

# Multi-Modal Multi-Task Federated Foundation Models for Next-Generation Extended Reality Systems: Towards Privacy-Preserving Distributed Intelligence in AR/VR/MR

Fardis Nadimi<sup>1</sup>, Payam Abdisarabshali<sup>1</sup>, Kasra Borazjani<sup>1</sup>, Jacob Chakareski<sup>2</sup>, and Seyyedali Hosseinalipour<sup>1</sup>

<sup>1</sup>University at Buffalo–SUNY, Department of Electrical Engineering, Buffalo, NY, USA

<sup>2</sup>New Jersey Institute of Technology (NJIT), Newark, New Jersey, USA

## ABSTRACT

Extended reality (XR) systems, which consist of virtual reality (VR), augmented reality (AR), and mixed reality (MR), offer a transformative interface for immersive, multi-modal, and embodied human-computer interaction. In this paper, we envision that multi-modal multi-task (M3T) federated foundation models (FedFMs) can offer transformative capabilities for XR systems through integrating the representational strength of M3T foundation models (FMs) with the privacy-preserving model training principles of federated learning (FL). We present a modular architecture for FedFMs, which entails different coordination paradigms for model training and aggregations. Central to our vision is the codification of XR challenges that affect the implementation of FedFMs under the **SHIFT** dimensions: (1) **S**ensor and modality diversity, (2) **H**ardware heterogeneity and system-level constraints, (3) **I**nteractivity and embodied personalization, (4) **F**unctional/task variability, and (5) **T**emporality and environmental variability. We illustrate the manifestation of these dimensions across a set of emerging and anticipated applications of XR systems. Finally, we propose evaluation metrics, dataset requirements, and design tradeoffs necessary for the development of resource-aware FedFMs in XR. This perspective aims to chart the technical and conceptual foundations for context-aware privacy-preserving intelligence in the next generation of XR systems.

## Introduction

Extended reality (XR) is a spectrum of immersive technologies (depicted in Fig. 1), comprising synthetic *virtual reality* (VR) scenes, *augmented reality* (AR) overlays anchored to the physical world, and *mixed/merged reality* (MR) where virtual and real objects coexist and interact with each other<sup>1</sup>. XR systems are envisioned as the 4<sup>th</sup> major disruptive technology wave after PC, Internet/Web, and smartphones<sup>2</sup>, and are expected to grow their market share from \$20B in 2025 to \$123B in 2032, affecting the smartphone market share along the way<sup>3</sup>. A defining feature of XR systems is their capacity to integrate *multiple sensory modalities* in real-time, such as vision, audio, haptics, motion, and spatial mapping<sup>4</sup>. In Table 1, we provide a structured representation of the current major XR devices and their capabilities in collecting various data modalities. This multi-modal fusion enables XR platforms to support a *diverse array of downstream tasks* across various domains. In particular, XR holds tremendous potential for revolutionizing healthcare (e.g., remote surgical guidance), education (e.g., virtual lab simulations), manufacturing (e.g., real-time equipment fault detection), and entertainment (e.g., adaptive gameplay) systems<sup>5</sup>.

Complementing the active research on enhancing the practicality of XR systems<sup>6–9</sup>, in this work, we propose a new direction that can catalyze a major leap forward in XR intelligence: *the integration of multi-modal multi-task federated foundation models*. To build a case for this integration, in the following, we articulate the motivation of this study through a curated set of guiding questions.

### **What Recent Advancements in ML Provide Unique Capabilities for Next-Generation XR Systems?**

Machine learning (ML) and subsequently multi-modal ML, techniques have become integral components of XR ecosystems<sup>10</sup>. To demonstrate this, in Table 2, we summarize different XR-specific ML tasks and their typically utilized data modalities. In recent years, ML has witnessed transformative progress, owing to the emergence of *foundation models* (FMs), which are large-scale models capable of fast adaptations to a wide range of downstream tasks (e.g., via zero/few-shot learning)<sup>11</sup>. These models encompass a spectrum of architectures, from *diffusion models*<sup>12</sup> and *large language models* (LLMs)<sup>13</sup> to the most recent *multi-modal multi-task* (M3T) FMs<sup>14</sup>. In particular, diffusion models (e.g., DALL·E 3) are generative models that can generate realistic data, such as images or audio. Also, LLMs such as GPT-3, BERT, and PaLM are trained on massive text corpora to predict or generate coherent language. Building on these, M3T FMs push the boundaries further by integrating *multiple input modalities* (e.g., vision, language, and audio) within a unified model architecture designed to concurrently perform a *variety*

**Table 1.** Representative AR/VR/MR devices, along with smart glasses, and their hardware and deployment characteristics. *Note:* Smart glasses are lightweight, often AI-powered wearables designed primarily for visual and auditory augmentation (e.g., media, notifications, translation), unlike AR/VR/MR headsets which offer fully immersive or spatially anchored environments with advanced interaction capabilities such as hand/eye tracking and spatial mapping.

XR Device	Type	Typical Deployment Environment	Modalities Collected	Notes
Apple Vision Pro	MR	Home, Offices	Visual, Audio, Eye Tracking, Hand Tracking, Spatial Mapping, Motion, Depth Sensing	Ultra-high resolution; advanced AR/VR passthrough
HTC Vive XR Elite	VR / MR	Home, Labs, Enterprise	Visual, Audio, Motion, Hand Tracking, Depth Sensing, Spatial Mapping	Modular standalone and PC VR; compact design
Magic Leap 2	AR	Industrial, Healthcare, Enterprise	Visual, Audio, Eye Tracking, Hand Tracking, Spatial Mapping, Motion, Depth Sensing	Lightweight AR with dynamic dimming
Meta Quest 3	VR / MR	Home, Gaming, Offices	Visual, Audio, Motion, Hand Tracking, Depth Sensing, Spatial Mapping	Affordable standalone with color passthrough
Meta Quest Pro	MR	Offices, Enterprise, Productivity	Visual, Audio, Eye Tracking, Hand Tracking, Spatial Mapping, Motion, Depth Sensing	Premium MR with enhanced passthrough and avatars
Microsoft HoloLens 2 (discontinued in 2024)	AR	Hospitals, Industrial, Education	Visual, Audio, Eye Tracking, Hand Tracking, Spatial Mapping, Motion, Depth Sensing	Advanced AR with spatial anchors
Ray-Ban Meta Glasses	Smart Glasses	Everyday, Mobile, Social	Visual, Audio	Stylish glasses for immersive sensing and AI interaction
RealWear Navigator 520	AR	Industrial, Fieldwork	Visual, Audio, Motion, Thermal	Rugged, hands-free AR with thermal imaging option for industrial use
Varjo XR-4	VR / MR	Enterprise, Simulation, Training	Visual, Audio, Eye Tracking, Hand Tracking, Spatial Mapping, Motion, Depth Sensing	Human-eye resolution VR/MR; used in aviation, design, and defense training
Vuzix Blade 2	AR	Industrial, Fieldwork	Visual, Audio, Motion	Hands-free smart glasses with enhanced security
XReal Air	AR	Home, Travel, Entertainment	Visual, Motion, Head Tracking	Lightweight high-resolution AR glasses; ideal for media consumption and productivity
Samsung Project Moohan (upcoming)	MR	Home, Enterprise, Productivity	Visual, Audio, Eye Tracking, Hand Tracking, Spatial Mapping, Motion, Depth Sensing	Lightweight Android XR headset with Gemini AI integration and Snapdragon XR2 Gen 2
Google AI Glasses (upcoming)	Smart Glasses	Everyday, Travel, Social	Visual, Audio	Stylish smart glasses with in-lens display; Gemini AI for live translation and contextual assistance
Apple Smart Glasses (upcoming)	Smart Glasses	Everyday, Productivity	Visual, Audio	AI-powered glasses with cameras, microphones, and speakers; photo capture, translation, and navigation



**Figure 1.** Various types of XR: (a) **Augmented Reality**, where digital elements are embedded into the environment surrounding a user, (b) **Virtual Reality**, where a user is fully immersed into a completely synthetic environment and interacts with various elements therein, and (c) **Mixed Reality**, which blends digital and physical worlds where a user can interact with digital objects.

of downstream tasks (e.g., text/image/audio generation, and visual classification). Notable examples include PaLM-E, which combines vision and language for robotics; CLIP, which jointly learns visual and textual representations to enable various tasks; and Gato, a generalist agent capable of handling a wide range of inputs and tasks.

M3T FMs align naturally with the demands of XR environments by offering capabilities for (i) *enhancing the user’s perception of their environment* and (ii) *enabling the XR system’s dynamic response to the user’s needs*. Focusing on the former, M3T FMs can enrich user perception through real-time contextual overlays and semantic labeling. For instance, by processing visual and spatial inputs, M3T FMs can identify and annotate objects in the user’s field of view, translate foreign text on signage, or highlight important elements in complex environments (e.g., an antique Persian rug at a carpet exhibition). Focusing on the

**Table 2.** Common ML tasks enabled by XR device data, along with associated modalities, applications, and real-time processing constraints.

ML Task	Primary Modalities Used	Typical Use Cases	Real-Time Constraints
Hand Gesture Recognition	Visual, Motion	Interaction control, virtual object manipulation	Low latency (real-time interaction)
Eye Gaze Estimation	Eye Tracking	Adaptive rendering, user experience optimization	Very low latency (frame-by-frame adjustment)
Scene Understanding	Visual, Spatial Mapping	Semantic relabeling of 3D space, navigation	Moderate latency (background processing possible)
Speech Recognition	Audio	Voice commands, accessibility	Low latency (near real-time)
Activity Recognition	Motion, Visual	Fitness apps, workplace monitoring	Low to moderate latency
Object Detection	Visual	AR overlays, safety alerts	Low latency (interactive rendering)
User Intent Prediction	Eye Tracking, Motion, Audio	Anticipatory UI, task adaptation	Very low latency (proactive response)
Environment Mapping	Visual, Spatial Mapping	3D reconstruction, virtual overlays	Moderate latency (updated in intervals)
Emotion Recognition	Audio, Visual	Mental health monitoring, immersive storytelling	Moderate latency (buffered analysis acceptable)
Augmented Object Anchoring	Visual, Spatial Mapping	Persistent AR experiences, spatial memory	Low latency (to maintain spatial stability)

latter, by continuously interpreting multi-modal embodied signals (e.g., micro-expressions, posture, or speech patterns), M3T FMs can anticipate user needs and provide proactive support.

