On the Generalization and Personalization Tradeoff in Agentic AI Networks

Xubo Li*, Yong Xiao * § ¶, Yingyu Li‡

*School of Elect. Inform. & Commun., Huazhong Univ. of Science & Technology, Wuhan, China

§Pengcheng National Laboratory, Shenzhen, China

¶Pazhou Lab, Guangzhou, China

‡School of Mech. Eng. and Elect. Inform., China University of Geosciences, Wuhan, China

IEEE GLOBECOM'25

January 17, 2026

Outline

- ☐ Background and Motivation
- ☐ System Model & Problem Formulation
- ☐ Workflow of MAN
- ☐ Theoretical & Experimental Results
- ☐ Conclusions

Background

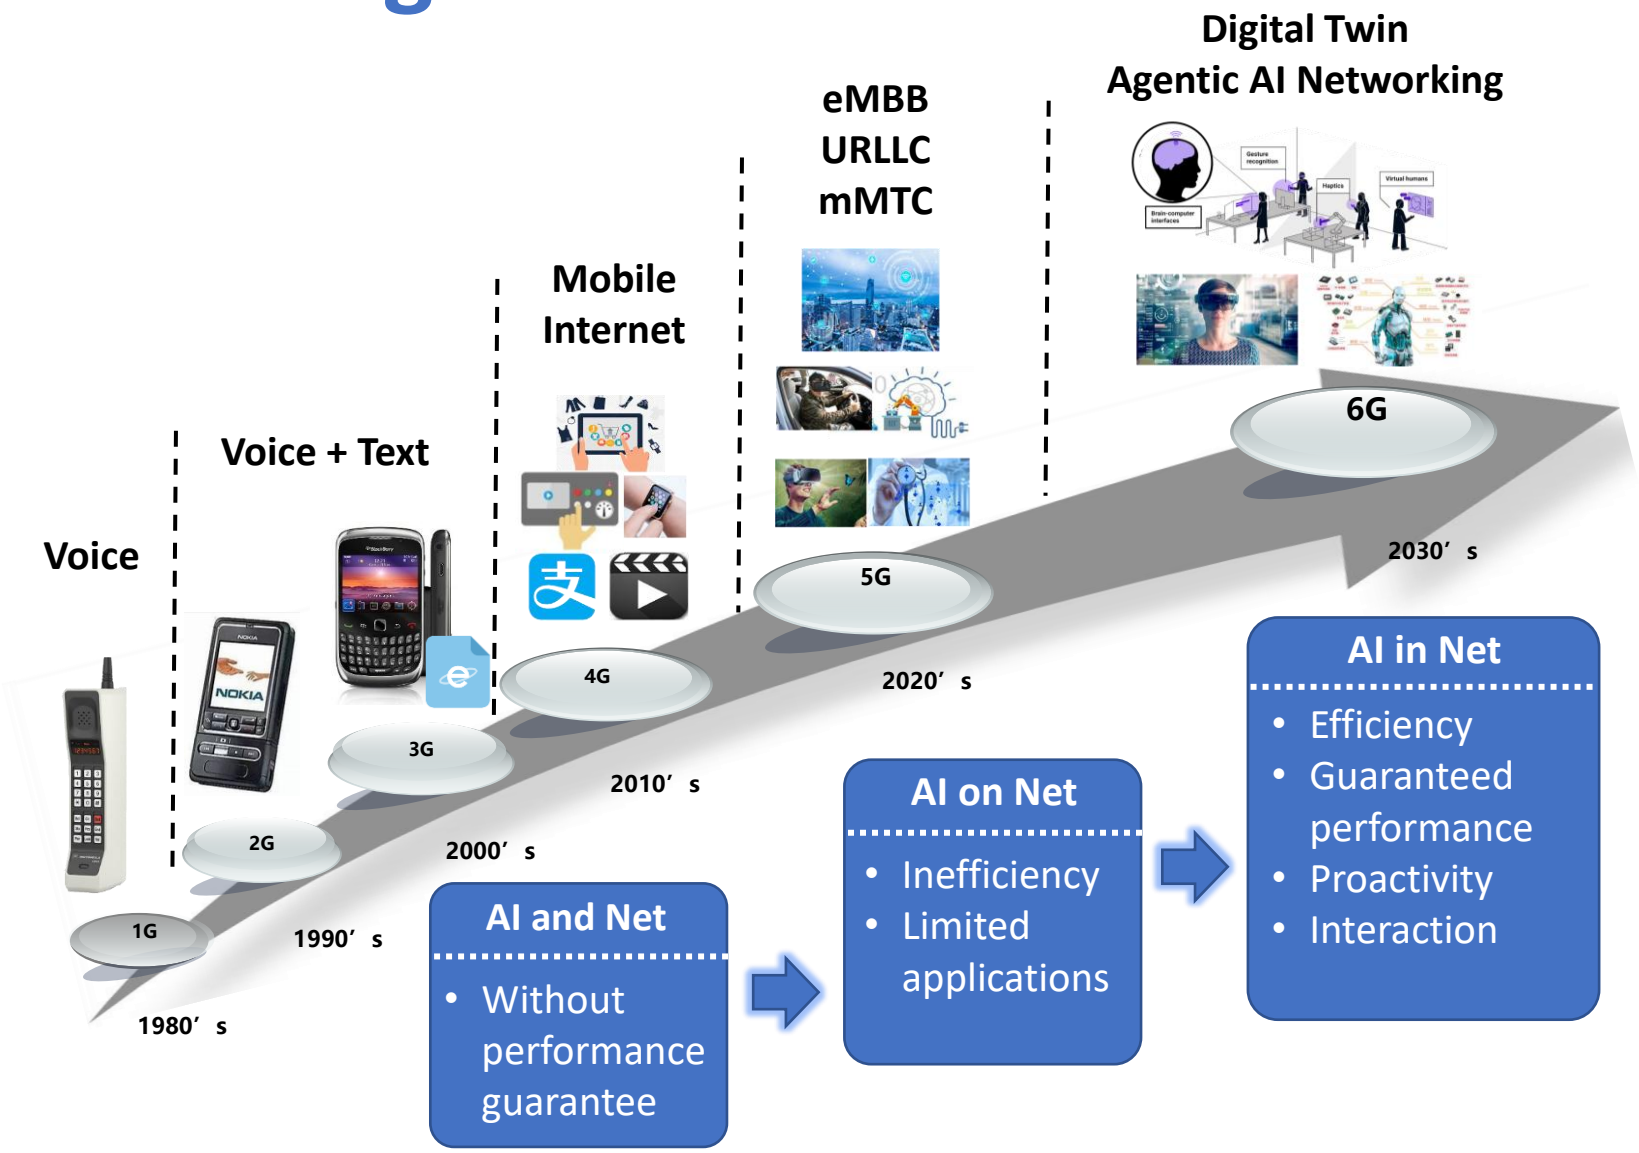

AI + Communication Networks:

- × Passive learning framework
- × AI-induced network congestion
- × Data privacy leakage risk
- × Without performance guarantee

Agentic AI Networks (AgentNet):

- ✓ Perceive dynamic environments
- ✓ Execute autonomous decision-making
- ✓ Refine operational strategies via continuous collaboration

Motivation

Two critical metrics for agentic AI networking:

