Mixture-of-Experts based Model Market

Yizhou Ma
School of Electronics and Computer Science, University of Southampton
Southampton, UK
ym7u22@soton.ac.uk

Xikun Jiang
Department of Computer Science, University of Copenhagen
Copenhagen, Denmark
xikun@di.ku.dk

Wenbo Wu
School of Electronics and Computer Science, University of Southampton
Southampton, UK
wenbo.wu@soton.ac.uk

Zhuoqin Yang
School of Computer Science, University of Nottingham,
Ningbo, China
scxzy5@nottingham.edu.cn

Luis-Daniel Ibáñez
School of Electronics and Computer Science, University of Southampton
Southampton, UK
l.d.ibanez@soton.ac.uk

ABSTRACT

The rapid expansion of AI services has exposed emerging limitations in existing model marketplaces, particularly in coordinating heterogeneous model providers, ensuring equitable participation, and maximizing combined capabilities. To address these issues, this paper proposes a novel marketplace architecture based on the Mixture of Experts (MoE) paradigm, utilizing a central gating network to orchestrate specialized experts or models. The gating network processes model service request and data, dynamically selecting suitable experts and allocating weighted data subsets accordingly. This enables efficient collaboration among experts, exploiting their complementary strengths to deliver optimized composite AI services to buyers. Specifically, we present the conceptual framework of the MoE-based marketplace and detail its fundamental operational principles. Furthermore, we discuss key research challenges, including designing an effective gating network that can intelligently route tasks to suitable experts and developing fair, efficient payment strategies. Together, these approaches aim to prevent expert overuse or underuse, thereby ensuring balanced expert engagement and fostering the inclusion of diverse contributors.

PVLDB Workshop Format:
Yizhou Ma, Xikun Jiang, Wenbo Wu, Zhuoqin Yang, and Luis-Daniel Ibáñez. Mixture-of-Experts based Model Market. VLDB 2025 Workshop: 3rd Data EConomy Workshop (DEC).

1 INTRODUCTION

In recent years, data marketplaces have emerged as a foundational component of the data economy, enabling organizations to monetize their data assets and allowing consumers to access high-quality datasets for diverse analytical and business purposes [7, 9, 14, 16]. However, the rising demand for AI and ML has shifted value from raw data to models trained on domain-specific datasets, capable of performing tasks such as classification, prediction, recommendation, and automated reasoning.

This evolution is giving rise to a new paradigm: machine learning model marketplaces [1, 2, 4, 5]. These platforms enable organizations to commercialize pre-trained models, either as a replacement for or a complement to traditional raw dataset sharing. Consumers can then buy these models without the burden of training them.

The convergence of data marketplaces and ML model marketplaces is therefore both natural and necessary. Models rely on high-quality, domain-specific data to achieve strong performance, while curated datasets gain practical value only when they can be used to train models that solve real-world tasks. In this sense, data and models form a mutually reinforcing ecosystem: better data enables better models, and useful models increase the demand for relevant, well-curated data.

In this context, we foresee the increasing adoption of the Mixture of Experts (MoE) [3, 23] pattern as a way to orchestrate different providers of multiple specialized models to meet complex requests from consumers. While prior studies [12, 19, 21] have explored competition and incentive alignment in relatively homogeneous settings, the effective integration and incentivization of heterogeneous AI service providers, with diverse architectures and competencies, remains a core challenge. Existing marketplace designs, whether centralized, decentralized [4, 11] or based on naive competition [12], often lack the mechanisms needed to support fine-grained cooperation among diverse service providers, limiting their ability to deliver efficient and equitable AI services.

To bridge this gap, we propose a novel Mixture-of-Experts (MoE) based model marketplace architecture. The MoE marketplace consists of a pool of experts, each specializing in heterogeneous machine learning services. A central gating network intelligently assigns tasks and allocates subsets of the training set to experts, the Gating Network assigns weights to the experts' contributions based on the specific characteristics of the incoming model requests. This structure facilitates effective cooperation among experts, maximizing their combined strengths to deliver a better model to buyers.