### ***Is the Seamless Integration of M3T FMs into the XR Ecosystem a Trivial Endeavor?***

Despite their promising applications, integration of M3T FMs into the XR ecosystem is far from trivial and certainly not without challenges. In particular, reaching the full potential of M3T FMs in XR often requires fine-tuning pre-trained M3T FMs (or training new ones), a process that is conventionally done at a *central server* with direct access to domain/user-specific data. Nevertheless, this domain-specific data is inherently collected in a *distributed* fashion at XR devices based on the real-time and continuous users' embodied interactions (e.g., gaze, voice, motion, and even physiological). Thus, to enable FM fine-tuning or adaptation, this high-dimensional data must be continuously transferred from XR devices to centralized cloud servers; a process that places a significant burden on wireless infrastructure and raises privacy concerns associated with data leaks.

Federated learning (FL) is a recent innovation in the area of distributed ML that offers a promising way to address the above concerns<sup>15</sup>. FL enables collaborative model training directly on user devices, eliminating the need to share their raw data with central servers. Integrating M3T FMs within the FL paradigm gives rise to a novel framework, called *M3T Federated Foundation Models (FedFMs)*, which enables collaboratively trained M3T FMs across various XR devices while alleviating the privacy concerns<sup>16,17</sup>. The decentralized training of M3T FedFMs across XR devices further unlocks the potential for *personalizing local M3T FMs to individual user and device characteristics*. Nevertheless, integrating M3T FedFMs into XR ecosystems would encounter multi-dimensional challenges, the identification of which is one of the goals of this paper.

### ***What Signals Point to the Feasibility of Integrating M3T FedFMs into the XR Ecosystem?***

We note that our vision for integrating M3T FedFMs into the XR ecosystem is motivated by the emerging body of work that demonstrates the successful implementation of their components (i.e., FL and M3T FedFMs). In particular, on one hand FL has gained traction in a variety of domains, such as mobile health<sup>18</sup>, Internet-of-Things (IoT)<sup>19</sup>, Industrial IoT (IIoT)<sup>20</sup>, vehicular networks<sup>21</sup>, and in non-terrestrial networks<sup>22</sup>. In XR settings, FL is particularly appealing: XR devices generate continuous streams of sensitive and personalized data (e.g., gaze trajectories, hand gestures, motion patterns), where FL can enable on-device model training while preserving user privacy. On the other hand, FMs and M3T FMs have emerged as transformative technologies both in XR and other domains<sup>23–25</sup>. Specifically, in XR, at the system level, early efforts from Meta Reality Labs<sup>26</sup> and NVIDIA Omniverse<sup>27</sup> demonstrate XR-focused M3T FM pipelines and simulation platforms. Yet all of these implementations still rely on conventional *centralized* M3T FMs. To further provide a better understanding of the uniqueness of XR settings, Table 3 provides an overview of the common limitations of the existing studies on M3T FM and FL when extended to M3T FedFM model training for XR applications.

It is worth mentioning that FedFMs and M3T FedFMs are becoming increasingly practical, as evident in other fields<sup>28–30</sup>, due to recent advances in resource-efficient local FM training/adaptation techniques. Further, in the AI/ML community, there is high interest in advancing parameter-efficient training/fine-tuning techniques, such as low rank adaptation (LoRA) and adapter modules, making the operations of M3T FedFMs over resource-constrained wireless devices increasingly practical<sup>31</sup>.

**Table 3.** Common limitations of FL and M3T FM approaches in XR contexts. Note that exceptions exist in each of the presented limitations, and the table rather presents the “common theme” of limitations.

Dimension	Traditional FL	Conventional M3T FMs	XR-Specific Challenges
Modality Diversity	Assumes uniform data modalities across clients; struggles with clients having different or missing modalities.	Pretrained on fixed modality combinations; lacks seamless adaptation to new or missing modalities without centralized fine-tuning.	XR devices capture diverse and dynamic modalities (e.g., vision, audio, haptics), which require handling modality heterogeneity and intermittent modality dropout during sessions.
Task Variability	Designed for homogeneous tasks; performance degrades with task heterogeneity among clients.	Fine-tuned for specific tasks; adapting to new tasks requires significant computational resources and data collection.	XR applications involve diverse and rapidly changing tasks (e.g., navigation, interaction, recognition), which demand adaptation to task variability across users and over time.
Real-Time Adaptation	Limited support for real-time model updates; communication overhead hinders low-latency requirements.	Large model sizes and inference times impede real-time responsiveness; not optimized for low-latency environments.	XR systems demand immediate adaptation to user interactions and environmental changes, which necessitate low-latency model updates.
Hardware Constraints	Often trains conventional ML models, where computation constraints may not be as restricting compared to FMs.	Require substantial memory, processing power, and energy supply; unsuitable for deployment on resource-limited devices.	XR devices often have limited computational capabilities, which makes it challenging to deploy and train/deploy large models efficiently.
Privacy Preservation	Provides data privacy by design; however, model updates can still leak sensitive information unless security countermeasures (e.g., differential privacy or encryption) are used.	Centralized training poses privacy risks.	XR applications collect sensitive user data (e.g., location), which requires robust privacy-preserving mechanisms during model training and inference.
Personalization	Personalization is limited and often requires additional mechanisms; struggling with balancing global and local model performance.	Fine-tuning for personalization is often not feasible due to not having access to the users’ private data; may not generalize well across users.	XR experiences are highly personalized, which calls for local models that can adapt to individual user preferences and behaviors in real-time.
Model Aggregation	Standard aggregation methods assume homogeneous models and data distributions; ineffective with heterogeneous client models and tasks.	Aggregation across different FMs is complex due to varying architectures and training objectives.	In XR, clients may have diverse models and tasks, which requires advanced aggregation techniques that can handle heterogeneity in both model structures and user objectives.
Communication Overhead	High communication costs due to frequent model updates.	Models are trained locally with no aggregations.	XR devices often generate high volumes of data and operate under battery constraints and limited bandwidth, which necessitates communication-efficient model update strategies.
Security and Robustness	Vulnerable to adversarial attacks and model poisoning; lacks robust mechanisms to ensure model integrity.	Often trained and kept in a central server equipped with firewalls	Distributed XR systems comprise unpredictable environments and users, which require models that are robust to adversarial conditions and can maintain integrity over time.

### The Vision and Goal of this Study

Despite their tremendous potentials and compelling motivations, the integration of M3T FedFMs within the XR ecosystem remains unexplored, primarily due to the emerging nature of M3T FedFMs themselves. This gap forms the central motivation for this study, in which we (1) identify application scenarios for M3T FedFMs in XR, (2) articulate the specific challenges facing the deployment of M3T FedFMs in XR environments, and (3) propose a performance evaluation framework for M3T FedFMs tailored to XR-specific needs, focusing on both the XR applications performance and the resource-efficiency of FedFM training/deployment. Collectively, our aim is to catalyze further research in this untapped space and lay the groundwork for realizing private, adaptive, and scalable intelligence within XR systems.

## M3T FedFMs: Model Architecture and Network Orchestration

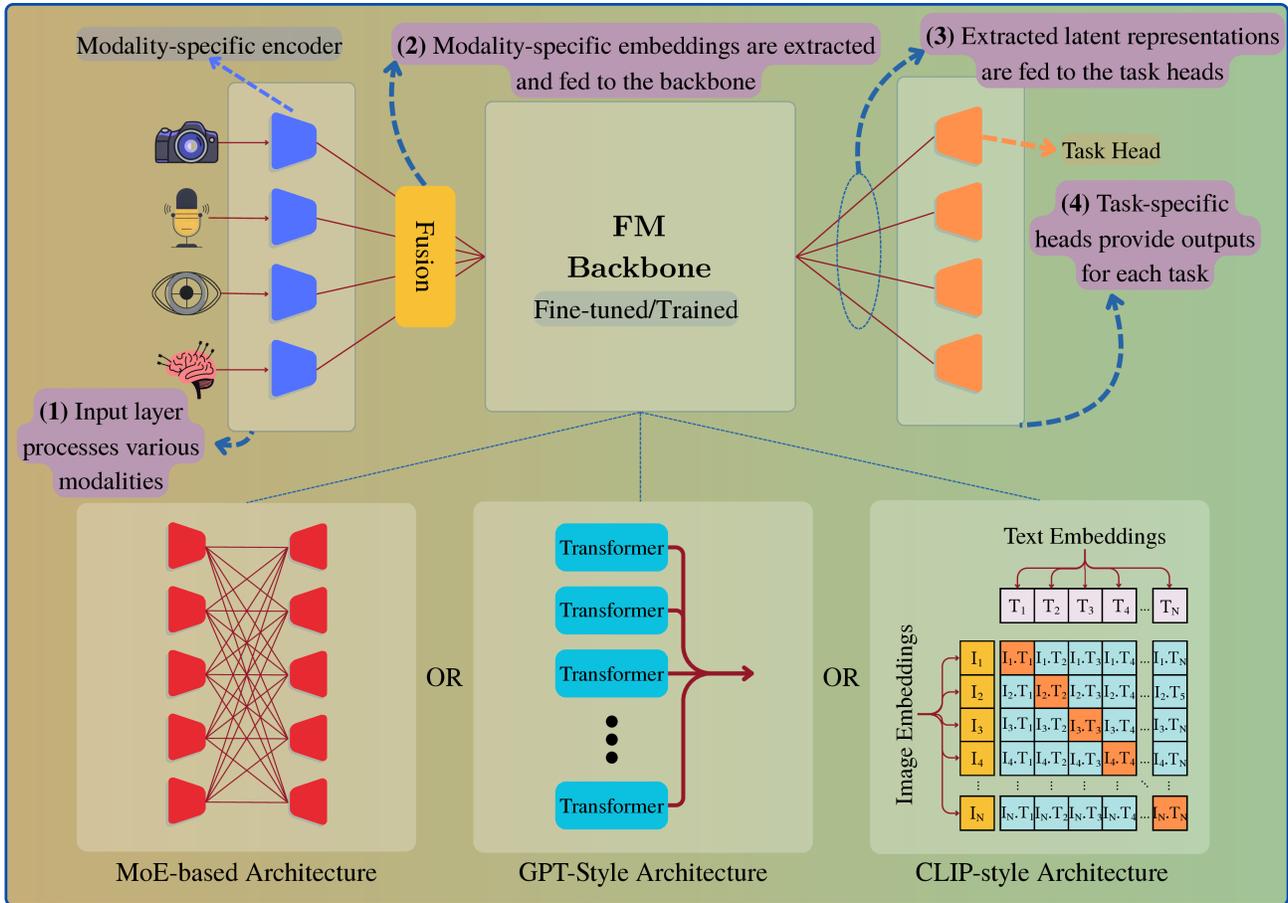
In its conventional form, FL conducts collaborative model training across devices by repeating two iterative operations until model convergence: (i) devices train their local models on their local datasets, and (ii) devices engage in transmitting their trained models (either to a server or among themselves) to aggregate their models and will use the aggregated model to continue their local training. Likewise, M3T FedFMs entail the orchestration of the training/fine-tuning of *local M3T FMs at the devices* through FL operations. Subsequently, understanding M3T FedFMs requires knowledge of M3T FMs and FL working mechanisms, which we provide next.