- ***Generalization:*** ensures each agent can apply previously learned knowledge or skills across diverse, unseen tasks and environments.
- ***Personalization:*** allows the agent to tailor its behavior, responses, or decision-making to individual users' personal preferences.

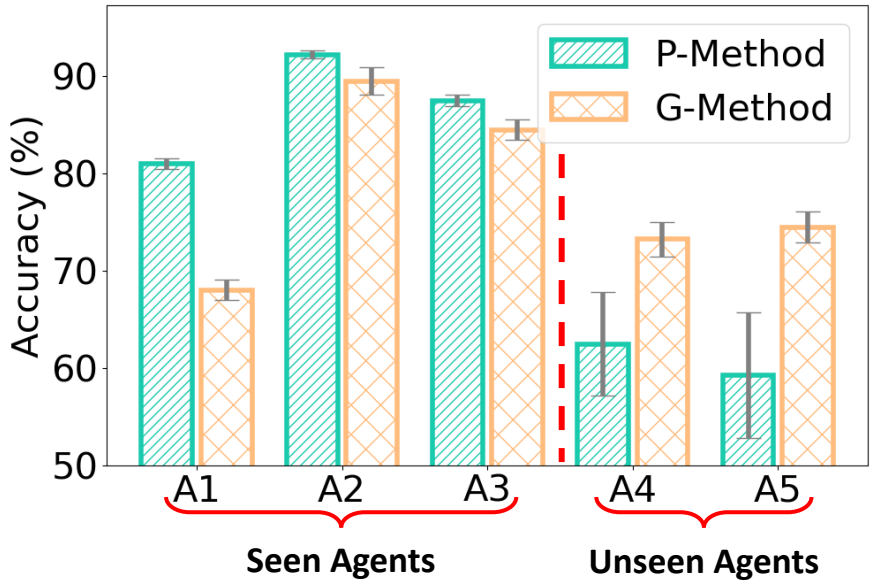

** A1 represents Agent 1, and so on*

Generalization ⟺ **Personalization**

Conflict

How can we effectively balance the tradeoff between generalization and personalization in agentic AI networking?

Key Contributions

- ✓ Propose a meta learning-based agentic AI networking framework (MAN) to strike an effective balance between global generalization and local personalization
- ✓ Derive asymptotic theoretical bounds for generalization and personalization errors of MAN respectively
- ✓ Conduct extensive measurements on real-world dataset to demonstrate that proposed MAN can improve the tradeoff between personalization and generalization in various environments and application scenarios.