In this paper we present an initial architecture of an MoE Market place and an early stage exploration of MoE-based marketplace. We aim to highlight open challenges in coordinating and incentivizing diverse experts and examine how cooperation and competition among heterogeneous providers can be balanced to ensure fair and high-quality service delivery. We further discuss the broader implications of this structure for the emerging AI service economy.

The remainder of this paper is structured as follows. Section 2 reviews related work, covering existing paradigms for data and model marketplaces, as well as the foundational concepts of the Mixture of Experts (MoE) pattern. Section 3 introduces our proposed MoE-based model market architecture, actors and operational workflow. Section 4 discusses the key technical and incentive-related challenges. Finally, Section 5 concludes the paper.

2 RELATED WORK

To contextualize our proposed Mixture-of-Experts (MoE) market architecture, we first review existing paradigms. This section begins with an overview of general data marketplaces, followed by a discussion on the model marketplaces, and finally, an introduction to the MoE pattern and its relevance.

Data Marketplaces: A data marketplace is broadly defined as an online venue that facilitates the commoditization of data among them. These platforms aim to strike a balance between individuals' potential privacy loss and the broader utility derived from the use of private data, often through the monetization of data.

The typical ecosystem of a data marketplace involves three main types of participants, as illustrated by Zhang *et al.* [24].

- *Data Owners:* These are entities, such as individuals or organizations, who possess data and are willing to monetize it. This can range from personal data (social, financial, health) to corporate data about users, their demographics, preferences, and behaviors.
- *Data Consumers:* These participants seek external data to enhance their decision-making, product design, service offerings, or customer management. They can include advertisers, software developers, retailers, and manufacturers.
- *Data Brokers:* Often acting as an intermediary, the data broker collects, integrates, stores, and redistributes data, generating profit from these operations. The broker can be an integral component of the marketplace platform or an external third party.

Data marketplaces can vary significantly in their market structure, which plays a key role in determining how data value is realized. Zhang *et al.* [24] categorize these structures primarily into sell-side markets, buy-side markets, and two-sided markets. In a sell-side market, data from multiple sources are integrated and sold to consumers, while a buy-side market enables entities to monetize their internal data by allowing brokers to procure it. Two-sided markets combine these aspects, with a data broker often facilitating transactions between owners and consumers. Such two-sided markets can be further distinguished as centralized, where all participants trade data through the broker, or decentralized, where the broker might provide a platform but data owners and consumers can transact more directly if they are members. These foundational structures set the stage for understanding more specialized marketplaces, such as those focusing on AI models.

Model Marketplaces: Recognizing that the value often lies not just in raw data but in the insights and predictive power derived from it, model marketplaces have emerged as a significant evolution. These platforms focus on the buying and selling of trained machine learning (ML) models as commodities.

Early conceptualizations of model marketplaces highlighted the need to transition from markets selling only data to those that can directly sell ML models. Chen *et al.* [4] introduced a formal Model-Based Pricing (MBP) framework, where, instead of pricing the data, ML model instances are directly priced based on their quality, typically accuracy. Their approach involves a seller (providing the dataset), a buyer, and a central broker who mediates the sale. The broker often trains an optimal model and then generates different versions (e.g., by injecting Gaussian noise) to offer various price-accuracy trade-offs to buyers. This MBP framework also formally addresses desiderata such as preventing arbitrage opportunities. Cong *et al.* [6] further survey pricing for ML models for end-users as a key step in ML deployment pipelines, noting challenges in versioning and arbitrage avoidance. Building on these foundational MBP concepts, Liu *et al.* [11] proposed a more comprehensive end-to-end model marketplace that considers the needs of all three entities: data owners, the broker, and model buyers. This "Dealer" framework allows data owners to set usage restrictions (e.g., via Differential Privacy) and receive fair compensation (e.g., using Shapley value) for their data's contribution to the models sold in the market. The broker, in this single-entity setup, trains a series of differentially private models and uses sophisticated algorithms for revenue maximization while respecting arbitrage-freeness.