We note that M3T FMs are not defined by a single unified architecture but rather encompass a family of designs under active exploration in both academia and industry. In the following, we abstract a set of popular M3T FM architectures (depicted in Fig. 2) that are inspired by current trends in the research area<sup>32,33</sup>, followed by the model training procedure of FL.

### 1. Modality Encoders:

M3T FMs use dedicated *modality encoders* that convert raw sensory inputs into standardized intermediate embeddings. These encoders operate *prior to fusion*, serving as the entry point for the model backbone.

**2. Shared Backbone:** After the initial sensory encoding, the intermediate modality-specific embeddings are passed to



**Figure 2.** Foundation model architecture containing the input layer (i.e., modality encoders), the FM backbone, and the output layer (i.e., task heads). The FM backbone can comprise varying architectures, such as Mixture-of-Experts (MoE), GPT-style (stacked transformers), and CLIP-style (contrastive learning-based, where transformers are used for image and text processing) architectures.

a shared backbone, which performs multi-modal fusion, reasoning, and contextual learning. This backbone is central to generalization across tasks and users, and it can take several architectural forms depending on the model’s design goal. These architectural forms are depicted at the bottom of Fig. 2 and detailed below:

- *Transformer Backbones (e.g., GPT-style)*<sup>34</sup>: A stack of transformers process all input embeddings.
- *Dual-Encoder or Multi-Encoder Fusion (e.g., CLIP-style)*<sup>35</sup>: Each modality is processed by its own encoder (or backbone stream, which can be a transformer), and the outputs are aligned or fused in a shared latent space, often via contrastive learning or cross-attention.
- *Mixture-of-Experts (MoEs)*<sup>32</sup>: MoEs refer to a collection of parallel expert sub-networks, often transformer blocks or neural networks, selectively activated depending on the input modality or the targeted tasks. MoEs play two roles: (1) selecting the most relevant set of modality combinations for processing, and (2) routing the resulting fused representations through task-relevant expert pathways before reaching the output heads. Inside the backbone, a gating network determines which subset of experts to activate per input instance, allowing the model to maintain computational efficiency.

**3. Task Heads:** Task heads in the M3T FM architecture aim to map the shared representations, outputted from the backbone, to task-specific outputs.

**4. Fine-Tuning Capabilities (Adapters and Prompts):** To support efficient personalization in resource-constrained devices, M3T FMs often integrate lightweight adaptation mechanisms, namely, *adapters* and *prompt-based tuning*, defined below.

- *Adapters* refer to small, trainable modules inserted between or within layers of the shared backbone. These modules can be trained at each device to fine-tune the entire FM rather than updating all of its parameters. Adapters enable each device to fine-tune the model’s behavior locally while keeping the majority of the backbone frozen, thereby reducing computation and energy burden<sup>36</sup>.

**Table 4.** Pairwise comparison of M3T FM local training strategies. Each cell explains how the row method differs from the column method in terms of model behavior modification, modality/task adaptability, and hardware feasibility.

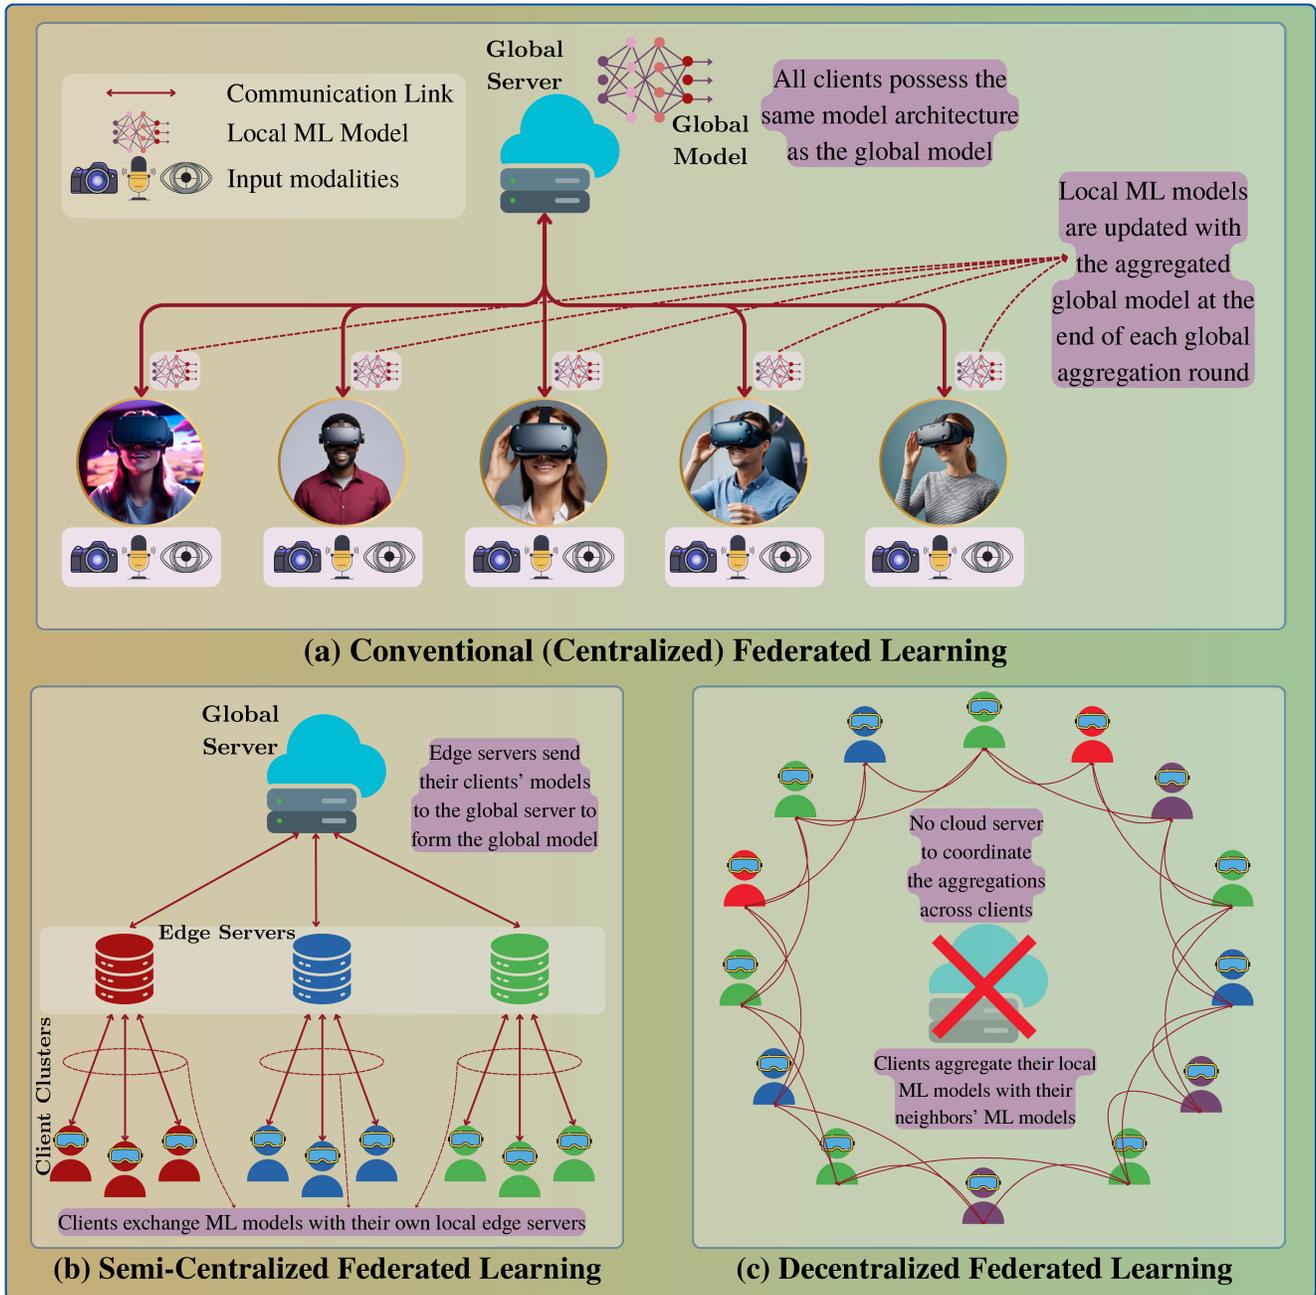
Compared To →	Prompt Tuning	Adapter Tuning	MoE Expert Tuning	Full Transformer / Backbone Tuning	LoRA Tuning
<b>Prompt Tuning</b>	—	Consumes fewer parameters and memory, but cannot significantly adapt to new modalities or tasks.	Lightweight but applies a fixed prompt across all inputs, limiting adaptation to diverse modalities.	Can run on resource-constrained devices, but lacks depth and task adaptation capacity provided by full fine-tuning.	Less expressive than LoRA, which introduces learnable rank-specific updates that better balance adaptability and footprint.
<b>Adapter Tuning</b>	Can better tailor responses to new tasks or sensors than prompt tuning, but needs slightly more compute/memory.	—	Applies uniformly across input, whereas MoE can adapt per sample with selective expert routing.	Offers reasonable task coverage with less resource overhead compared to tuning all transformer weights.	LoRA is lighter than adapters in compute and memory (and inference) but slightly less modular for some downstream tasks.
<b>MoE Expert Tuning</b>	Dynamically activates specialized sub-modules per input type, boosting adaptability at the cost of routing complexity and memory.	More modular and task-adaptive in theory, but needs well-designed expert architecture to be efficient.	—	Often achieves similar performance to full model tuning if expert selection is well optimized, but with significantly lower resource costs.	LoRA is more memory-efficient but lacks the task-specific routing of MoE; MoE suits structured modality switching better.
<b>Full Transformer / Backbone Tuning</b>	Deeply retrains all parameters for strong adaptation, but impractical for edge device constraints.	Too resource-heavy for real-time or on-device use; adapter tuning is more suitable for low inference-time adaptation.	Consumes more memory/compute than MoE without necessarily improving adaptability in structured tasks.	—	LoRA achieves comparable performance in many cases with drastically fewer trainable parameters and better deployability.
<b>LoRA Tuning</b>	Provides rank-based updates that adapt better than fixed prompts while retaining low overhead.	Often more efficient in training/inference compared to adapters, but slightly less flexible for multi-modal fusion.	Lacks dynamic expert routing but excels in memory- and compute-limited settings.	Can mimic some full-tuning benefits with reduced cost, making it viable for on-device learning.	—

