From Few-Shot Learning to Data-Efficient Intelligence

Yaqing Wang

Beijing Institute of Mathematical Sciences and Applications
wangyaqing@bimsa.cn

Modern artificial intelligence (AI) excels in data-rich environments but continues to struggle with learning and reasoning from only a few examples—a hallmark of human intelligence. My research aims to understand and enable *data-efficient generalization*—the ability of models to learn, adapt, and transfer knowledge under limited supervision. This long-standing challenge has guided the evolution from classical meta-learning to in-context learning in large language models (LLMs), and further to data-efficient adaptation in agent systems. My work provides conceptual, methodological, and application-level advances that unify these paradigms toward more human-like learning.

Few-Shot Learning as a Core Learning Problem. Few-shot learning (FSL) considers a fundamental question: *how can a system learn reliably from limited experience?* My early work (Wang et al. 2020) established a systematic taxonomy, revealing why empirical risk minimization fails under sparse data and how prior knowledge can guide more efficient learning. This perspective frames few-shot generalization as a cornerstone problem that underlies modern efforts toward data-efficient intelligence.

Paradigms for Solving Data-Efficient Generalization. Different learning paradigms have emerged to address this challenge. Our research investigates and connects three major families:

Meta-Learning. Meta-learning, or "learning to learn," enables models to generalize across tasks by leveraging prior experience for rapid adaptation. We develop meta-learning methods (Wu, Wang, and Yao 2025a) that exploit task identity and representation-level supervision to enhance generalization across related tasks, providing a principled mechanism for fast adaptation in low-data regimes.

In-Context Learning in LLMs. In-context learning (ICL) refers to the ability of LLMs to perform new tasks directly from examples presented in their input context, without parameter updates. Our theoretical analysis (Wu, Wang, and Yao 2025b) shows that ICL is not merely analogous to meta-learning but can be viewed as a *data-dependent optimal learning process* implicitly embedded within transformers. We prove that transformer-based ICL can represent the algorithms of classical meta-learners, revealing that it learns optimal adaptation strategies from data. Building on this connection, we further explore how example selection, retrieval, and reasoning-aware prompting improve the data efficiency and stability of ICL, illuminating how LLMs internalize and generalize learning behavior.

Data-Efficient Agents. When AI systems act as agents, the few-shot challenge manifests as the need to rapidly align with user goals, preferences, and evolving contexts. We design data-efficient agent architectures (Nie et al. 2025) that enable such personalized adaptation from minimal interactions, dynamically balancing long-term and short-term user signals. These studies demonstrate how few-shot principles naturally extend into agentic settings, paving the way for responsive and trustworthy AI assistants.

Real-World Scenarios Where Data-Efficient Generalization Matters. Although the underlying problem is conceptual, its impact is grounded in two representative real-world scenarios:

Scientific Scarcity. In domains such as drug discovery, data are inherently scarce due to the high cost, risk, and complexity of experiments. We develop methods (Wang et al. 2021) that learn property-aware molecular representations from minimal samples, demonstrating how data-efficient generalization enables progress in AI for Science.

Efficient Data Utilization. In applications such as recommendation, search, and user modeling, data are abundant but sparsely labeled or dynamically evolving. We design models (Wang et al. 2024) that efficiently utilize limited supervision, adapt as new data arrive, and capture user-, item-, and scenario-specific behaviors. These systems show that data-efficient learning principles remain central even when data volume is high but costly to use effectively.

Future Outlook. Data-efficient generalization remains one of the key challenges in AI, bridging model architectures, learning paradigms, and application domains. My future research aims to advance a unified understanding of data-efficient learning by: (1) incorporating mathematical and physical priors to build interpretable and stable models for scientific discovery, and (2) developing data-efficient approaches for leveraging LLMs and agents that integrate human priors for reliable adaptation. Together, these directions move toward human-like, trustworthy intelligence that learns efficiently from limited experience.

References

Nie, H.; Wang, Y.; Zhou, M.; Pan, F.; Yao, Q.; and Wang, Z. 2025. AdaPA-Agent: A Personalized Agent with Adaptive Preference Arithmetic for Dynamic Preference Modeling. In *Advances in Neural Information Processing Systems*.

Wang, Y.; Abuduweili, A.; Yao, Q.; and Dou, D. 2021. Property-aware relation networks for few-shot molecular property prediction. In *Advances in Neural Information Processing Systems*, 17441–17454.

Wang, Y.; Piao, H.; Dong, D.; Yao, Q.; and Zhou, J. 2024. Warming up cold-start CTR prediction by learning item-specific feature interactions. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 3233–3244.

Wang, Y.; Yao, Q.; Kwok, J. T.; and Ni, L. M. 2020. Generalizing from a few examples: A survey on few-shot learning. *ACM Computing Surveys*, 53(3): 1–34.

Wu, S.; Wang, Y.; and Yao, Q. 2025a. Learning to Learn with Contrastive Meta-Objective. In *Advances in Neural Information Processing Systems*.

Wu, S.; Wang, Y.; and Yao, Q. 2025b. Why In-Context Learning Models are Good Few-Shot Learners? In *International Conference on Learning Representations*.