Outline

- ☐ Background and Motivation
- ☐ System Model & Problem Formulation
- ☐ Workflow of MAN
- ☐ Theoretical & Experimental Results
- ☐ Conclusions

System Model

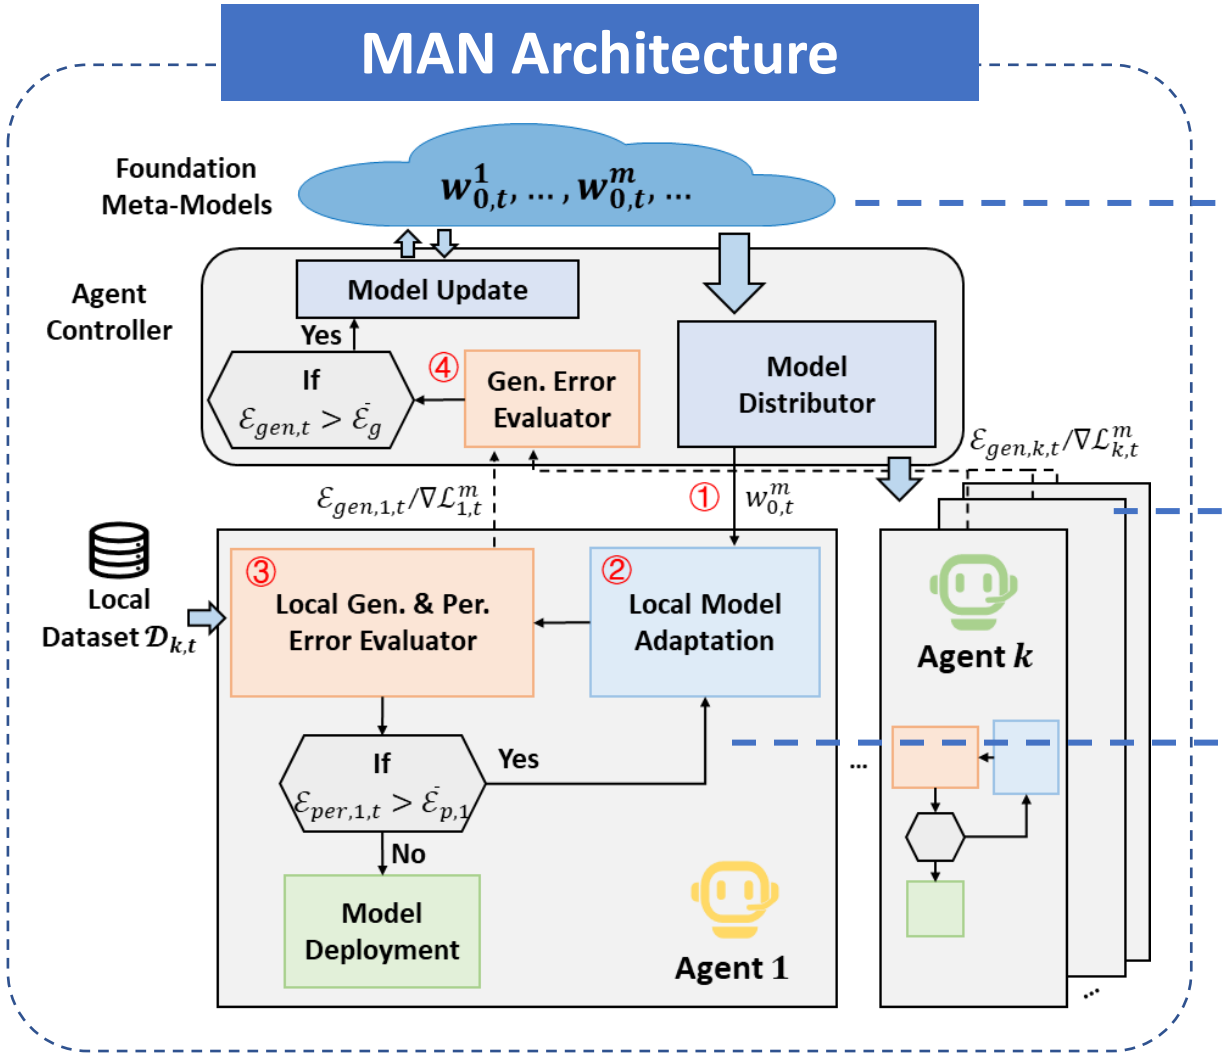

Foundation models are deployed at the cloud targeting at different tasks and constructed and maintained by all agents for **better generalization capability**.

The environment of each agent **change** with time.

Each agent **finetune the meta-model** according to its local dataset to enhance the **personalization performance**.