- *Prompts* are learnable input tokens, obtained via *prompt-tuning*, prepended to the sequence of modality embeddings or inserted at key positions within the network. Rather than modifying internal weights, prompts steer model behavior to personalize the local model to the device, making them ideal for low-latency on-device updates<sup>37</sup>.

**5. FL Coordination Paradigm:** Following FL principle, in M3T FedFM, each device performs local training over its dataset to update its local M3T FM modules/parameters; however, *this process usually does not entail the training of the entire M3T FM*, as the burden of such training exceeds the capabilities of resource-constrained devices. Instead, the local training is often performed in the form of prompt-tuning, learning the adapter parameters (e.g., via LoRA-based methods), or updating relevant M3T FM components, such as local experts. In Table 4, we provide a pairwise comparison between the techniques that can be used during the local training, identifying their unique advantages and disadvantages.

Followed by local training, the devices can use three various FL coordination paradigms (depicted in Fig. 3), which are detailed below and compared against each other in Table 5, to improve their local model generalizability and performance:

- **Centralized Aggregation:** A cloud server collects model updates from all participating devices. It then aggregates these updates to form a global model, typically via weighted averaging, and redistributes/broadcasts the global model back to the devices. This style of aggregation, although it offers good scalability, engages devices in energy-intensive uplink communications, making it less ideal for implementation over resource-limited systems<sup>15</sup>.
- **Fully Decentralized Aggregation:** In the absence of any central server, devices collaborate purely through energy-efficient, short-range device-to-device (D2D) model exchanges, the aggregation is often done through distributed consensus mechanisms. While this aggregation method provides maximum autonomy and energy-efficiency, it introduces challenges in convergence and is often not suitable for large-scale networks as it requires a connected network substrate formed by the D2D links<sup>40</sup>.
- **Semi-Decentralized Aggregation:** Devices are organized into clusters based on spatial proximity, connectivity, or task similarity. Within each cluster, devices engage in energy-efficient D2D communications to reach a local model consensus. Afterwards, one device, often selected dynamically, acts as a cluster representative and performs an uplink to the central server. The server then forms the global model through the received updates, as in the centralized aggregation, and broadcasts the global model back to the devices. This approach reduces uplink communication overhead and can promote local personalization of models at the cluster level; however, it is only deployable when devices are capable of D2D communications<sup>41</sup>.



**Figure 3.** Various model aggregation and network coordination strategies in FL: **(a) Conventional FL** employs a centralized model aggregation strategy where clients perform local training and send their model updates to a central server. The server aggregates these updates into a global model and broadcasts it back to the clients. **(b) Semi-decentralized FL** introduces a hierarchical structure where clients are organized into clusters. Within each cluster, clients exchange models via D2D communications and reach local consensus<sup>38,39</sup>. A selected client from each cluster then sends the aggregated model to the central server, which constructs a global model and distributes it across all clients. **(c) Fully decentralized FL** operates without a central server; model updates are exchanged entirely over D2D links, where each client updates its local model based on models received from its neighbors (e.g., using distributed consensus algorithms).

### M3T FedFMs Use Cases in the XR Ecosystem

In the following, we highlight a range of scenarios where M3T FMs have the potential to redefine how XR systems perceive, adapt, and interact, spanning domains such as healthcare, education, manufacturing, and entertainment. These scenarios, schematically visualized in Fig. 4, represent a blend of evolved versions of existing XR deployments and forward-looking use cases that push the boundaries of what is possible. Notably, with the recent wave of on-device FM integration into smartphones, particularly LLMs<sup>42</sup>, it is increasingly clear that no major player in the XR space, including tech giants such as Apple, Google, and Meta can afford to ignore the paradigm shift that the recently emerged M3T FMs, and by extension, M3T FedFMs, are poised to bring to this field.

**Table 5.** Comparison of data/model aggregation mechanisms for M3T FedFMs in XR ecosystems.

Strategy	Definition	Benefits	Drawbacks
Centralized Training	All training data is collected and processed on a central server.	Simplifies coordination; supports large-scale pretraining; ensures global consistency.	Raises privacy concerns; assumes constant connectivity; not ideal for on-device personalization.
Centralized FL	Devices train locally and send model updates to a central server for aggregation.	Supports privacy; adapts to local data distributions; scalable.	Relies on server availability; Stragglers (e.g., low-compute users) can affect the model convergence.
Fully Decentralized FL	Devices average their models via D2D consensus without relying on any central server.	Resilient against server failures/downtimes; useful in ad-hoc or low-connectivity networks.	Slow convergence; harder to maintain model consistency across devices.
Semi-Decentralized FL	Devices form clusters to reach local D2D consensus on their models before one representative communicates with the server.	Reduces uplink bandwidth; supports localized adaptation; balances load across the network.	Adds coordination complexity within clusters.

We emphasize that in presenting the scenarios, our focus is on justifying the *use of well-trained/fine-tuned* M3T FMs within each context. The discussion on *how these models are obtained in the first place* follows later in the section, anchored by a key observation: M3T FMs intended for deployment on XR devices cannot be feasibly trained via conventional centralized pipelines due to constraints in privacy, bandwidth, and personalization. Instead, these models must be trained through FL approaches, highlighting the essential role of M3T FedFMs across all the scenarios presented.

**(Scenario 1) Cognitive XR Augmentation: Real-Time Perceptual Overlays**

Cognitive XR augmentation entails using the multi-modal embodied data to generate useful, real-time, perceptual overlays for various downstream tasks<sup>43</sup>. In these settings, the multi-modal embodied data includes visual (e.g., scene context and handwriting), audio (e.g., spoken commands), and user-specific cues (e.g., emotional tone). Also, the typical downstream tasks are live translation, contextual labeling of objects, and environment-aware recommendations.

In such environments, M3T FMs are particularly potent due to their ability to handle various modalities and adapt to multiple concurrent tasks. For instance, a user might look at a handwritten draft and instantly receive suggestions on clarity, grammar, or tone, all rendered directly within their visual field through a well-trained M3T FM<sup>44</sup>. Similarly, in industrial settings, a technician wearing an XR headset may receive dynamic overlays pointing out abnormal machine vibrations, heat signatures, or wear patterns, all interpreted and delivered by an M3T FM that feeds the overlays to the headset<sup>45</sup>.

