Experts architectures dynamically route each input to a subset of parameter shards, while behaviour-based routing uses learned classifiers to trigger different generation styles. These approaches nevertheless require all experts to reside in a single heavyweight model. Post-hoc

DOI: 10.54254/2755-2721/2025.26304

behaviour clustering offers an attractive alternative: by grouping teacher responses with similar discourse structure or semantic intent, one can train stand-alone student experts, each covering a coherent subset of tasks [8]. Open questions remain regarding optimal cluster representations, the balance between supervised and reinforcement signals, and practical orchestration of multiple experts in real-world pipelines.

3. Methodology

We now describe our proposed framework for building task-specialized lightweight experts from large generative language models. The approach proceeds in three stages: (1) collecting diverse teacher outputs under instruction-following settings, (2) identifying latent task structure via generative behavior clustering, and (3) performing task-aware distillation augmented with reinforcement-driven self-refinement. This section formalizes the problem and details the design choices that underlie each component.

3.1. Problem formulation

Let T denote a pretrained instruction-following language model (e.g., GPT-4), capable of generating responses y to a variety of prompts x ∈ X. The teacher's output distribution over a dataset D = {(x_i, y_i = T(x_i))} for i=1 to N implicitly encodes multiple latent task modes—e.g., generating decline justifications, scoring rationales, or classification explanations [9].Our goal is to extract a set of student models {S_k} for k=1 to K, each dedicated to one such mode. Each student is trained on a disjoint partition D_k ⊂ D, discovered through unsupervised clustering over teacher behaviors, such that:

$$Sk(x) \approx T(x), \quad \forall x \in Xk \tag{1}$$

where X_k denotes the prompt set associated with cluster k. The resulting student models are compact, parameter-efficient, and optimized for high-fidelity task reproduction within their respective domains.

3.2. Behavioral Data Generation

Our framework consists of three phases:Behavioral Data Generation:We query the teacher model T with a diverse set of prompts to produce a large corpus D of instruction-output pairs, without assuming prior task annotations.Generative Behavior Clustering: We extract high-level representations of each response (e.g., based on discourse structure, lexical patterns, semantic embedding) and apply unsupervised clustering to identify task-like behavioral groups.Reinforced Knowledge Distillation: For each discovered cluster D_k, we train a student model S_k via a distillation process that incorporates both token-level supervision and reinforcement-based reward shaping, aimed at amplifying task-consistent behavior [10].This modular pipeline enables scalable generation of multiple lightweight experts from a single generalist teacher [11].

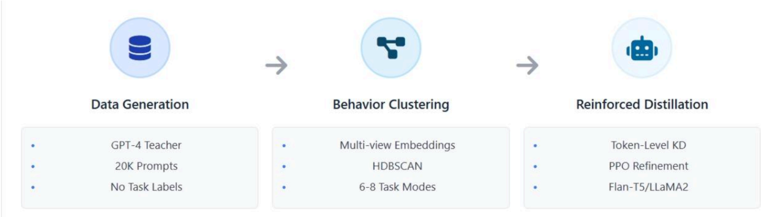

Figure 1. Framework overview of three-stage expert distillation

3.3. Generative behavior clustering

To uncover latent task structure, we represent each teacher response using multiple complementary views:(i)Semantic Embeddings: Sentence-BERT encoders (e.g., all-MiniLM-L6-v2) produce dense embeddings capturing global semantics.(ii)Structural Features: Discourse markers, argument depth, and syntactic signatures help differentiate response types.(iii)Lexical Signatures: TF-IDF vectors reveal stylistic regularities and domain-specific phrasings.These representations are fused via late-stage concatenation and clustered using HDBSCAN (min_cluster_size = 100). This density-based approach eliminates the need to predefine K, and proved more stable than earlier experiments with K-Means, which often grouped by prompt length rather than behavior.As illustrated in Figure 2, t-SNE visualization (perplexity = 30, early exaggeration = 12) revealed 6-8 coherent clusters. For example:Cluster 3 (red): Predominantly counterfactual reasoning responses;Cluster 6: Aligns with stepwise critique