The advent of Federated Learning (FL) introduced a new dimension to model marketplaces by enabling collaborative model training without direct raw data exchange, thus addressing some of the critical privacy concerns inherent in centralized data collection. Sun *et al.* [20] proposed DEVELOP, a broker centric FL model marketplace that incentivizes data owners to participate in Differentially Private FL (DPFL) by explicitly accounting for their privacy costs. The broker in DEVELOP then undertakes optimal model versioning and pricing to maximize its profit from selling these collaboratively trained models. Similarly, Pan *et al.* [15] addressed the challenge of non-IID data and diverse customer requirements in FL model marketplaces by proposing a privacy-preserving trainer recruitment mechanism. Their hierarchical mechanism selects optimal trainers based on task preferences, data distributions (approximated using coarse-grained information and compared against a customer's template dataset using KL divergence), and data sizes, aiming to improve the quality of purchased models.

Mixture of Experts: The Mixture of Experts (MoE) architecture is a machine learning paradigm that employs a "divide and conquer" strategy to handle complex and diverse datasets [23]. The basic design of an MoE model involves multiple specialized sub-models, termed "experts," and a "gating function". This gating function dynamically selects and activates the most relevant subset of experts to process a given input data. This selective activation allows MoE models to significantly expand their capacity and handle diverse knowledge domains without a proportional increase in computational costs, thereby achieving a balance between performance and efficiency [13]. The MoE pattern has proven effective in various machine learning paradigms, including continual learning [10], meta-learning [8], multi-task learning [17], reinforcement learning [22], and federated learning [18].

However, to the best of our knowledge, the direct application of the MoE pattern—with its distinct components of specialized

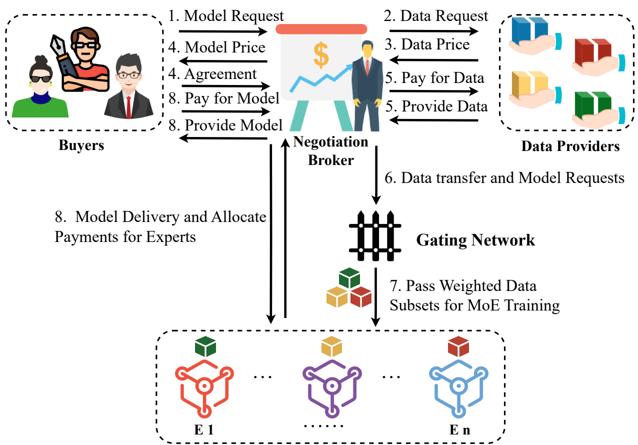

Figure 1: the structure of MoE Model Market

experts and a selective gating mechanism—as an overarching architectural principle for designing and operating a marketplace of autonomous, heterogeneous service providers remains largely unexplored. In typical MoE applications, the "experts" and the "gating network" are integral, co-trained components of a unified system designed to perform a specific learning task. Our work distinguishes itself by transposing this powerful pattern to a market context. We envision a marketplace where the "experts" are independent, specialized service-providing brokers, and the "gating network" functions as an intelligent market coordinator that dynamically orchestrates collaboration among these brokers to fulfill complex buyer requests. This novel application of the MoE concept to a multi-level market structure forms the core of our proposed research.

3 MOE-BASED MODEL MARKET

This section details our proposed marketplace architecture, designed to facilitate the trading of AI-driven model services by leveraging a Mixture of Experts (MoE) paradigm. The architecture, as depicted in Figure 1, integrates several key participants and a structured workflow to orchestrate the delivery of specialized AI services and manage the associated transactions. Our primary aim with this design is to create a flexible and potentially more efficient market by intelligently coordinating diverse expert capabilities.

3.1 Main Actors and Their Roles

Our proposed MoE-based model marketplace (Figure 1) is conceptualized as an ecosystem in which Buyers request bespoke AI model services. These services are delivered through the coordinated efforts of multiple Experts ($E_1, \cdots, E_n$), orchestrated by a central Negotiation Broker. A critical component in this architecture is the Gating Network, which processes model service request and data, dynamically selecting suitable experts and allocating weighted data subsets accordingly. This enables efficient collaboration among experts, exploiting their complementary strengths to deliver optimized composite AI services to buyers. The ultimate goal is to generate a tailored model for the Buyer, with the Negotiation Broker managing the entire pipeline from service request to revenue distribution of experts.