**(Scenario 2) Neuroadaptive Interfaces: XR with Brain-Computer Integration**

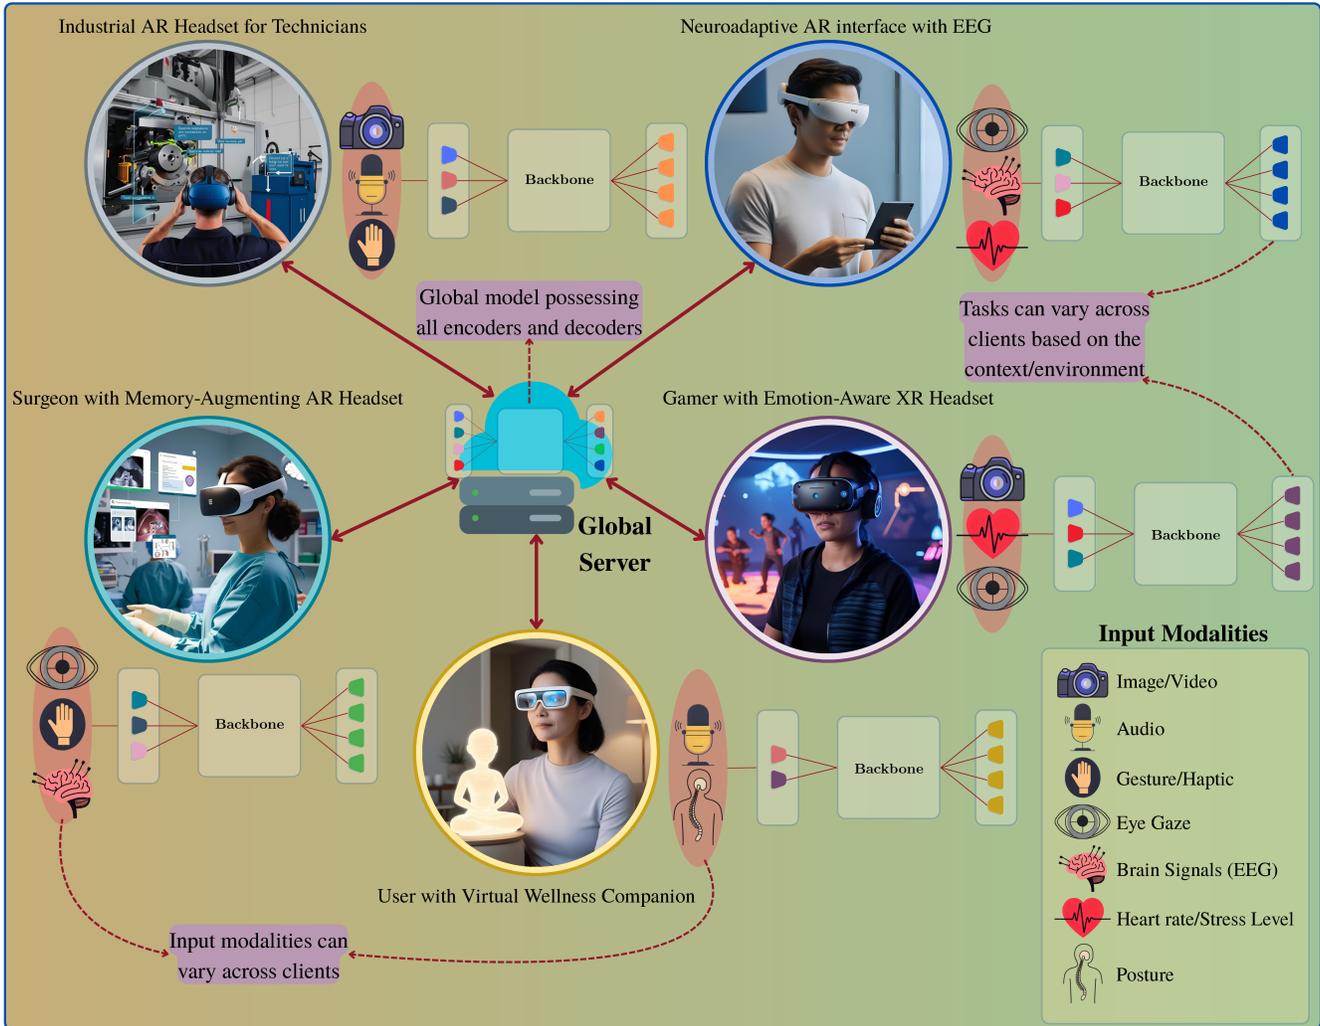
Neuroadaptive XR interfaces represent the next frontier in embodied interaction, where XR systems respond to both the external behavior of users and their internal cognitive and emotional states. These interfaces have come closer to reality owing to the advances in neural sensing technologies, including wearable electroencephalogram (EEG) devices, functional near-infrared spectroscopy (fNIRS) sensors, and implantable brain-computer interfaces (BCIs) such as Neuralink<sup>46</sup>, which collectively make the integration of brain activity as a live input stream within XR environments a feasibility.

In such settings, M3T FMs can enable truly *intention-driven experiences*, an emerging paradigm toward which both industry and academia are increasingly converging<sup>47</sup>. For instance, a user might simply *think* about a command on opening an application, and the interface seamlessly responds without requiring any spoken or gestural command. In the medical domain, this capability holds transformative potential<sup>48</sup>. For example, a patient with motor impairments might imagine moving a paralyzed limb. Subsequently, the M3T FM, trained to process multi-modal embodied inputs (e.g., neural signals, visual context, and spatial mapping), can interpret this mental intent and translate it into the movement of a virtual limb within the XR environment<sup>49</sup>.

**(Scenario 3) Embodied AI Companions: Adaptive Agents**

Another future application of XR systems, actively pursued in both academia and industry, is the creation of embodied AI companions<sup>50</sup>. These AI companions are intelligent agents that evolve alongside the user. These agents can manifest as visual avatars or operate invisibly in the background, providing contextual support such as reminders, co-navigation, or behavioral nudges<sup>51</sup>. The operation of these systems is grounded in the continuous acquisition of multi-modal embodied data streams, including prosodic/rhythmic features of speech (e.g., pitch, cadence) and ocular patterns. Also, the computational tasks associated with such systems encompass emotional state classification, fine-grained intent recognition, longitudinal behavior modeling, and real-time context-aware recommendation generation<sup>52</sup>.

In these systems, M3T FMs can lead to unprecedented advancements. For example, a companion powered by an M3T FM may detect a gradual decline in the user’s emotions, moods, or feelings (e.g., based on combined analysis of speech, gaze, and posture), and preemptively offer mental health resources. Also, by leveraging contextual cues such as eye-gaze and user behavior, AI companions can provide suggestions in terms of overlays to improve the user’s productivity<sup>53</sup>.



**Figure 4.** FedFM deployment visualization in an XR system, under FL with centralized model aggregation strategy. Different deployment environments of XR devices will result in varying tasks and modalities across the clients.

**(Scenario 4) Emotionally Resonant Storytelling in XR Gaming**

With modern XR headsets equipped with sensors for gaze tracking, pupil dilation, facial micro-expressions, and vocal stress detection, *emotion-aware gaming* can transition from a futuristic concept to an emerging reality, the promise of which has triggered a major interest in the gaming industry<sup>54</sup>. XR devices can continuously collect multi-modal embodied signals that reflect the player’s internal/emotional state in real-time. Building on this capability, adaptive storytelling is emerging as a new frontier in emotion-aware gaming, where narratives evolve not only in response to a player’s actions, but also in response to their physiological states<sup>55,56</sup>. This enhances the emotional depth that XR games can offer, fostering a more profound and personalized connection between players and virtual environments. Notably, significant efforts have been directed toward dynamically eliciting/triggering complex emotions such as fear, creating deeply immersive and psychologically resonant gaming experiences<sup>57,58</sup>.

In this context, M3T FMs’ ability to integrate diverse sensory inputs and simultaneously handle multiple downstream tasks makes them ideal for dynamically adjusting the pace, intensity, and structure of game narratives. Specifically, a well-trained M3T FM can recognize when a player is deeply immersed, emotionally overwhelmed, or disengaged, and adapt the storytelling rhythm accordingly.

## Why in the Above (and All Comparable) Scenarios M3T FMs Must be Accompanied by FL?

A common thread across all the scenarios discussed above is the geo-distributed nature of data generation. XR systems deployed in homes, hospitals, factories, classrooms, and entertainment spaces continuously collect embodied, multi-modal data, which is generated/collected locally on XR devices. Such a geo-distributed nature of data generation poses a major challenge to the integration of M3T FMs into XR scenarios: traditional *centralized* M3T FM training paradigms require transferring/pooling data from all XR devices into a centralized server to train a globally shared model. Yet, in the context of XR, the centralization of XR data is impeded by three key limitations:

- **Privacy:** XR data inherently includes deeply personal and behavioral information (e.g., emotional states, cognitive intentions, and biometric patterns) that are not suitable for off-device transmission or cloud storage.
  - **Bandwidth and Latency:** Continuously uploading high-dimensional sensory data from distributed XR devices imposes unsustainable demands on wireless infrastructure, introducing unacceptable delays in real-time experiences.
  - **Personalization:** Centralized training typically produces generic models that fail to reflect local variations in user behavior and context – critical for applications such as neuroadaptive interfaces or emotion-aware storytelling.
- ★ Leveraging the distributed model training capabilities of FL, **M3T FedFMs** offer a compelling alternative that can address these challenges, making them a major component of next-generation intelligent XR systems.

★ In the following, as our focus is solely on ‘M3T FedFMs’, we use a shorter abbreviation ‘FedFMs’ to refer to them.

## SHIFT Dimensions and their Manifestation in XR Ecosystems

The implementation of FedFMs in the XR domain, despite their high potential, faces a variety of nuanced challenges that must be carefully characterized in order to guide future research, system design, and deployment strategies. Motivated by the need to isolate and understand the most relevant obstacles, we categorize the core challenges associated with training, adapting, and personalizing FedFMs in XR systems under **SHIFT** dimensions: (1) **S**ensor and modality diversity, (2) **H**ardware heterogeneity and system-level constraints, (3) **I**nteractivity and embodied personalization, (4) **F**unctional/task variability, and (5) **T**emporality and environmental variability. It is worth mentioning that these challenges/dimensions are complementary to those of conventional FL (e.g., data heterogeneity<sup>59</sup>).

Below, we elaborate on each of **SHIFT** dimensions, highlighting how it imposes constraints that can disrupt standard FedFM learning assumptions and necessitates a rethinking of how FedFMs are designed, trained, and deployed in XR environments.

### **(Dimension 1) Sensor and Modality Diversity**

XR devices may *differ vastly* in the types and combinations of sensors they possess, ranging from lightweight AR glasses with vision-only input to immersive VR headsets equipped with eye-tracking, haptic feedback, spatial audio, and depth sensing. This sensor heterogeneity leads to non-overlapping feature spaces across devices. Moreover, modalities *may dynamically appear or disappear* within an XR session (e.g., a muted microphone, disabled haptics, or environmental occlusion). Addressing these characteristics in conjunction requires enabling XR devices with *incomplete modality views* and *temporal modality dropouts and emergence* to still contribute meaningfully to FedFMs, which are not conventionally designed to handle such characteristics.

### **(Dimension 2) Hardware Heterogeneity and System-Level Constraints**

XR systems span a wide range of hardware from AR glasses with minimal compute and memory to high-end MR headsets equipped with advanced computing units. This hardware heterogeneity leads to *asymmetric participation in model training*, where not all XR devices can host full models or conduct local model updates frequently. While conventional FL methods, and subsequently FedFMs, support asynchronous and partial device participation, the strict latency and reliability requirements of XR, such as sub-50ms response<sup>1</sup> for gesture-based interaction, amplify the impact of hardware heterogeneity.