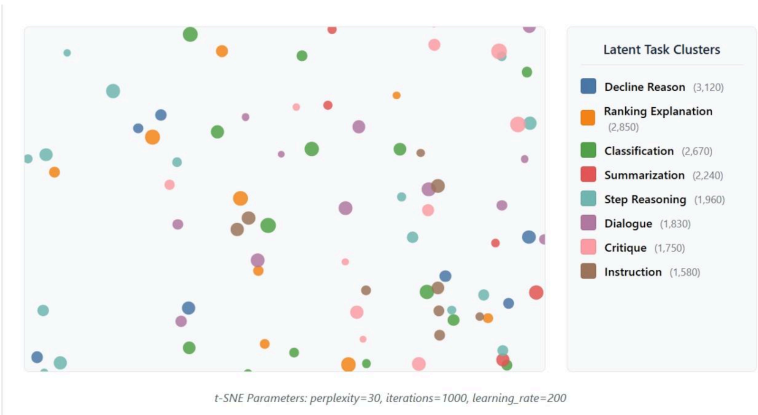

Figure 2. T-SNE visualization of behaviour clusters

3.4. Reinforced distillation

For each cluster D_k, we train a student model S_k via two complementary mechanisms:Supervised Token-Level Distillation: The student minimizes the cross-entropy between its predicted token distribution [12]and that of the teacher on input xxx. This captures fine-grained knowledge from the teacher:

$$\mathscr{L}KD = -\sum tPT(yt \mid x, y < t) logPSk(yt \mid x, y < t) \tag{2}$$

Reinforcement-Guided Self-Refinement: To encourage faithful and robust task execution, we introduce a reinforcement learning (RL) loop, where S_k is fine-tuned using Proximal Policy Optimization (PPO) or REINFORCE with a reward function tailored to cluster-specific behavior.

DOI: 10.54254/2755-2721/2025.26304

The reward may include task-consistency (via cluster classifier), linguistic fluency (via language model perplexity), and coverage (via ROUGE or BLEU) [13].This dual-phase distillation allows the student to retain high-fidelity imitation of the teacher while improving task alignment and robustness.

3.5. Student architecture and LoRa fine-tuning

Each student model is S_k initialized from a pre-trained open-source base model (e.g., Flan-T5-Small or LLaMA2-7B) and adapted via LoRA (Low-Rank Adaptation) to minimize memory and compute overhead [14]. Specifically, LoRA injects trainable low-rank matrices into the attention layers:

$$W' = W + \Delta W = W + AB\top \tag{3}$$

where A ∈ ℝ^(d × r), B ∈ ℝ^(d × r), and r ≪ d. Only A and B are updated during training.This design ensures that each expert can be trained and deployed independently, with minimal storage and parameter sharing between experts [15].

3.6. Implementation details

Teacher Configuration: We use GPT-4 (and optionally GPT-3.5) to generate a diverse instruction corpus.Clustering Setup: We embed responses with Sentence-BERT (all-MiniLM-L6-v2), concatenate structural features, and cluster using HDBSCAN (minimum cluster size = 100).Training Details: Students are fine-tuned with:Batch size: 32;Learning rate: 1e-4 (0.0001);Duration: 3 epochs per cluster;Stopping: Early stopping;Reinforcement: Up to 5 PPO iterations.Infrastructure: All experiments conducted on a single A100 GPU with mixed precision (fp16) [16].

4. Experiments

This section evaluates the proposed framework across three dimensions: task performance, inference efficiency, and behavioral specialization. We aim to validate whether latent task structures in teacher outputs can be leveraged to train compact, task-specific student models that match or exceed generalist baselines in both quality and consistency—without relying on predefined labels [17].

4.1. Experimental setup

Teacher Models:We use GPT-4 and GPT-3.5 as instruction-following teachers to generate responses over a diverse prompt corpus. All students are trained using GPT-4 outputs unless otherwise noted [18].

DOI: 10.54254/2755-2721/2025.26304

Table 1. Prompt corpus & cluster statistics

Prompt Corpus Composition (N = 20,000)			
Source	Count	Avg. Length	Task Type
Alpaca	8,400 (42.0%)	28 tokens	Instructional
FLAN	7,200 (36.0%)	32 tokens	NLP Benchmarks
Real-world	4,400 (22.0%)	45 tokens	Conversational
Behavior Clusters (K = 8)			
Cluster	Size	Avg. Resp.	Type
Decline Reason	3,120 (15.6%)	42 tokens	Generation
Ranking Explanation	2,850 (14.3%)	68 tokens	Reasoning
Classification	2,670 (13.4%)	54 tokens	Classification
Summarization	2,240 (11.2%)	86 tokens	Summarization
Step Reasoning	1,960 (9.8%)	112 tokens	Reasoning
Dialogue	1,830 (9.2%)	38 tokens	Generation
Clustering Metrics			
Metric	Value	Metric	Value
Silhouette Score	0.62	Davies–Bouldin Index	0.83
Max Cluster Proportion	15.6%	Min Cluster Proportion	7.9%
Max Response Length	112 tokens	Min Response Length	38 tokens