We identify the following actors:

- *Buyers:* Entities that commission specific AI-powered services or model outputs. For example, a buyer could be a financial institution requesting a model for credit risk assessment, or a healthcare provider seeking diagnostic predictions from medical imaging data. As shown in Figure 1, Buyers initiate the process by submitting a *Model Request* to the Negotiation Broker (Step 1), often accompanied by expectations regarding the *Model Price* (Step 4).
- *Negotiation Broker:* The central coordinator responsible for: **(i) Interfacing with Buyers**: receiving requests, negotiating pricing, delivering the final model, and processing payments (Steps 1, 8); **(ii) Managing Data Acquisition**: issuing data requests to Data Providers, negotiating prices, handling payment, and collecting datasets (Steps 2–5); **(iii) Interfacing with the Experts**: forwarding acquired data to the Gating Network (Step 6); and **(iv) Delivery and Revenue Allocation**: delivering the model and distributing payments among the contributing experts based on their contribution to the final output.(Step 8).
- *Data Providers:* Entities that supply the data necessary to train the models to fulfill buyers' requests. They receive *Data Requests* (Step 2) and respond with *Data Price* quotes (Step 3). Upon agreement and payment (Step 5), they deliver the required data to the Negotiation Broker.
- *Gating Network:* An intelligent routing component owned and operated by the Negotiation Broker. Upon receiving data from the broker (Step 6), the Gating Network analyzes the information and assigns *Weighted Data Subsets* to selected experts within the MoE block (Step 7). This mechanism determines each expert's relevance and scope of participation, effectively customizing the data pipeline to match expert capabilities.
- *Experts:* A pool of heterogeneous models or specialized service modules. Experts receive task-specific weighted data from the Gating Network (Step 7) and collaboratively generate the final AI output. Only experts selected by the Gating Network participate in the processing and are eligible for compensation. Experts not selected (i.e., with zero weight) do not receive data for that task and are excluded from both execution and payment for that request (Step 8).

3.2 Marketplace Workflow and Interactions

The operational flow of the MoE-based service trading marketplace, as illustrated in Figure 1, unfolds through the following sequence of interactions:

- *Model Request Initiation:* A Buyer initiates the process by submitting a *Model Request* to the Negotiation Broker (Step 1). A model request may include quality expectations and maximum budget, that are input to steps 2-4 of the Figure 1.
- *Data Sourcing Strategy:* Based on the model request, the Negotiation Broker identifies the required data and sends

a *Data Request* to relevant Data Providers (step 2). Data providers return a price quotation of the desired data.

- *Service Agreement and Data Acquisition:* The Negotiation Broker evaluates the *Data Price* quoted by the Data Providers and the specifications in the Buyer's *Model Request* to propose an overall *Model Price* to the Buyer (step 4). If the Buyer does not agree with the offered price, the transaction is terminated. Upon reaching an agreement (step 4) with the model buyer, the negotiation broker triggers the acquisition of data from the data providers (step 5), the Negotiation Broker completes the transaction by issuing a *Pay for Data* transfer to the Data Providers, who then *Provide Data* in return (step 5).
- *Data Preprocessing and Routing:* The acquired data and requested model specification are passed to the Gating Network (step 6), which is responsible for interpreting the task context and preprocessing the data accordingly.
- *Expert Engagement:* The Gating Network analyzes the data and distributes *Weighted Data Subsets for MoE Training* to selected experts ($E_1 \ldots E_n$) in the Mixture-of-Experts block (step 7). This step is critical, as it operationalizes the MoE principle by assigning task-specific data to the most relevant experts. Experts not selected (i.e., assigned zero weight) do not participate in this training instance, nor do they receive compensation for it.
- *Service Delivery and Financial Settlement:* The resulting model or service output is delivered to the Buyer (*Provide Model*, step 8), who completes the transaction through payment (*Pay for Model*, step 8). Simultaneously, the Negotiation Broker performs the *Model Delivery and Allocate Payments for MoE* (step 8), compensating participating experts based on their assigned contribution weights and the revenue generated.

