

◆ **Projet 5 : SMARTNET : décembre 2023 – décembre 2027**



Yassine HADJADJ AOUL

Full Professor, Univ Rennes, Inria,
CNRS, IRISA Ermine Team-project

SMARTNET: Proactive Multi-Domain Orchestration for 6G Networks

Parsa Rajabzadeh, Abdelkader Outtagarts (Nokia Bell Labs)

Yassine Hadjadj-Aoul, Soraya Ait-Chellouche (Inria, Univ. Rennes, CNRS, IRISA)

Project context — DEFI SMARTNET

Research context

Contribution 1 — HERO

Contribution 2 — NeuroPulse

Summary

SmartNet: AI methods for smart network management


Why

- **Exploding Complexity:** Massive network slicing across multi-tenant, heterogeneous administrative domains.
- **Resource Contention:** Dynamic scaling under stringent SLA targets & energy limits.
- **Observability Bottlenecks:** Limited inter-domain visibility due to strict business and data privacy
- **Network management becomes the bottleneck**

Project Vision: Transitioning from human-centric control toward **fully autonomous, Network Management.**

Who

- Six Inria team-projects (Aptikal, Coati, Ermine, Neo, Spades, Stack)
- Nokia Bell Labs (ML & Systems; Network Services & Security)



Research Context: 6G Demands End-to-End Multi-Domain Intelligence

- **6G (~2030):** ubiquitous connectivity integrating terrestrial, aerial, space, underwater networks.
- Network slices with stringent resources requirements must operate **simultaneously across administrative domains**.
- **Vision:** connected intelligence – self-optimizing, autonomous, **AI-enabled orchestrators** exchanging information to provide E2E network intelligence.

Problem: today's orchestration is mostly centralized, single-domain, reactive.

Core tension: A global view is needed for optimal E2E decisions, but domains are *administratively separated*, partially observable, and unwilling to expose private data.

Why multi-domain orchestration is difficult

Dynamical coupling:

Slice traffic loads vary continuously; traffic with changing magnitude *crosses* multiple domains.


An action of orchestrator i (e.g., scaling NFs) changes the future state of domain $j \neq i$.

Slices may be *shared* between domains $\Rightarrow$ hidden contention.

Limited information:

Each orchestrator fully observes *only its own* domain.

Cross-domain information is either unavailable or **private** (competing operators hide infrastructure).




Why multi-domain orchestration is difficult (continued)

Scale & combinatorics:

Service placement = mapping constrained VNF graphs onto a resource graph: NP-hard.

Edge deployment adds heterogeneity and tight resource budgets.

μ -services, dynamic scaling, topology changes
 $\Rightarrow$ static models break.



Consequence: Decision quality, response time, cost and energy targets cannot be met by myopic, single-domain, reactive policies $\Rightarrow$ **orchestrators must be proactive, holistic, and privacy-preserving.**

When Network Functions Become Intelligent: Orchestration Gets Harder

From NF to VNIF:


NFs embed AI models and become **Network Intelligent Functions** (VNIFs), with four roles: data generation, prediction, decision making.

Slices now carry SLA targets on the **AI itself**: accuracy, energy efficiency, responsiveness.

Added complexity:

Two coupled decisions: *how many NFs and which AI profile each one runs.*

Resource-intelligence coupling: a heavier model raises CPU/memory load, hence scaling needs; scaling and load change latency, energy and accuracy.



joint decision under SLA constraints

Scaling decisions in real time, under changing load and SLA priorities. New SLA priorities arrive anytime; **exhaustive profiling is impossible.** One under-performing VNIF can **break** the SLA of the whole chain.


Research problem: proactive multi-domain orchestration

Orchestration task

Scale resources (NFs) up/down per slice as traffic varies, in several coupled domains.

Why hard

- Traffic changes in real time; one domain's decision shifts load onto the next
- No central vantage point; domains do not share raw data
- Once NFs embed AI models, KPIs must also be met by the **AI itself**




Two questions structuring this talk

RQ1 (HERO) How many resources, and when? Proactive, holistic, privacy-preserving multi-domain scaling.