### **(Dimension 3) Interactivity and Embodied Personalization**

XR environments are not passive; instead, they are inherently interactive, responding to user’s motion, attention, speech, and preference in real-time. This interactivity, combined with the embodied nature of XR (i.e., user movement, posture, and environment are part of the interaction loop), introduces new challenges for FedFM personalization: FedFMs must adapt continuously in XR devices, rather than post-hoc. Nevertheless, conventional FedFM personalization techniques, such as local fine-tuning, personalized output heads or adapters, are often resource-demanding or not quick/agile enough for XR scenarios, where real-time feedback and low-latency inference are critical.

<sup>1</sup>Such latency requirements align with AR/VR studies<sup>60</sup>, which show that delays over 50ms noticeably degrade system responsiveness and user comfort.

**Table 6.** Evaluation metrics for FedFMs in XR ecosystems. These metrics reflect both generic ML objectives and XR-specific requirements and are categorized by the relevant SHIFT dimensions.

Metric	Description	How to Measure	SHIFT Dimensions
Inference Latency	Time between the reception of sensing inputs and model output generation; critical for responsiveness.	Log time between input event (e.g., gesture, gaze) and output generation.	H - I - T
Personalization Effectiveness	Performance gain after local personalized (e.g., using prompt tuning or local training) vs. generic global model.	Compare task accuracy or user satisfaction rate before and after personalization.	I - F
Personalization Necessity Rate	Proportion of users or tasks where global FedFM meets performance threshold, avoiding need for adaptation.	Count users/tasks for which the global model exceeds a baseline accuracy (e.g., 90%) without fine-tuning.	H - I - F
Cross-Context Generalization	The model's ability to generalize effectively across previously unseen users, environments, or tasks.	Test model on data from new users or contexts not seen during training.	S - F - T
Task-Specific Continuity	Retaining performance on previously learned tasks during adaptation on the new tasks.	Evaluate drop in earlier task accuracy post new-task fine-tuning (i.e., model forgetting measure under continual learning).	F - T
Submodule Activation Efficiency	Average number of active FedFM modules per task; reflects model modularity.	Count number of modules activated per task or per user interaction.	H - F
Model Storage Footprint	Total model size deployed on XR device, including adapters or quantized formats.	Sum memory usage in MB for all active modules on device.	H
Communication Overhead	Amount of data transferred per aggregation round or XR session.	Monitor the bits exchanged during global aggregation rounds or XR sessions.	H - T
Battery Impact / Energy Use	Device power consumption while running/adapting FedFM.	Measure power draw in mW or energy used (Joules) via onboard or external sensors.	H - T
User-Reported QoE	Subjective user-perceived quality of experience, focusing on responsiveness and interaction fluidity.	Collect Likert-scale ratings, qualitative feedback, or lag reports through user surveys.	I - T
Model-Inferred QoE	System-side estimation of perceived interaction quality without explicit user input.	Monitor implicit behavioral signals such as gaze jitter, task repetition, latency-induced hesitation, or input aborts.	I - T
On-Device Training Time	Duration required for local model adaptation, tuning, or fine-tuning directly on the XR device.	Record duration of local fine-tuning routine on device per adaptation event.	H - I - T
Privacy Leakage Risk	Potential for data leakage through observation of model updates or representations during global aggregations.	Run membership inference attacks, gradient inversion, or compute differential privacy bounds.	S

#### (Dimension 4) Functional/Task Variability

In a single session/time-window of XR interactions, the user's tasks may vary from identifying objects and following spatial navigation cues, to engaging in conversations with a virtual assistant and responding to real-time gesture-based commands. These tasks differ not only in output format but also in temporal scale (e.g., real-time vs. near-real-time), relevant modality, and computational demand. This intra-user multi-tasking in XR is further compounded by *inter-user and inter-application variability*: a user in a manufacturing setting may need real-time error detection and instructions, while a student in an educational XR system might engage in immersive simulations and voice-based tutoring. This breadth of functionality, which we refer to as *functional variability*, breaks the common assumption in conventional FedFMs that devices often share a fixed set of downstream tasks or training objectives.

#### (Dimension 5) Temporality and Environmental Variability

XR systems operate in environments that may change rapidly over time (e.g., users move through spaces with shifting lighting). Therefore, federated adaptations in FedFMs for XR must account for such rapid temporal variability while *not* forgetting the previously learned knowledge, resembling the *catastrophic forgetting* issue in conventional deep learning models<sup>61</sup>. Nevertheless, addressing this challenge is underexplored in the literature on FedFMs.

Addressing the aforementioned **SHIFT** dimensions necessitates the development of novel techniques across training, fine-tuning, model storage/compression, and model communication within FedFM workflows. In the following, we discuss how such future techniques should be evaluated in terms of their performance. While this work does not aim to prescribe detailed technical research directions, we will later outline a set of high-level solution strategies that can serve as a basis for future exploration.

## Realization of Resource-Aware FedFMs for XR: Metrics, Tradeoffs, and Datasets

Unlocking the potential of FedFMs in XR environments requires more than the development of novel algorithms; it demands a robust and interoperable ecosystem that reflects the unique challenges posed by the **SHIFT** dimensions. We offer our perspective on the various aspects of this ecosystem below.

**Table 7.** Key design tradeoffs in deploying FedFMs in XR ecosystems. Each tradeoff highlights competing objectives that must be balanced to meet XR-specific constraints.

Design Tradeoff	Options	Why It Matters in XR
Local Personalization vs. Global Generalization	<b>Option 1:</b> Fine-tune locally for user-specific behavior <b>Option 2:</b> Rely on a shared global FedFM	Personalization improves responsiveness and user alignment but increases on-device compute and energy consumption that may penalize the XR application. Global models without personalization may lack specificity for XR contexts.
Update Frequency vs. Communication Overhead	<b>Option 1:</b> Frequent federated updates for real-time adaptation <b>Option 2:</b> Sparse, event-triggered updates to save bandwidth	Frequent updates accommodate behavioral drift in the global model but increase on-device energy and bandwidth use that may penalize the XR application that runs on it. Sparse updates conserve resources but risk model obsolescence, degrading user experience over time.
Centralized Aggregation vs. Decentralized D2D Aggregation	<b>Option 1:</b> Aggregation via a central server (e.g., cloud-based aggregator) <b>Option 2:</b> D2D aggregation directly among nearby XR devices	Centralized aggregation ensures global model consistency and convergence, yet, it is dependent on cloud connectivity and may introduce latency. D2D aggregation allows opportunistic learning, but risks slow model convergence.
User-Driven QoE Optimization vs. System-Driven QoE Optimization	<b>Option 1:</b> Collect explicit feedback from users <b>Option 2:</b> Infer experience quality from behavioral signals	Explicit feedback is precise but intrusive; implicit inference is unobtrusive but error-prone.
Modular Model Deployment vs. Monolithic Model Deployment	<b>Option 1:</b> Use lightweight, shallow adapters/prompt-tuners for fast tuning <b>Option 2:</b> Fine-tune deeper layers of the local FM for higher fidelity	Shallow tuning conserves energy, while deep tuning (especially on rich data) enhances personalization at the cost of increased power and thermal stress.
Privacy Preservation vs. Model Performance	<b>Option 1:</b> Enforce strong privacy using differential privacy or encryption <b>Option 2:</b> Allow model updates and aggregation with no privacy protection	Using privacy-protection methods guards XR users' confidentiality but may limit model quality and add resource overhead (e.g., when encryption methods are applied). Looser policies improve learning but raise concerns about data leakage.

**FedFM Performance Evaluation in XR Ecosystems:** To support the design and benchmarking of FedFMs in XR, we envision a targeted set of evaluation metrics that reflect the unique demands of XR systems, which are further aligned with the **SHIFT** dimensions. In Table 6, we provide a structured presentation of these metrics that can serve as a guide for developers and researchers to assess FedFM behavior across diverse XR settings.

**Major FedFM Performance Tradeoffs in XR Ecosystems:** Designing FedFMs for XR environments involves navigating a set of *competing objectives/trade-offs* that emerge from the unique constraints of embodied interaction, mobile hardware, and real-time responsiveness. Subsequently, in Table 7, we outline the key tradeoffs that researchers can consider when deploying FedFMs over XR devices. It is worth mentioning that recent studies have started to explore such fundamental tradeoffs between computation, communication, and quality of experience for ML<sup>62,63</sup> and other tasks<sup>64,65</sup> in XR systems; however, extending them to FedFM settings remains an open direction.

**Datasets for FedFMs in XR:** A core challenge in advancing FedFMs for XR is the lack of standardized datasets that reflect the embodied, multi-modal, and distributed nature of real-world XR systems. Table 8 summarizes both existing and our perspective on the needed/missing datasets, with a focus on their modality coverage, task diversity, and high-level description. Our goal in presenting this table is twofold: (1) to identify available XR datasets that can be used for FedFM training and evaluation, and (2) to outline datasets that do not exist but are essential for benchmarking FedFMs under realistic XR conditions, helping future data collection efforts.

## Towards Resource-Aware and High-Performance FedFMs for XR under SHIFT Dimensions

Effectively addressing the **SHIFT** dimensions requires a concrete multi-faceted research agenda tailored to the operational realities of XR environments, investigation of which deserves a separate study and is outside the scope of this perspective paper. Nevertheless, to stimulate further investigations in this untapped area we provide a set of *high-level research directions* below.

**(Solution 1) Addressing Sensor and Modality Diversity:** To address incomplete or intermittent modality signals, future works can explore the following directions:

- **Robust Multimodal Fusion under Missing Inputs:** Develop FedFM models that can operate effectively with partial, missing, or dynamically changing sensory streams by leveraging redundancy across modalities.
- **Asymmetric Input Adaptation:** Design FedFM training strategies where devices with fewer modalities (e.g., vision-only AR glasses) can improve the quality of their local FMs through knowledge sharing with rich-modality devices (e.g., VR headsets with gaze and audio) through techniques such as cross-modal distillation or embedding alignment.

**(Solution 2) Addressing Hardware Heterogeneity and System-Level Constraints:** To address the broad range of compute, battery, and thermal limitations across XR devices, future works can explore the following directions:

- **Resource-Aware Inference Paths:** Design model components (e.g., MoEs) that dynamically adjust execution based on real-time resource states (e.g., battery, compute).