This architecture establishes a dynamic and modular environment in which heterogeneous AI expertise can be efficiently coordinated and monetized. The separation of negotiation, data acquisition, intelligent gating, and expert execution enables specialized roles at each stage.

4 CHALLENGES AND FUTURE DIRECTIONS

While the proposed Mixture-of-Experts (MoE) marketplace architecture presents a promising framework for facilitating collaboration among heterogeneous experts and delivering high-quality AI services,its implementation raises several challenges:

Revenue Allocation: A primary challenge lies in establishing a fair and efficient profit allocation mechanism among diverse experts who collaboratively fulfill a buyer's request. In such heterogeneous settings, accurately quantifying each expert's contribution to the final aggregated output is inherently difficult. Unlike homogeneous model training scenarios, where performance metrics such as task-specific accuracy, confidence calibration, computational efficiency, robustness, and complementarity are more directly comparable, MoE marketplaces must account for varying capabilities, costs, and roles. Traditional metrics, such as individual accuracy or independently computed Shapley values, often fail to reflect experts' true added value within the context of the collaborative generation of a requested model within the collaborative operation. An effective allocation mechanism must go beyond simplistic measures, capturing nuanced contributions while incentivizing sustained high quality participation from all experts, including those with niche specializations who may be invoked less frequently but are critical in certain requests. Failure to design such a mechanism may result in detrimental effects such as "free-riding" by underperforming experts or reward concentration among a few dominant contributors, ultimately destabilizing participation. The key challenge is thus to reward impactful contributors while avoiding the exclusion of less frequently selected, yet potentially valuable, experts.

Gating Network Design: An essential component of our model is the Gating Network. The broker must solve the problem of training a gating network that, given a buyer's model request, a heterogeneous pool of experts and the data acquired from data providers, outputs a weight distribution over the experts and weighted subdatasets. Training such a Gating Network is challenging for several reasons: First, it must accurately evaluate the relative strengths of diverse experts across a wide variety of tasks, despite differences in architecture, specialization, and performance scales. Second, defining an effective supervision method is non-trivial, especially when the optimal combination of experts is unknown or changes dynamically with incoming requests. Third, the network must avoid developing persistent biases that over-allocate tasks to a small subset of "star" experts, which risks excluding less frequently activated—but potentially critical—participants and eroding the marketplace's diversity and flexibility. Beyond these learning challenges, the negotiation broker must manage operational concerns such as onboarding new experts, enforcing quality control, and maintaining robustness under dynamic conditions, including shifting requests and expert turnover. These requirements collectively impose significant demands on the learning strategy.

Validation Methods: Finally, the design of a proper empirical validation of an MoE market is challenging. Accurately evaluating expert quality is inherently multidimensional, encompassing not only accuracy, but also confidence calibration, computational efficiency, robustness, and complementarity. Developing standardized benchmarks and online metrics that comprehensively reflect these aspects is crucial to evaluate both the intelligence of the Gating Network and the fairness of any profit allocation scheme. Furthermore, experimental design must include tasks and datasets that genuinely require heterogeneous collaboration, avoiding scenarios where a single expert could trivially outperform the ensemble. Validating the effectiveness of the Gating Network demands more than overall performance metrics, it requires isolating and quantifying the contribution of its decision making relative to simpler heuristics, particularly under uncertainty. Simulating long-term market dynamics, such as broker exclusion trends, shifts in participation under varying allocation policies, and scalability under an increasing expert population, further complicates the evaluation. These experiments require precise control, realistic assumptions, and substantial computational resources to faithfully capture the complex and evolving nature of the MoE-based marketplace.

5 CONCLUSION

In this paper, we introduced a conceptual framework for a novel Mixture-of-Experts (MoE)-based model marketplace, designed to overcome limitations in existing market paradigms by fostering effective collaboration and ensuring fair participation among heterogeneous AI service providers. We outlined an architecture where specialized experts are intelligently coordinated by a central gating network to deliver composite AI services. The core innovation lies in extending the MoE pattern, traditionally applied within individual models, to a broader inter-agent market structure, aiming to harness the collective intelligence of diverse AI capabilities.