RQ2 (NeuroPulse) What intelligence runs inside the NFs? Generating AI profiles that satisfy SLA-derived KPIs.

From orchestrating resources to orchestrating intelligence



Key message: NeuroPulse is built *on top of* the autoscaling framework of HERO: HERO decides how many resource units to run; NeuroPulse decides what AI profile each unit runs.

HERO: Holistic Envisioned RL Multi-Domain Orchestration

Three elements per orchestrator

1. Predictive element:

Latent-ODE auto-encoder;
shares a *noised latent code*,
never raw history


2. Focus element:

gated
attention fuses all predicted
futures into one *holistic*
vision, biased to own domain

3. Decision-making unit:

actor-critic (A2C) acting on
current state + holistic future


Plus **federated training** of the RL
models, aggregated with attention
weights.



HERO: Experimental Setup

Emulation platform: Containernet + k8s-netSim (Knetsim)

Multi-domain network infrastructure



Illustrative topology; evaluated with 2, 3, and 5 clusters


Scale configurations		
Clusters	Min. CNFs	Max. CNFs
2	10	50
3	15	75
5	25	125

Training and evaluation	
Training duration	400 episodes
Episode length	200 steps
Reward range	[−100, +100]
Base attention γ	0.75


Autoscaling objective: Minimize the number of deployed CNF containers while keeping CPU and memory usage within $[\text{Threshold}_{\min}, \text{Threshold}_{\max}]$.

HERO: Performance and Scale Comparison

Scenario comparison




Federated models across scales



Key message: A2C + Latent ODE already achieves strong performance **without federation**; Fed A2C + Latent ODE achieves the **best average scores** in the tested scenarios.

NeuroPulse: High-Level Architecture




Key message: **Generate** an AI profile and **estimate** its KPIs before deployment; use **SLA-based reward and live feedback** to improve subsequent profiles.

NeuroPulse: Experimental Testbed & Setup

Emulation platform: Containernet + k8s-netSim (Knetsim) | **Traffic:** Iperf

From packet streams to AI workloads



Recognition example: sample an N -point waveform from packet windows and classify it as **normal or abnormal**.

Evaluation objective: compare generated AI profiles under live traffic, balancing **accuracy, energy, and latency** against service SLA priorities.


Candidate architecture search space

Parameter	Candidates
Hidden layers	1–5
Neurons per layer	16–256
Activations	ReLU, Tanh, Sigmoid

Models and evaluation

Generators: MetaRNN; auto-regressive RNN and MLP (REINFORCE).
Estimators: KNN, MLP, RF, GB, SVM, linear regression; best validation model retained.
Training: 5 services with distinct SLA priorities.


NeuroPulse: Performance and Slice Profiles



Key message: NeuroPulse generates **slice-specific AI profiles**, adapting the balance between accuracy, energy efficiency, and responsiveness to the distinct SLA priorities of each service.

Summary

SMARTNET vision: from fragmented, reactive management to coordinated AI-native network autonomy



HERO orchestrates where resources are needed; NeuroPulse orchestrates which intelligence should run.

Thank you for your attention!

References

-  P. Rajabzadeh, A. Outtagarts, Y. Hadjadj-Aoul, and S. Ait-Chellouche, *Reliable Feedback-Driven Federated Learning for Distributed NWDAFs at the Edge*, Computer Networks, vol. 277, article 111972, 2026.
-  P. Rajabzadeh, A. Outtagarts, Y. Hadjadj-Aoul, and S. Ait-Chellouche, *NeuroPulse: AI Empowered Network Functions via AI-Native Multi-Stakeholders Orchestration*, IEEE CCNC, pp. 1–6, 2026.
-  P. Rajabzadeh, A. Outtagarts, Y. Hadjadj-Aoul, and S. Ait-Chellouche, *HERO: Holistic Envisioned Reinforcement Learning Multi-Domain Orchestration with Latent ODE*, IEEE CCNC, pp. 1–7, 2025.
-  P. Rajabzadeh, A. Outtagarts, Y. Hadjadj-Aoul, and S. Ait-Chellouche, *Next-Gen AI Orchestration Platform*, ICIN, pp. 80–82, 2025.