**Table 8.** Extended list of XR-relevant datasets for training and evaluating FedFMs. This table includes both real-world and needed datasets across multiple modalities and XR-specific challenges.

Dataset	Type	Modalities	Primary Tasks	Description
xR-EgoPose <sup>66</sup>	Existing	Egocentric RGB	3D human pose estimation	Head-mounted fisheye video dataset enabling full-body 3D pose recovery for XR applications.
HOI4D <sup>67</sup>	Existing	RGB-D, 3D hand/object pose	Human-object interaction, affordance learning	Large-scale egocentric dataset with 2.4M frames covering daily indoor tasks and object manipulation.
CWIPC-SXR <sup>68</sup>	Existing	RGB-D, point clouds, audio	Social interaction modeling, real-time rendering	Dynamic point cloud dataset capturing 45 sequences of human interactions in social XR scenarios; includes RGB-D, audio, etc.
HEADSET <sup>69</sup>	Existing	RGB-D, point cloud, light field	Emotion recognition under occlusion	Volumetric capture of emotional facial expressions with head-mounted gear and occlusion scenarios.
AMIS <sup>70</sup>	Existing	Audio, video	Multimodal communication modeling	Open-source audiovisual dataset for analyzing immersive communication and social XR interaction.
Federated XR Simulators	Needed	Multimodal (vision, audio, haptics, speech, motion)	Multimodal fusion, personalized policy learning	Simulated XR environments (e.g., Unity/Omniverse) with structured profiles for device-user-task modeling in FL settings.
Federated Role-Specific Benchmarks	Needed	Multi-modal with user metadata	Role-conditioned intent prediction, task alignment	Pre-partitioned benchmarks reflecting user roles, hardware diversity, and task heterogeneity.
Longitudinal XR Interaction Logs	Needed	Vision, gaze, audio, task feedback, physiological signals	Continual learning, drift detection	Tracks evolving user behavior across time for learning drift, retention, and personalization analysis.
Social and Collaborative XR Benchmarks	Needed	Gaze, speech, multi-user context	Multi-agent coordination, shared task grounding	Captures collaborative and role-aware interactions in XR environments to train FedFM-based assistants.
Cognitive Load-Aware Datasets	Needed	Gaze, micro-pauses, EEG/fNIRS (optional)	Cognitive load estimation, pacing adjustment	Supports user state estimation via behavioral and physiological patterns for adaptive interface behavior.

- **Collaborative Cross-Device Training:** Enable extremely resource-limited devices (e.g., smart glasses) to securely offload their data to more capable co-owned devices (e.g., XR headsets, smartphones, or edge hubs), where the amount of shared data and device assignments are optimizable variables. This paradigm leverages user-owned compute availability to maintain personalization while satisfying device constraints.

*(Solution 3) Addressing Interactivity and Embodied Personalization:* To address the need for fast user-specific model adaptations in dynamic XR settings, future works can explore the following directions:

- **Composable User-Specific Modules:** Optimize the selection and use of the plug-and-play components (e.g., adapters or task heads) that specialize rapidly to a user’s context without modifying the global model backbone.
- **Behavior-Conditioned Personalization Signatures:** Develop low-overhead user “signatures” (e.g., in the form of contextual prompts) derived from embodied interaction traces (e.g., motion rhythm, gaze dynamics, voice cadence) that are fed to the model and condition its responses without requiring full parameter updates, enabling personalization that is both efficient and privacy-preserving.

*(Solution 4) Addressing Functional and Task Variability:* To address the wide spectrum of tasks present in XR applications, future works can explore the following directions:

- **Task-Isolated Update Paths:** Implement task-specific modules (e.g., separate heads, adapter stacks, or routing gates) that can carefully isolate gradients during training, ensuring that updates from one task do not unintentionally degrade performance on others.
- **Dynamic Task Module Selection:** Develop strategies for deciding whether to support a new task by extending an existing module, such as stacking an adapter on top of a previously trained task head, or by creating a new dedicated task head within the FedFM. These decisions should consider factors, including representational similarity between tasks, available device resources (e.g., memory, compute), and empirical indicators such as model convergence speed.

*(Solution 5) Addressing Temporality and Environmental Variability:* To address the environmental changes and shifts in user behavior and physical surroundings, future works can explore the following directions:

- **Temporal Context Signaling Modules:** Introduce lightweight modules that explicitly encode temporal drift signals, such as a shift in gaze patterns, degraded gesture precision, or increased latency in user feedback, and feed them into the FedFM. These modules can act as control gates that adjust model behavior based on the real-time collected signals from the user and environment, while preventing the model from forgetting its past learned knowledge.

- **Continual Contextual Adaptation:** Apply continual learning algorithms on various FedFM modules (e.g., encoder, task heads, MoEs) that balance quick adaptation to new scenarios, while retaining prior context for long-term stability.

It is worth mentioning that following the development of each of the solutions, a comprehensive evaluation is essential: one that rigorously assesses performance using the previously defined metrics and design trade-offs.

## Conclusion

In this perspective paper, we highlighted the motivations for utilizing M3T FedFMs in XR ecosystems and the unique challenges of this vision framed by the **SHIFT** dimensions. We also provided a framework for performance evaluation to enable the resource-effective integration of M3T FedFMs into XR ecosystems. We finally envisioned a set of research directions, aiming to motivate further investigations in this underexplored area. We believe that challenges ahead are substantial, but so too are the potentials and benefits: intelligent XR systems that respect privacy, adapt seamlessly to users and environments, and generalize across a wide spectrum of devices, applications, and modalities.

## Author contributions

F. Nadimi and S. Hosseinalipour conceptualized the paper, prepared the first draft of the paper, and refined the paper. P. Abdisarabshali and K. Borazjani contributed to the ideas, editing the text, and drawing the diagrams. J. Chakareski contributed to the ideas and editing the text. The final version of the paper was critically reviewed and approved by all authors.

## Acknowledgements

Not applicable

## Funding

S. Hosseinalipour was partially funded by NSF 2421761 while working on this paper.

## Competing interests

The authors declare no competing interests.

## Additional information

Correspondence and requests for materials should be addressed to S. Hosseinalipour.

## References

1. Rauschnabel, P. A., Felix, R., Hinsch, C., Shahab, H. & Alt, F. What is XR? Towards a framework for augmented and virtual reality. *Comput. human behavior* **133**, 107289 (2022).
2. Merel, T. Virtual, augmented and mixed reality are the 4<sup>th</sup> wave. The Wrap News Inc. (2016).
3. Fortune Business Insights. Virtual and augmented reality market size research reports. [Online] (2023). <https://www.fortunebusinessinsights.com/industry-reports/virtual-reality-market-101378>, <https://www.fortunebusinessinsights.com/augmented-reality-ar-market-102553>.
4. Kourtesis, P. A comprehensive review of multimodal XR applications, risks, and ethical challenges in the metaverse. *Multimodal Technol. Interact.* **8**, 98 (2024).
5. Ziker, C., Truman, B. & Dodds, H. Cross reality (xr): Challenges and opportunities across the spectrum. *Innov. learning environments STEM higher education: Oppor. challenges, looking forward* 55–77 (2021).
6. Chakareski, J. & Khan, M. Live 360° video streaming to heterogeneous clients in 5G networks. *IEEE Trans. Multimed.* **26**, 8860–8873 (2024).
7. Chakareski, J., Khan, M., Ropitault, T. & Blandino, S. Millimeter wave and free-space-optics for future dual-connectivity 6DOF mobile multi-user VR streaming. *ACM Trans. Multimed. Comput. Commun. Appl.* **19**, 57:1–25 (2023).
8. Gupta, S., Chakareski, J. & Popovski, P. mmWave networking and edge computing for scalable 360-degree video multi-user virtual reality. *IEEE Trans. Image Process.* **32**, 377–391 (2023).

9. Badnava, B., Chakareski, J. & Hashemi, M. Neural-enhanced rate adaptation and computation distribution for emerging mmWave multi-user 3D video streaming systems. *IEEE Trans. Multimed.* (2025). Accepted.
10. Wang, Z., Rao, M., Ye, S., Song, W. & Lu, F. Towards spatial computing: Recent advances in multimodal natural interaction for XR headsets. *arXiv preprint arXiv:2502.07598* (2025).
11. Bommasani, R. *et al.* On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258* (2021).
12. Yang, L. *et al.* Diffusion models: A comprehensive survey of methods and applications. *ACM Comput. Surv.* **56**, 1–39 (2023).
13. Naveed, H. *et al.* A comprehensive overview of large language models. *arXiv preprint arXiv:2307.06435* (2023).
14. Luo, S. *et al.* Delving into multi-modal multi-task foundation models for road scene understanding: From learning paradigm perspectives. *IEEE Transactions on Intell. Veh.* (2024).
15. McMahan, B., Moore, E., Ramage, D., Hampson, S. & y Arcas, B. A. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, 1273–1282 (PMLR, 2017).
16. Qiao, Y. *et al.* Towards artificial general or personalized intelligence? A survey on foundation models for personalized federated intelligence. *arXiv preprint arXiv:2505.06907* (2025).
17. Chen, H., Zhang, Y., Krompass, D., Gu, J. & Tresp, V. FedDAT: An approach for foundation model finetuning in multi-modal heterogeneous federated learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, 11285–11293 (2024).
18. Wang, T., Du, Y., Gong, Y., Choo, K.-K. R. & Guo, Y. Applications of federated learning in mobile health: Scoping review. *J. Med. Internet Res.* **25**, e43006 (2023).
19. Savazzi, S., Nicoli, M. & Rampa, V. Federated learning with cooperating devices: A consensus approach for massive IoT networks. *IEEE Internet Things J.* **7**, 4641–4654 (2020).
20. Nguyen, D. C. *et al.* Federated learning for industrial internet of things in future industries. *IEEE Wirel. Commun.* **28**, 192–199 (2021).
21. Elbir, A. M., Soner, B., Çöleri, S., Gündüz, D. & Bennis, M. Federated learning in vehicular networks. In *IEEE Int. Mediterranean Conf. Commun. Netw. (MeditCom)*, 72–77 (IEEE, 2022).
22. Han, D.-J., Hosseinalipour, S., Love, D. J., Chiang, M. & Brinton, C. G. Cooperative federated learning over hybrid terrestrial and non-terrestrial networks. In *IEEE Int. Conf. Commun. (ICC)*, 2992–2997 (IEEE, 2024).
23. Moor, M. *et al.* Foundation models for generalist medical artificial intelligence. *Nature* **616**, 259–265 (2023).
24. Afzal, M. Z. *et al.* Next generation XR systems-large language models meet augmented and virtual reality. *IEEE computer graphics applications* (2025).
25. Li, X., Tian, Y., Ye, P., Duan, H. & Wang, F.-Y. A novel scenarios engineering methodology for foundation models in metaverse. *IEEE Transactions on Syst. Man, Cybern. Syst.* **53**, 2148–2159 (2022).
26. Meta. Reality Labs. <https://tech.facebook.com/reality-labs/> (2025). Accessed: May, 2025.
27. NVIDIA. Omniverse. <https://www.nvidia.com/en-us/omniverse/> (2025). Accessed: May, 2025.
28. Borazjani, K. *et al.* Multi-modal multi-task (M3T) federated foundation models for embodied AI: Potentials and challenges for edge integration. *arXiv preprint arXiv:2505.11191* (2025).
29. Du, J. *et al.* Distributed foundation models for multi-modal learning in 6G wireless networks. *IEEE Wirel. Commun.* **31**, 20–30 (2024).
30. Li, X., Peng, L., Wang, Y.-P. & Zhang, W. Open challenges and opportunities in federated foundation models towards biomedical healthcare. *BioData Min.* **18**, 2 (2025).
31. Zhuang, W., Chen, C. & Lyu, L. When foundation model meets federated learning: Motivations, challenges, and future directions. *arXiv preprint arXiv:2306.15546* (2023).
32. Chen, J. & Zhang, A. On disentanglement of asymmetrical knowledge transfer for modality-task agnostic federated learning. In *Proc. AAAI Conf. Artificial Intell.*, vol. 38, 11311–11319 (2024).
33. Xiao, C. *et al.* Configurable foundation models: Building LLMs from a modular perspective. *arXiv preprint arXiv:2409.02877* (2024).
34. Brown, T. *et al.* Language models are few-shot learners. *Adv. Neural Inf. Process. Syst.* **33**, 1877–1901 (2020).