System Model

MAN Architecture

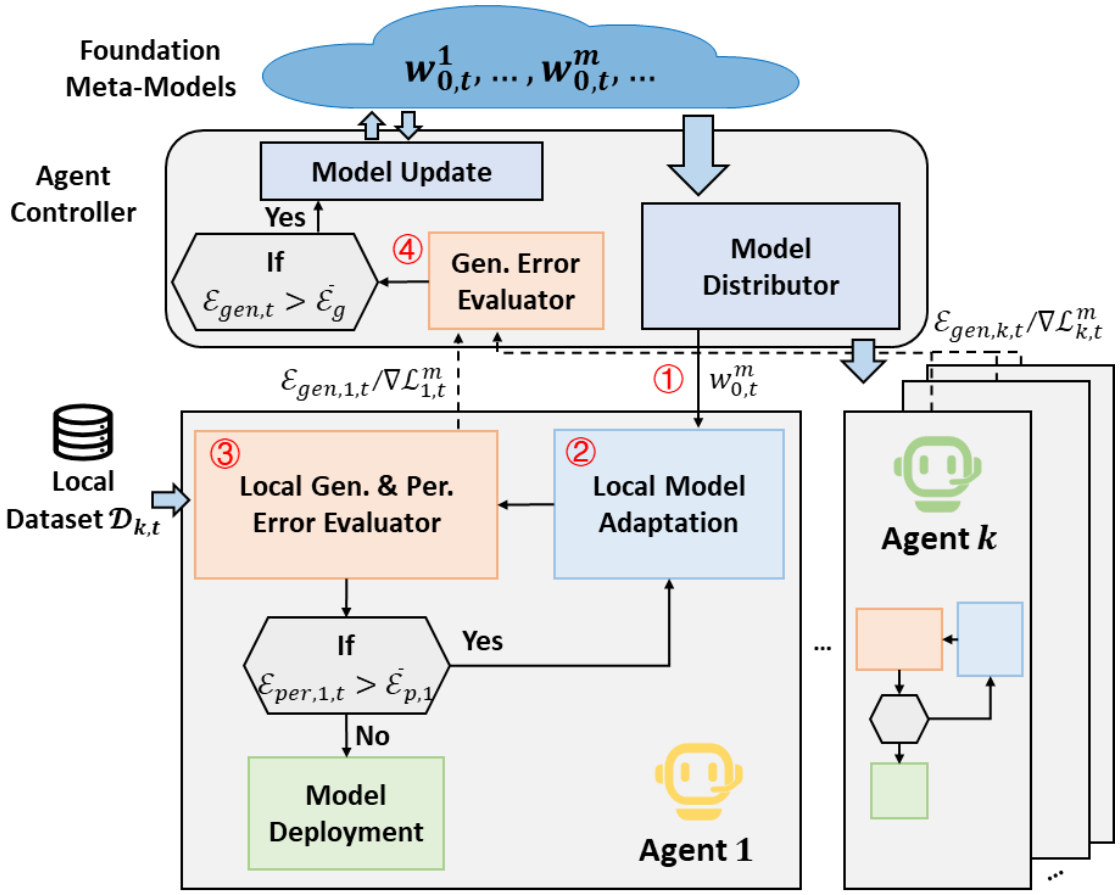

Key Procedures

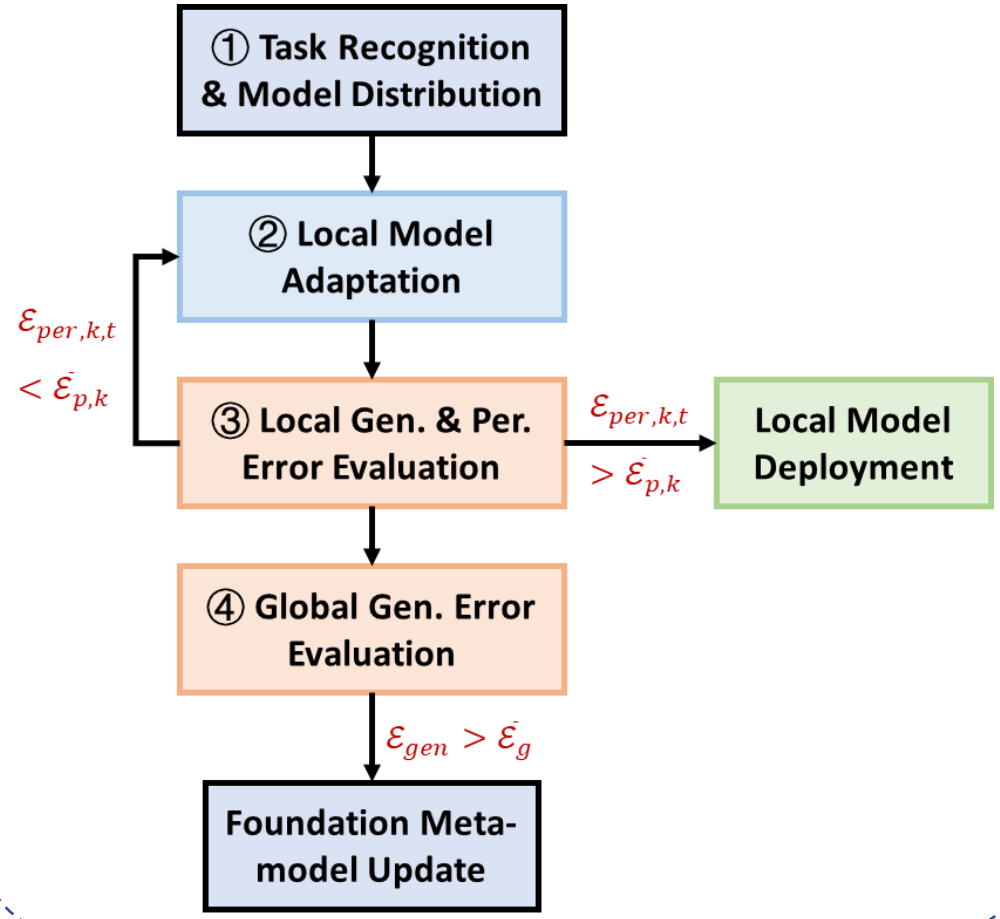

Problem Formulation

Definition1 (Generalization Error)