We identified several key challenges in realizing this vision, including the design and training of an effective gating network capable of accurately assessing and routing requests to heterogeneous experts and allocating weighted data subsets accordingly, as well as the development of fair and efficient revenue allocation mechanisms that reward diverse contributions while preventing market imbalances. We also discussed the complexities of empirically validating such a dynamic and multifaceted marketplace.

As this research is in its early stages, the proposed framework and outlined challenges are intended to serve as a foundation for further development and exploration. Our primary goal is to share our vision for a more cooperative and specialized AI service economy and to invite feedback from the broader research community. We believe that insights into robust gating network strategies, equitable profit-sharing mechanisms, and rigorous experimental methodologies will be critical to advancing this work.

ACKNOWLEDGMENTS

This work was partially supported by the European Union's Horizon Europe research and innovation actions under Grant Agreement Nº 101093216 (UPCAST project).

REFERENCES

[1] Anish Agarwal, Munther Dahleh, and Tuhin Sarkar. 2019. A marketplace for data: An algorithmic solution. In *Proceedings of the 2019 ACM Conference on Economics and Computation*. 701–726. https://doi.org/10.1145/3328526.3329589

[2] Santiago Andrés Azcoitia and Nikolaos Laoutaris. 2022. A survey of data marketplaces and their business models. *ACM SIGMOD Record* 51, 3 (2022), 18–29. https://doi.org/10.1145/3572751.3572755

[3] Weilin Cai, Juyong Jiang, Fan Wang, Jing Tang, Sunghun Kim, and Jiayi Huang. 2025. A survey on mixture of experts in large language models. *IEEE Transactions on Knowledge and Data Engineering* (2025). https://doi.org/10.1109/TKDE.2025.3554028

[4] Lingjiao Chen, Paraschos Koutris, and Arun Kumar. 2019. Towards model-based pricing for machine learning in a data marketplace. In *Proceedings of the 2019 international conference on management of data*. 1535–1552. https://doi.org/10.1145/3299869.3300078

[5] Lingjiao Chen, Hongyi Wang, Leshang Chen, Paraschos Koutris, and Arun Kumar. 2019. Demonstration of nimbus: Model-based pricing for machine learning in a data marketplace. In *Proceedings of the 2019 International Conference on Management of Data*. 1885–1888. https://doi.org/10.1145/3299869.3320231

[6] Zicun Cong, Xuan Luo, Jian Pei, Feida Zhu, and Yong Zhang. 2022. Data pricing in machine learning pipelines. *Knowledge and Information Systems* 64, 6 (2022), 1417–1455. https://doi.org/10.1007/s10115-022-01679-4

[7] Weiqi Dai, Chunkai Dai, Kim-Kwang Raymond Choo, Changze Cui, Deiqing Zou, and Hai Jin. 2019. SDTE: A secure blockchain-based data trading ecosystem. *IEEE Transactions on Information Forensics and Security* 15 (2019), 725–737. https://doi.org/10.1109/TIFS.2019.2928256

[8] Yongxing Dai, Xiaotong Li, Jun Liu, Zekun Tong, and Ling-Yu Duan. 2021. Generalizable person re-identification with relevance-aware mixture of experts. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 16145–16154.

[9] Raul Castro Fernandez, Pranav Subramaniam, and Michael J Franklin. 2020. Data Market Platforms: Trading Data Assets to Solve Data Problems. *Proceedings of the VLDB Endowment* 13, 11 (2020).

[10] Heinke Hihn and Daniel A Braun. 2021. Mixture-of-variational-experts for continual learning. *arXiv preprint arXiv:2110.12667* (2021). https://doi.org/10.48550/arXiv.2110.12667

[11] Jinfei Liu, Jian Lou, Junxu Liu, Li Xiong, Jian Pei, and Jimeng Sun. 2021. Dealer: An end-to-end model marketplace with differential privacy. *Proceedings of the VLDB Endowment* 14, 6 (2021).