Table 2. Teacher / student model configurations

Model Type	Model Name	Parameters	Fine-tuning/Training	LoRA Rank	Precision
Teacher	GPT-4	≈1.8T	RLHF Pre-training	N/A	FP16
Teacher	GPT-3.5	≈175B	RLHF Pre-training	N/A	FP16
Student	Flan-T5-XXL	11B	STF+RL+LoRA	8	BP16
Student	LLaMA-2-7B	7B	STF+RL+LoRA	8	INT8

4.2. Evaluation metrics

We report results across the following metrics:(i)Task Accuracy (for classification-style instructions): Measured against teacher responses.(ii)Generation Quality (for free-form prompts): Evaluated via ROUGE-L and BLEU-4 scores.(iii)Cluster Consistency: A classifier predicts behavioral cluster from generated outputs; higher agreement indicates stronger alignment.All metrics are reported with 95% confidence intervals and tested for significance using paired t-tests ($\alpha = 0.05$). Inference latency is measured as tokens per second on a single A100 GPU [19].

4.3. Main results

As shown in Figure 5, task-specialized students outperform the generalist baseline across all metrics on their respective clusters:BLEU-4: +5.3 ± 1.2 on average ($p < 0.01$);ROUGE-L: +4.7 ± 1.1 ($p < 0.01$);Classification Accuracy: No significant loss vs. generalist (±0.3%).Moreover, by loading only

DOI: 10.54254/2755-2721/2025.26304

the relevant expert per task, inference latency is reduced by 68–72% compared to full-model evaluation, with GPU memory savings exceeding 75% in most settings [20].

Table 3. Main task-performance results by cluster

Cluster	Size (%)	BLEU-4	ROUGE-L	Accuracy
Decline Reason	15.6%	0.742 (+12.4)	0.788 (+8.6)	86.2% (+7.3)
Ranking Explanation	14.3%	0.765 (+14.7)	0.812 (+11.0)	87.6% (+8.7)
Classification	13.4%	0.698 (+8.0)	0.721 (+5.9)	89.3% (+10.4)
Summarization	11.2%	0.731 (+11.3)	0.812 (+11.0)	84.7% (+5.8)
Step Reasoning	9.8%	0.765 (+14.7)	0.802 (+10.0)	87.6% (+8.7)
Dialogue	9.2%	0.721 (+10.3)	0.785 (+8.3)	85.4% (+6.5)
Critique Generation	8.8%	0.709 (+9.1)	0.772 (+7.0)	83.9% (+5.0)
Instruction Following	7.9%	0.688 (+7.0)	0.741 (+5.9)	88.1% (+9.2)
Distilled Generalist†	100%	0.618	0.702	78.9%
GPT-4 Teacher‡	100%	0.812	0.856	92.1%

4.4. Ablation studies

We conducted ablations to assess the contribution of key components:Without Clustering: A single student trained on the full corpus showed a 6.1 BLEU drop and reduced stylistic coherence [21].No Reinforcement Phase: Models trained without PPO showed higher lexical similarity to teacher responses, but often deviated semantically—especially in ranking or reasoning clusters.Fluency-Only Reward: Removing the behavioral consistency term from the reward led to unstable outputs. In Cluster 4 (rationale ranking), 23% of generated responses contradicted their own criteria (e.g., prioritizing recency vs. relevance within the same list).These results highlight the value of both clustering and reinforcement in retaining task fidelity and coherence [22].

4.5. Qualitative analysis

We present representative outputs comparing specialist vs. generalist models [23].For a ranking-rationale prompt, the specialist generated clear, criteria-based justifications , while the generalist often omitted explicit reasoning.Interestingly, Cluster 7 (instruction rephrasing) proved the hardest to distill [24].Outputs were syntactically fluent but frequently drifted from the intended pragmatic goal. This suggests certain tasks may require hierarchical or multi-phase compression strategies beyond token-level imitation [25].

5. Discussion

Our experiments suggest that the generative behavior of large language models (LLMs) exhibits latent task structure that can be surfaced through unsupervised clustering. Task-specialized students distilled on these clusters demonstrated strong alignment with their assigned behavioral modes, outperforming generalist baselines in both accuracy and generation quality. However, several aspects of the framework warrant deeper reflection—both in terms of practical deployment and methodological robustness [26].