Define the generalization error $\mathcal{E}_{gen}$ of all the agents as the overall performance discrepancy between the loss obtained by model $w_{k,t}^m$ trained based on the dataset $\mathcal{D}_{k,t}$, denoted as $\hat{\mathcal{L}}_k^m(w_{k,t}^m, \mathcal{D}_{k,t})$, and that developed based on the ground truth distribution of the environmental state, denoted as $\mathcal{L}^m(w_{k,t}^m)$, given by

$$\mathcal{E}_{gen} \coloneqq \frac{1}{K}\sum_{k\in\mathcal{K}} \mathbb{E}_{\mathcal{D}_{k,t}}\left[\mathcal{L}^m\left(w_{k,t}^m\right) - \hat{\mathcal{L}}_k^m\left(w_{k,t}^m, \mathcal{D}_{k,t}\right)\right]$$

***Notation:** m - index of task; k - index of agent; t - time slot

- To improve generalization performance, it requires to learn common patterns and features from a wide range of diverse datasets

Problem Formulation

Definition2 (Personalization Error)

Suppose that for each agent k, there exists a locally optimal model $w_{k,t}^{m*}$. We assume each agent k can adopt a personalization solution, denoted as $w_{k,t}^{m} = g_k(w_{0,t}^{m}, \mathcal{D}_{k,t})$ where $w_{0,t}^{m}$ denotes the foundation model obtained from the agent controller at the beginning of time slot t for task m. Define the personalization error as

$$\mathcal{E}_{per} := \left\| \mathcal{L}_k^m\left(w_{k,t}^{m*}\right) - \mathcal{L}_k^m\left(g_k(w_{0,t}^{m}, \mathcal{D}_{k,t})\right) \right\|$$

***Notation:** m - index of task; k - index of agent; t - time slot

- To enhance the personalization capability, it requires the model focus on the unique preferences and personalized behaviors of a single agent.

Problem Formulation

Main Objective

To construct task-oriented models that minimize their local task-specific losses and at the same time ensure both generalization and personalization errors are below tolerable levels, we formulate the optimization problem as:

$$\min_{w_{k,t}^m} \quad \hat{\mathcal{L}}_k^m\left(w_{k,t}^m, \mathcal{D}_{k,t}\right), \qquad \forall k \in \mathcal{K}$$

$$\text{s.t.} \quad \mathcal{E}_{gen} \leq \bar{\mathcal{E}}_g \quad \text{and} \quad \mathcal{E}_{per} \leq \bar{\mathcal{E}}_p$$

Maximum tolerable levels of gen. and per. errors

Outline

- ☐ Background and Motivation
- ☐ System Model & Problem Formulation
- ☐ Workflow of MAN
- ☐ Theoretical & Experimental Results
- ☐ Conclusions

Workflow of MAN

Algorithm 1: MAN Workflow

Input: Inner layer step size β; Outer layer step size γ.

Output: Meta-model $w_{0,t}^{m}$; Personalized models for each agent $\langle w_{k,t}^{m} \rangle_{k \in \mathcal{K}}$.

// **Each Agent**

for *agent* $k \in \mathcal{K}$ **do**

  while $\mathcal{E}_{per,k} > \bar{\mathcal{E}}_{per,k}$ **do**

    Sample a subset of data $\mathcal{D}_{k,t}^{in}$;

    Compute adapted model by using (5);

  Periodically evaluate local generalization error by (1);

  Upload $\mathcal{E}_{gen,k}$ to agent controller;

// **Agent Controller**

Periodically evaluate generalization error;

for $\mathcal{E}_{gen} > \bar{\mathcal{E}}_{g}$ **do**

  Update meta-model by using (6);

Workflow of MAN

Algorithm 1: MAN Workflow

Input: Inner layer step size β; Outer layer step size γ.

Output: Meta-model $w_{0,t}^m$; Personalized models for each agent $\langle w_{k,t}^m \rangle_{k \in \mathcal{K}}$.