[12] Yizhou Ma, Xikun Jiang, Evan W Wu, Luis-Daniel Ibáñez, and Jian Shi. 2024. Model-based Data Markets: A Multi-Broker Game Theoretic Approach. In *2024 IEEE 23rd International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom)*. IEEE, 2293–2301. https://doi.org/10.1109/TrustCom63139.2024.00316

[13] Siyuan Mu and Sen Lin. 2025. A comprehensive survey of mixture-of-experts: Algorithms, theory, and applications. *arXiv preprint arXiv:2503.07137* (2025). https://doi.org/10.48550/arXiv.2503.07137

[14] Thanh Linh Nguyen, Lam Nguyen, Thong Hoang, Dilum Bandara, Qin Wang, Qinghua Lu, Xiwei Xu, Liming Zhu, and Shiping Chen. 2025. Blockchain-empowered trustworthy data sharing: Fundamentals, applications, and challenges. *Comput. Surveys* 57, 8 (2025), 1–36. https://doi.org/10.1145/3718082

[15] Yanghe Pan, Zhou Su, and Yuntao Wang. 2023. Privacy-Preserving Trainer Recruitment in Model Marketplace of Federated Learning. In *2023 IEEE Intl Conf on Dependable, Autonomic and Secure Computing, Intl Conf on Pervasive Intelligence and Computing, Intl Conf on Cloud and Big Data Computing, Intl Conf on Cyber Science and Technology Congress (DASC/PiCom/CBDCom/CyberSciTech)*. IEEE, 0362–0366. https://doi.org/10.1109/DASC/PiCom/CBDCom/Cy59711.2023.10361386

[16] Elena G Popkova and Kantoro Gulzat. 2019. Technological revolution in the 21st century: digital society vs. artificial intelligence. In *Institute of Scientific Communications Conference*. Springer, 339–345. https://doi.org/10.1007/978-3-030-32015-7_38

[17] Moritz Reuss, Jyothish Pari, Pulkit Agrawal, and Rudolf Lioutikov. 2024. Efficient Diffusion Transformer Policies with Mixture of Expert Denoisers for Multitask Learning. *arXiv preprint arXiv:2412.12953* (2024). https://doi.org/10.48550/arXiv.2412.12953

[18] Shangchao Su, Bin Li, and Xiangyang Xue. 2023. One-shot federated learning without server-side training. *Neural Networks* 164 (2023), 203–215. https://doi.org/10.1016/j.neunet.2023.04.035

[19] Peng Sun, Haoxuan Che, Zhibo Wang, Yuwei Wang, Tao Wang, Liantao Wu, and Huajie Shao. 2021. Pain-FL: Personalized privacy-preserving incentive for federated learning. *IEEE Journal on Selected Areas in Communications* 39, 12 (2021), 3805–3820. https://doi.org/10.1109/JSAC.2021.3118354

[20] Peng Sun, Xu Chen, Guocheng Liao, and Jianwei Huang. 2022. A profit-maximizing model marketplace with differentially private federated learning. In *IEEE INFOCOM 2022-IEEE Conference on Computer Communications*. IEEE, 1439–1448. https://doi.org/10.1109/INFOCOM48880.2022.9796833

[21] Xuezhen Tu, Kun Zhu, Nguyen Cong Luong, Dusit Niyato, Yang Zhang, and Juan Li. 2022. Incentive mechanisms for federated learning: From economic and game theoretic perspective. *IEEE transactions on cognitive communications and networking* 8, 3 (2022), 1566–1593. https://doi.org/10.1109/TCCN.2022.3177522

[22] Ziqiang Wei, Deng Chen, Yanduo Zhang, Dawei Wen, Xin Nie, and Liang Xie. 2025. Deep reinforcement learning portfolio model based on mixture of experts. *Applied Intelligence* 55, 5 (2025), 347. https://doi.org/10.1007/s10489-025-06242-6

[23] Seniha Esen Yuksel, Joseph N Wilson, and Paul D Gader. 2012. Twenty years of mixture of experts. *IEEE transactions on neural networks and learning systems* 23, 8 (2012), 1177–1193. https://doi.org/10.1109/TNNLS.2012.2200299

[24] Mengxiao Zhang, Fernando Beltrán, and Jiamou Liu. 2023. A survey of data pricing for data marketplaces. *IEEE Transactions on Big Data* 9, 4 (2023), 1038–1056. https://doi.org/10.1109/TBDATA.2023.3254152