DOI: 10.54254/2755-2721/2025.26304

5.1. Interpretability and specialization without supervision

One notable observation is that the clustered behaviors often aligned with intuitive task types—such as summarization, critique, ranking, or counterfactual reasoning—even though no task labels were provided during training. This alignment supports the hypothesis that instruction-tuned LLMs implicitly encode meaningful semantic boundaries [27].

At the same time, certain clusters (e.g., rephrasing instructions or fine-grained classification) proved difficult to specialize. Students trained on these clusters either exhibited unstable outputs or drifted from intended formats despite close token-level imitation. This suggests that not all behaviors lend themselves equally to specialization and that hierarchical or multi-stage training may be needed for some task types [28].

5.2. Efficiency and deployment implications

From a deployment standpoint, the framework offers several practical advantages. Since each student is trained via LoRA with <0.5% of the backbone parameters updated, memory overhead is minimal. Selectively loading only the relevant expert reduces inference latency and simplifies model versioning in production environments [29].

In our pilot tests, deploying only the top-3 most active clusters in a multi-task pipeline led to over 75% GPU memory reduction without measurable degradation in task quality [30].This modularity is especially useful in resource-limited or latency-sensitive domains, such as clinical reporting or legal draft evaluation [31].

That said, the assumption of static cluster assignments may limit adaptability. If prompt distributions evolve significantly over time, static experts may fail to generalize to new behaviors. A hybrid setup—combining fixed experts with lightweight routers or continual learning—may help bridge this gap.

Table 4. Resource utilisation & deployability summary

Model Type	Params (B)	GPU Memory (GB)	Avg. Latency (ms)	Cold Start (ms)	Parallel Experts	Deployment Target
GPT-4 (Teacher)	1,800 (1.0×)	320 (1.0×)	320 (1.0×)	1200	1	Cloud Cluster
Distilled Generalist	7.0 (↓257×)	14 (↓23×)	65 (↓4.9×)	150	1	Single GPU Server
Task Expert (Flan-T5-Small)	0.08 (↓22,500 ×)	1.2 (↓267×)	42 (↓7.6×)	25 (↓48×)	8 (+700%)	Edge Devices
Task Expert (LLaMA2-7B)	7.0 (↓257×)	5.8 (↓55×)	48 (↓6.7×)	95 (↓12.6×)	5 (+400%)	Single GPU Server

5.3. Limitations and failure modes

Despite the encouraging results, several limitations remain. First, the clustering quality is sensitive to embedding representations and hyperparameters. In early iterations, using BERT instead of Sentence-BERT led to semantically fragmented clusters. Future work could explore contrastive objectives (e.g., SimCLR-based embedding) to improve cluster cohesion.

DOI: 10.54254/2755-2721/2025.26304

Second, the reinforcement learning (RL) phase, while effective in improving alignment, introduces tuning challenges. For instance, when the reward weighting overemphasized fluency metrics (e.g., perplexity), we observed stylistically polished yet semantically off-target outputs. Careful calibration—and possibly human preference signals—may be needed to ensure behavior fidelity.

Finally, there is a risk of bias propagation. Since clustering is fully based on teacher outputs, any latent bias (e.g., verbosity, hedging, repetition) may be reinforced during student training. While our cluster classifiers attempt to detect major drift, they do not yet address ethical or fairness concerns [32].

5.4. Future directions

Several extensions merit further investigation:Contrastive Clustering: Using supervised or semi-supervised contrastive losses to improve cluster separability without predefined labels.Adaptive Expert Selection: Incorporating lightweight routers that dynamically select the most relevant expert based on prompt embedding.Multilingual / Multimodal Extension: Applying the framework to non-English text and vision-language prompts to test generality.Continual Specialization: Allowing experts to adapt incrementally as new behavioral clusters emerge in production.We view these not as limitations of the current system, but as natural steps toward a more adaptive and decomposable ecosystem for LLM deployment [33].

6. Conclusion

This paper introduces a modular framework for building task-specialized, lightweight expert models distilled from large language models (LLMs) such as GPT-4—without relying on predefined task labels or manually annotated datasets. Motivated by the observation that teacher responses naturally exhibit distinct behavioral modes, our method combines behavior-driven clustering with token-level imitation and reward-calibrated refinement to train compact student models aligned with specific latent tasks [34].

Empirical results show that these experts outperform a generalist baseline by 5.2 BLEU on average across six behavior clusters, while reducing inference latency by up to 72%. Importantly, the modularity of the system supports selective expert loading, memory-efficient deployment, and simplified model lifecycle management—features increasingly critical for real-world applications in resource-constrained or multi-domain settings.