// **Each Agent**

for *agent* $k \in \mathcal{K}$ **do**

  while $\mathcal{E}_{per,k} > \bar{\mathcal{E}}_{per,k}$ **do**

    Sample a subset of data $\mathcal{D}_{k,t}^{in}$;

    Compute adapted model by using (5);

  Periodically evaluate local generalization error by (1);

  Upload $\mathcal{E}_{gen,k}$ to agent controller;

// **Agent Controller**

Periodically evaluate generalization error;

for $\mathcal{E}_{gen} > \bar{\mathcal{E}}_g$ **do**

  Update meta-model by using (6);

Local model adaptation by

$$w_{k,t}^m = g_k\left(w_{0,t}^m, \mathcal{D}_{k,t}^{in}\right) = w_{0,t}^m - \beta \nabla \hat{\mathcal{L}}_k^m\left(w_{0,t}^m, \mathcal{D}_{k,t}^{in}\right)$$

Workflow of MAN

Algorithm 1: MAN Workflow

Input: Inner layer step size β; Outer layer step size γ.

Output: Meta-model $w_{0,t}^m$; Personalized models for each agent $\langle w_{k,t}^m \rangle_{k \in \mathcal{K}}$.

// **Each Agent**

for *agent* $k \in \mathcal{K}$ **do**

  while $\mathcal{E}_{per,k} > \bar{\mathcal{E}}_{per,k}$ **do**

    Sample a subset of data $\mathcal{D}_{k,t}^{in}$;

    Compute adapted model by using (5);

  Periodically evaluate local generalization error by (1);

  Upload $\mathcal{E}_{gen,k}$ to agent controller;

// **Agent Controller**

Periodically evaluate generalization error;

for $\mathcal{E}_{gen} > \bar{\mathcal{E}}_g$ **do**

  Update meta-model by using (6);

Local model adaptation by

$$w_{k,t}^m = g_k\left(w_{0,t}^m, \mathcal{D}_{k,t}^{in}\right) = w_{0,t}^m - \beta \nabla \hat{\mathcal{L}}_k^m\left(w_{0,t}^m, \mathcal{D}_{k,t}^{in}\right)$$

Meta-model optimization by

$$w_{0,t}^m \leftarrow w_{0,t}^m - \frac{\gamma}{K} \nabla \hat{\mathcal{L}}_k^m\left(w_{k,t}^m, \mathcal{D}_{k,t}^{out}\right)$$

Outline

- ☐ Background and Motivation
- ☐ System Model & Problem Formulation
- ☐ Workflow of MAN
- ☐ Theoretical & Experimental Results
- ☐ Conclusions

Theoretical Results

Assumptions

Loss function $\mathcal{L}(w; x, y)$ satisfies

- L-smooth
- μ-strongly convex
- ρ-Lipschitz continuous
- G-gradient bounded

Upper Bound of Generalization Error

With $\beta \leq \min\left\{\frac{1}{2L}, \frac{\mu}{8\rho G}\right\}$ and $\gamma \leq \frac{1}{4L+2\beta\rho G}$, the generalization error can be upper bounded by

$$\mathcal{E}_{gen} \leq \mathcal{O}\left(\frac{G^2(1+\beta LB)S}{KD}\right)\left[1-\left(1-\gamma c_g\right)^{ET}\right]$$

$$c_g = \frac{2\mu(2L+\rho\beta G)}{16(2L+\rho\beta G)+\mu}$$

***Notation:**

B – batch size, i.e., the size of $\mathcal{D}_{k,t}^{in}$ and $\mathcal{D}_{k,t}^{out}$;

E – local adaptation step;

D – size of local dataset;

S – number of collaborative agents ;

T – global iterations

β & γ – learning rates

Theoretical Results

Assumptions

Loss function $\mathcal{L}(w; x, y)$ satisfies

- L-smooth
- μ-strongly convex
- ρ-Lipschitz continuous
- G-gradient bounded
- Bounded gradient variance

$\mathbb{E}[\|\nabla\mathcal{L}(w; x, y) - \mathbb{E}[\nabla\mathcal{L}(w; x, y)]\|^2] \leq \sigma^2$

Upper Bound of Personalization Error

Define the gap between the foundation meta-model w_0^m and the optimal model of target agent w_k^{m*} as c_I,

i.e., $\|w_0^m - w_k^{m*}\| \leq c_I$, with $\beta \leq \frac{2}{\mu(E+1)}$, we have

$$\mathcal{E}_{per} \leq \mathcal{O}\left(\frac{\sigma^2}{\mu E}\right) + \mathcal{O}\left(\frac{L}{2} c_I^2 \left(1 - \frac{\mu}{L}\right)^E\right)$$

Discussion of tradeoff

Global collaboration view

Iterations T, collaborative agents S

Better Generalization — Small — Large — **Better Personalization**

Adaptation steps E, step size β

Local adaptation view

Experiment Setup

Datasets

- **MNIST**[1]

  Handwritten digit image dataset where each image consists of 28×28 gray-scale pixels

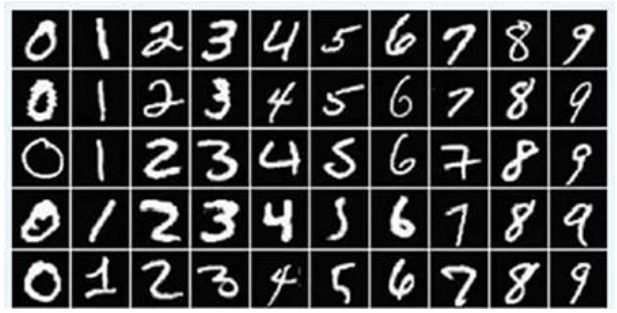

- **Widar 3.0**[2]