35. Radford, A. *et al.* Learning transferable visual models from natural language supervision. In *Int. Conf. Machine Learn.*, 8748–8763 (PmLR, 2021).
36. Long, G., Shen, T., Jiang, J., Blumenstein, M. *et al.* Dual-personalizing adapter for federated foundation models. *Adv. Neural Inf. Process. Syst.* **37**, 39409–39433 (2024).
37. Jia, M. *et al.* Visual prompt tuning. In *European conference on computer vision*, 709–727 (Springer, 2022).
38. Parasnis, R., Hosseinalipour, S., Chu, Y.-W., Chiang, M. & Brinton, C. G. Energy-efficient connectivity-aware learning over time-varying D2D networks. *IEEE J. Sel. Top. Signal Process.* (2024).
39. Abdisarabshali, P., Kim, K. T., Langberg, M., Su, W. & Hosseinalipour, S. Dynamic D2D-assisted federated learning over O-RAN: Performance analysis, MAC scheduler, and asymmetric user selection. *arXiv preprint arXiv:2404.06324* (2024).
40. Zehtabi, S., Han, D.-J., Parasnis, R., Hosseinalipour, S. & Brinton, C. G. Decentralized sporadic federated learning: A unified methodology with generalized convergence guarantees. *arXiv e-prints arXiv-2402* (2024).
41. Lin, F. P.-C., Hosseinalipour, S., Azam, S. S., Brinton, C. G. & Michelusi, N. Semi-decentralized federated learning with cooperative D2D local model aggregations. *IEEE J. on Sel. Areas Commun.* **39**, 3851–3869 (2021).
42. Apple. Introducing apple’s on-device and server foundation models. <https://machinelearning.apple.com/research/introducing-apple-foundation-models> (2025).
43. Rajan, S. K., Prabhakar, G. *et al.* Heads up eXperience (HUX): Always-on AI companion for human computer environment interaction. *arXiv preprint arXiv:2407.19492* (2024).
44. Xu, J., He, S., Jiang, H., Yang, Y. & Cai, S. Design and implementation of an english lesson based on handwriting recognition and augmented reality in primary school. *Int. Assoc. for Dev. Inf. Soc.* (2019).
45. Raj, S. *et al.* Augmented reality and deep learning based system for assisting assembly process. *J. on Multimodal User Interfaces* **18**, 119–133 (2024).
46. Musk, E. *et al.* An integrated brain-machine interface platform with thousands of channels. *J. medical Internet research* **21**, e16194 (2019).
47. Zhu, H. Y., Hieu, N. Q., Hoang, D. T., Nguyen, D. N. & Lin, C.-T. A human-centric metaverse enabled by brain-computer interface: A survey. *IEEE Commun. Surv. & Tutorials* (2024).
48. Vourvopoulos, A. *et al.* Effects of a brain-computer interface with virtual reality (VR) neurofeedback: A pilot study in chronic stroke patients. *Front. human neuroscience* **13**, 210 (2019).
49. Vourvopoulos, A. *et al.* Efficacy and brain imaging correlates of an immersive motor imagery BCI-driven VR system for upper limb motor rehabilitation: A clinical case report. *Front. human neuroscience* **13**, 244 (2019).
50. Yousri, R. *et al.* Illusionx: An LLM-powered mixed reality personal companion. *arXiv preprint arXiv:2402.07924* (2024).
51. Lima, R., Chirico, A., Gaggioli, A., Gamboa, H. & i Badia, S. B. Real-time emotion regulation in virtual reality: An adaptive experience using breathing biofeedback. *preprint (Version 1) available at Res. Sq.* (2025).
52. Merrill Jr, K., Kim, J. & Collins, C. AI companions for lonely individuals and the role of social presence. *Commun. Res. Reports* **39**, 93–103 (2022).
53. Bovo, R. *et al.* Embardiment: an embodied AI agent for productivity in XR. In *IEEE Conf. Virtual Reality 3D User Interfaces (VR)*, 708–717 (IEEE, 2025).
54. Lopes, P., Fachada, N. & Fonseca, M. Closing the loop: A systematic review of experience-driven game adaptation. *arXiv preprint arXiv:2505.01351* (2025).
55. Savaş, E. H. *et al.* A biofeedback based virtual reality game for pediatric population (BioVirtualPed): A feasibility trial. In *Seminars in Oncology Nursing*, vol. 40, 151615 (Elsevier, 2024).
56. Klein Haneveld, L. *et al.* Exploring the added value of virtual reality biofeedback game deep in forensic psychiatric inpatient care—a qualitative study. *Front. psychology* **14**, 1201485 (2023).
57. de Lima, E. S., Silva, B. M. & Galam, G. T. Towards the design of adaptive virtual reality horror games: A model of players’ fears using machine learning and player modeling. In *Brazilian Symp. Comput. Games Digital Entertainment (SBGames)*, 171–177 (IEEE, 2020).
58. Zaib, S. E. & Yamamura, M. Using heart rate and machine learning for VR horror game personalization. In *2022 IEEE Conference on Games (CoG)*, 213–220 (IEEE, 2022).

59. Borazjani, K., Abdisarabshali, P., Khosravan, N. & Hosseinalipour, S. Redefining non-IID data in federated learning for computer vision tasks: Migrating from labels to embeddings for task-specific data distributions. *arXiv preprint arXiv:2503.14553* (2025).
60. Di Luca, M. & Mahnan, A. Perceptual limits of visual-haptic simultaneity in virtual reality interactions. In *2019 IEEE World Haptics Conference (WHC)*, 67–72 (IEEE, 2019).
61. Kirkpatrick, J. *et al.* Overcoming catastrophic forgetting in neural networks. *Proc. national academy sciences* **114**, 3521–3526 (2017).
62. Wang, L., Singh, S., Chakareski, J., Hajiesmaili, M. & Sitaraman, R. K. BONES: Near-optimal neural-enhanced video streaming. In *Proc. SIGMETRICS/Performance Int’l Conf. on Measure. Analysis Comput. Syst.* (ACM/IFIP, Venice, Italy, 2024).
63. Wang, L., Singh, S., Chakareski, J., Hajiesmaili, M. & Sitaraman, R. K. Bones: Near-optimal neural-enhanced video streaming. *Proc. ACM Meas. Analysis Comput. Syst.* **8**, 1–28 (2024).
64. Chakareski, J. Vr/ar immersive communication: Caching, edge computing, and transmission trade-offs. In *Proc. Workshop Virtual Reality Augment. Reality Netw.*, 36–41 (2017).
65. Chakareski, J., Khan, M. & Yuksel, M. Toward enabling next-generation societal virtual reality applications for virtual human teleportation: A novel future system concept and computation-communication-signal representation trade-offs. *IEEE Signal Process. Mag.* **39**, 22–41 (2022).
66. Tome, D., Peluse, P., Agapito, L. & Badino, H. xR-EgoPose: Egocentric 3D human pose from an HMD camera. In *Proc. IEEE/CVF Int. Conf. Comput. Vision*, 7728–7738 (2019).
67. Liu, Y. *et al.* Hoi4d: A 4D egocentric dataset for category-level human-object interaction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 21013–21022 (2022).
68. Reimat, I. *et al.* CWIPC-SXR: Point cloud dynamic human dataset for social XR. In *Proc. 12th ACM Multimedia Syst. Conf.*, 300–306 (2021).
69. Lohesara, F. G., Freitas, D. R., Guillemot, C., Eguiazarian, K. & Knorr, S. Headset: Human emotion awareness under partial occlusions multimodal dataset. *IEEE Trans. Vis. Comput. Graph.* **29**, 4686–4696 (2023).
70. Bhattacharya, A. *et al.* AMIS: An audiovisual dataset for multimodal XR research. In *Proc. 16th ACM Multimedia Syst. Conf.*, 291–297 (2025).