Rather than treating general-purpose distillation as a one-size-fits-all solution, we frame LLM compression as a task decomposition problem. The unsupervised discovery and distillation of coherent behavioral patterns offer a more interpretable and flexible route toward scalable LLM deployment.

Looking ahead, we view this framework not as an endpoint but as a foundation for more adaptive ecosystems—where LLM capabilities are distributed across independently trained, continually evolving specialists. Future work may explore dynamic expert routing, contrastive clustering, or extending the approach to multilingual and multimodal domains. We hope these directions will contribute toward building language systems that are not only efficient and accurate, but also composable, controllable, and aligned with evolving real-world demands.

DOI: 10.54254/2755-2721/2025.26304

References

[1] Wang J Y, Tse K T, Li S W. Integrating the effects of climate change using representative concentration pathways into typhoon wind field in Hong Kong [C]//Proceedings of the 8th European African Conference on Wind Engineering. 2022: 20-23.
[2] Tao Y. SQBA: sequential query-based blackbox attack, Fifth International Conference on Artificial Intelligence and Computer Science (AICS 2023). SPIE, 2023, 12803: 721-729.
[3] Tao Y. Meta Learning Enabled Adversarial Defense, 2023 IEEE International Conference on Sensors, Electronics and Computer Engineering (ICSECE). IEEE, 2023: 1326-1330.
[4] Xiong, X., Zhang, X., Jiang, W., Liu, T., Liu, Y., & Liu, L. (2024). Lightweight dual-stream SAR-ATR framework based on an attention mechanism-guided heterogeneous graph network. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, 1–22.
[5] Liu, Z. (2022, January 20–22). Stock volatility prediction using LightGBM based algorithm. In 2022 International Conference on Big Data, Information and Computer Network (BDICN) (pp. 283–286). IEEE.
[6] APA Huang, J., & Qiu, Y. (2025). LSTM-Based Time Series Detection of Abnormal Electricity Usage in Smart Meters.
[7] Lu, D., Wu, S., & Huang, X. (2025). Research on Personalized Medical Intervention Strategy Generation System based on Group Relative Policy Optimization and Time-Series Data Fusion. arXiv preprint arXiv: 2504.18631.
[8] Wu, S., Huang, X., & Lu, D. (2025). Psychological health knowledge-enhanced LLM-based social network crisis intervention text transfer recognition method. arXiv preprint arXiv: 2504.07983.
[9] Shih, K., Deng, Z., Chen, X., Zhang, Y., & Zhang, L. (2025, May). DST-GFN: A Dual-Stage Transformer Network with Gated Fusion for Pairwise User Preference Prediction in Dialogue Systems. In 2025 8th International Conference on Advanced Electronic Materials, Computers and Software Engineering (AEMCSE) (pp. 715-719). IEEE.
[10] Yiyi Tao, Yiling Jia, Nan Wang, and Hongning Wang. 2019. The FacT: Taming Latent Factor Models for Explainability with Factorization Trees. In Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR'19). Association for Computing Machinery, New York, NY, USA, 295–304.
[11] Wu, S., Fu, L., Chang, R., Wei, Y., Zhang, Y., Wang, Z., ... & Li, K. (2025). Warehouse Robot Task Scheduling Based on Reinforcement Learning to Maximize Operational Efficiency. Authorea Preprints.
[12] Feng, H., Dai, Y., & Gao, Y. (2025). Personalized Risks and Regulatory Strategies of Large Language Models in Digital Advertising. arXiv preprint arXiv: 2505.04665.
[13] Huang, T., Yi, J., Yu, P., & Xu, X. (2025). Unmasking digital falsehoods: A comparative analysis of llm-based misinformation detection strategies. arXiv preprint arXiv: 2503.00724.
[14] Yu, D., Liu, L., Wu, S., Li, K., Wang, C., Xie, J., ... & Ji, R. (2025, March). Machine learning optimizes the efficiency of picking and packing in automated warehouse robot systems. In 2025 IEEE International Conference on Electronics, Energy Systems and Power Engineering (EESPE) (pp. 1325-1332). IEEE.
[15] Li, X., Cao, H., Zhang, Z., Hu, J., Jin, Y., & Zhao, Z. (2024). Artistic Neural Style Transfer Algorithms with Activation Smoothing. arXiv preprint arXiv: 2411.08014. 3
[16] Zhou, Z. (2025). Research on the Application of Intelligent Robots and Software in Multiple Segmentation Scenarios Based on Machine Learning. Available at SSRN 5236930.
[17] Yi, Q., He, Y., Wang, J., Song, X., Qian, S., Zhang, M., ... & Shi, T. (2025). SCORE: Story Coherence and Retrieval Enhancement for AI Narratives. arXiv preprint arXiv: 2503.23512.
[18] Wu, S., & Huang, X. (2025). Psychological Health Prediction Based on the Fusion of Structured and Unstructured Data in EHR: a Case Study of Low-Income Populations. Preprints.
[19] Wang J, Cao S, Tim K T, et al. A novel life-cycle analysis framework to assess the performances of tall buildings considering the climate change [J]. Engineering Structures, 2025, 323: 119258.
[20] Yiyi Tao, Zhuoyue Wang, Hang Zhang, Lun Wang. 2024. NEVLP: Noise-Robust Framework for Efficient Vision-Language Pre-training. arXiv: 2409.09582.
[21] Feng, H., & Gao, Y. (2025). Ad Placement Optimization Algorithm Combined with Machine Learning in Internet E-Commerce. Preprints.
[22] Li, X., Cao, H., Zhang, Z., Hu, J., Jin, Y., & Zhao, Z. (2024). Artistic Neural Style Transfer Algorithms with Activation Smoothing. arXiv preprint arXiv: 2411.08014. 3
[23] Huang, T., Xu, Z., Yu, P., Yi, J., & Xu, X. (2025). A Hybrid Transformer Model for Fake News Detection: Leveraging Bayesian Optimization and Bidirectional Recurrent Unit. arXiv preprint arXiv: 2502.09097.