  RF-based human gesture recognition dataset recorded by 6 distributed Wi-Fi receivers

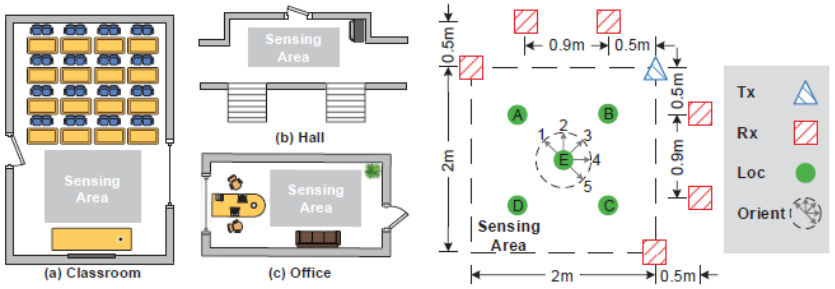

Default Parameters

Dataset	MNIST	Widar 3.0
Number of agents (K)	10	
Number of collaborative agents (S)	5	
Mini-batch size (B)	32	
Learning rate	Inner layer (β): 1e-3 Outer layer (γ): 5e-3	
Collaboration rounds (T)	1,000	200
Local adaptation steps (E)	30	

[1] Y. LeCun, "The mnist database of handwritten digits," http://yann.lecun.com/exdb/mnist/, 1998.

[2] Y. Zheng, Y. Zhang, K. Qian, G. Zhang, Y. Liu, C. Wu, and Z. Yang, "Zero-effort cross-domain gesture recognition with wi-fi," in ACM MobiSys, Seoul, Korea, Jun. 2019.

Performance of MAN

TABLE I: Recognition accuracy of agents deployed at different environments trained by different algorithms, including local training, joint training, and the proposed MAN based on (a) MNIST and (b) Widar3.0.

Dataset	Eval. Condition	Local Training	Joint Training	Ours
MNIST	worst seen	96.17±0.27	95.59±0.45 ↓	97.28±0.04 ↑
	best seen	98.39±0.08	97.88±0.24 ↓	99.33±0.03 ↑
	worst unseen	/	93.52±0.82	96.43±0.28 ↑
	best unseen	/	96.85±0.73	97.98±0.26 ↑
Widar3.0	worst seen	81.03±0.53	68.06±1.01 ↓	85.06±0.78 ↑
	best seen	96.22±0.43 ↑	84.50±1.42 ↓	95.11±0.69
	worst unseen	/	83.25±1.77 ↑	82.79±1.12
	best unseen	/	84.51±1.62	86.69±0.48 ↑

worst

best

10 E-Agents

5 E-Agents

personalization performance

generalization performance

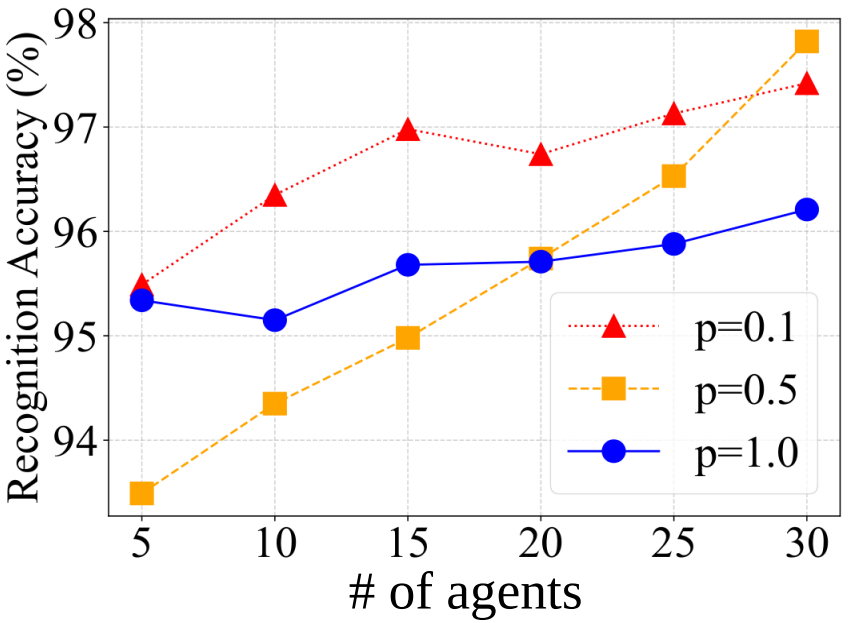

** p denotes the non-i.i.d. degree*

Observation 1

- ✓ The proposed MAN demonstrates superior performance in both personalization and generalization, achieving a favorable tradeoff.