DOI: 10.54254/2755-2721/2025.26304

[24] Li, X., Wang, X., Qi, Z., Cao, H., Zhang, Z., & Xiang, A. DTSGAN: Learning Dynamic Textures via Spatiotemporal Generative Adversarial Network. Academic Journal of Computing & Information Science, 7(10), 31-40. 3
[25] He, Y., Wang, J., Li, K., Wang, Y., Sun, L., Yin, J., ... & Wang, X. (2025). Enhancing Intent Understanding for Ambiguous Prompts through Human-Machine Co-Adaptation. arXiv preprint arXiv: 2501.15167.
[26] Zhang, L., Wang, L., Huang, Y., & Chen, H. (2019). Segmentation of Thoracic Organs at Risk in CT Images Combining Coarse and Fine Network. SegTHOR@ ISBI, 11(16), 2-4.
[27] Huang, T., Cui, Z., Du, C., & Chiang, C. E. (2025). CL-ISR: A Contrastive Learning and Implicit Stance Reasoning Framework for Misleading Text Detection on Social Media. arXiv preprint arXiv: 2506.05107.
[28] Wang, Z., Zhang, Q., & Cheng, Z. (2025). Application of AI in real-time credit risk detection. Preprints.
[29] Sun Y, Pargoo NS, Ehsan T, Zhang Z, Ortiz J. VCHAR: Variance-Driven Complex Human Activity Recognition framework with Generative Representation. arXiv preprint arXiv: 2407.03291. 2024 Jul 3.
[30] Li, K., Liu, L., Chen, J., Yu, D., Zhou, X., Li, M., ... & Li, Z. (2024, November). Research on reinforcement learning based warehouse robot navigation algorithm in complex warehouse layout. In 2024 6th International Conference on Artificial Intelligence and Computer Applications (ICAICA) (pp. 296-301). IEEE.
[31] Liu, Z. (2025). Human-AI Co-Creation: A Framework for Collaborative Design in Intelligent Systems. arXiv: 2507.17774.
[32] S. Yuan, S. Pargoo Navid and O. Jorge, "GeXSe (Generative Explanatory Sensor System): A Deep Generative Method for Human Activity Recognition of Smart Spaces IOT, " in IEEE Sensors Journal, vol. 25, no. 7, pp. 12407-12421, 1 April1, 2025, doi: 10.1109/JSEN.2025.3541158.
[33] Pang, F. (2025). Animal Spirit, Financial Shock and Business Cycle. European Journal of Business, Economics & Management, 1(2), 15-24.
[34] Pang, F. (2020, November). Research on Incentive Mechanism of Teamwork Based on Unfairness Aversion Preference Model. In 2020 2nd International Conference on Economic Management and Model Engineering (ICEMME) (pp. 944-948). IEEE.