Observation 2

- ✓ Accuracy always improves as the number of agents increases.
- ✓ This improvement is most pronounced in moderately heterogeneous scenarios.

Verification of Theoretical Results

Global collaboration view

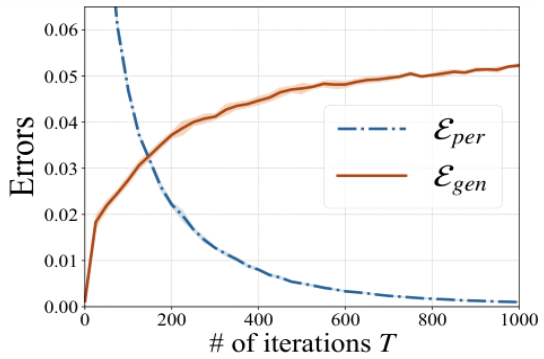

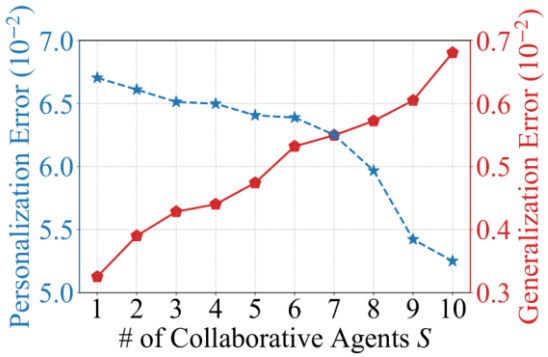

Local Adaptation view

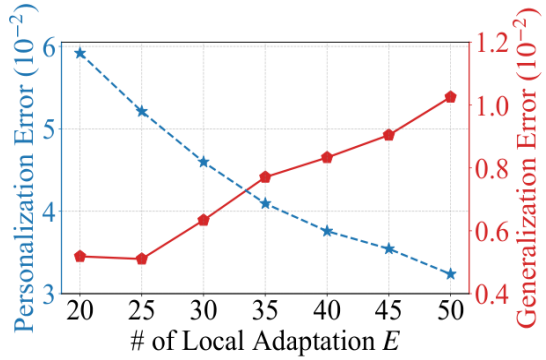

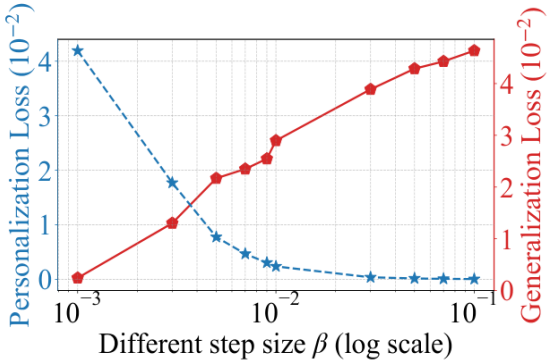

Theoretical bounds for generalization

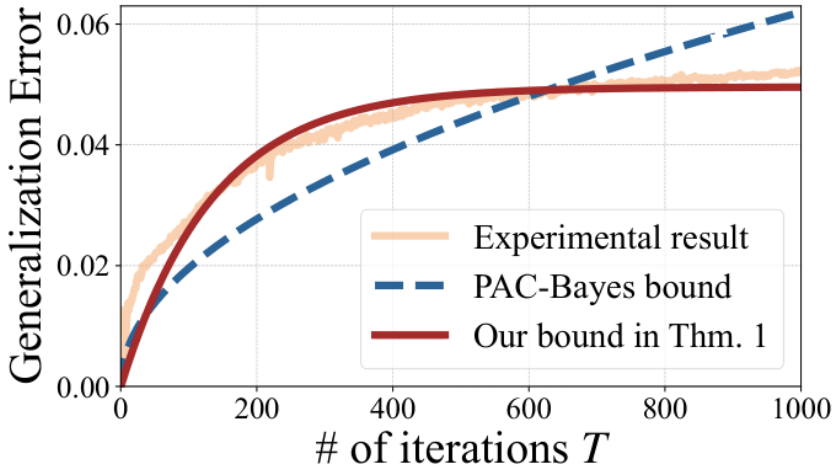

Observation 3

- ✓ Our theoretical results hold generally.
- ✓ The tradeoff can only be optimized but not entirely eliminated.

Observation 4

- ✓ The asymptotic stability of our bound enables a better characterization of generalization performance, especially in large T regime.

Outline

- ☐ Background and Motivation
- ☐ System Model & Problem Formulation
- ☐ Workflow of MAN
- ☐ Theoretical & Experimental Results
- ☐ Conclusions

Conclusions

- ✓ Proposed a meta learning-based agentic AI networking (MAN) framework to efficiently balance generalization and personalization in multi-agent networks
- ✓ Derived asymptotic theoretical bounds on generalization and personalization errors respectively
- ✓ Conducted extensive measurements on real-world dataset to verify the theoretical results

Thank You!

For more information, please contact:
Xubo Li
(xuboli@hust.edu.cn)
Yong Xiao
（yongxiao@hust.edu.cn）