AGENTIFICATION OF NUCLEAR SENSING DOMAINS FOR COLLECTIVE SPECTRAL REASONING

Anonymous authors
Paper under double-blind review

ABSTRACT

Modern machine learning methods have substantially advanced sensing by enabling increasingly sophisticated interpretation of acquired signals. Yet most existing approaches represent a sensing domain as a single computational object, producing monolithic vectors or sequences while leaving decomposition of the domain implicit. I introduce *agentification*, a learning paradigm in which a structured sensing domain is explicitly decomposed into localized subdomains and each subdomain is instantiated as a computational agent with explicit responsibility for local observation, learning, inference, and, optionally, interaction. Agentification therefore goes beyond partitioning by transforming domain components into computational entities that independently learn local behavior and contribute locally generated evidence to collective inference. I instantiate this paradigm in gamma-ray spectroscopy by transforming a 1,024-channel spectrum into 1,024 spectral agents, each using a Gaussian process to learn normal background behavior and quantify anomalous local observations. This setting targets a central nuclear-security challenge: weak source-induced counts can be masked by naturally fluctuating radiation background. The detector is source agnostic: Cs-137 and Co-60 response functions are used only to generate controlled evaluation anomalies and are never supplied to the inference framework. Across ten repeated 70/30 holdouts, agentification improves source-agnostic AUC over PCA, a nonlinear autoencoder, and the strongest monolithic baseline in the weakest source regime, while the advantage diminishes or reverses as anomalies become stronger. Agent-level evidence also provides intrinsic spectral localization, with channel-level localization AUCs of 0.718/0.789 for Cs-137/Co-60 at the weakest tested level and 0.842/0.883 at the strongest. Under partial sensing-domain degradation, agentification remains competitive and can reverse the intact-domain ordering under substantial contiguous channel loss. Simple neighboring-agent communication is not consistently beneficial, demonstrating that agent formation and communication are distinct design dimensions. These results establish agentification as a general approach for transforming structured sensing observations into populations of locally responsible computational entities and provide a computational foundation for Agentic Nuclear Sensing Systems (ANSS).

1 INTRODUCTION

Machine learning methods have been widely used in sensing to analyze measurements and extract useful inferences. A common paradigm across sensing applications is to represent acquired measurements as a monolithic vector, tensor, image, or sequence and process them as a unified computational object (Kamencay et al. (2024)). While computationally convenient, this representation can obscure the inherent structure of a sensing domain. Many observations contain naturally distinguishable spatial, temporal, spectral, or functional subdomains with different local characteristics and dependencies. Treating such domains as a single object therefore leaves their internal organization largely implicit.

Capturing this organization suggests a shift from treating constituent subdomains merely as features toward assigning explicit computational responsibility to localized regions. I introduce *agentification* as the transformation of a structured sensing domain into a population of computational agents through domain decomposition and the assignment of localized sensing responsibility. Each agent

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

learns the behavior of its assigned subdomain, evaluates new observations locally, and contributes its evidence to collective inference (Acharya et al. (2025)). Agentification therefore changes the computational organization of the sensing problem rather than merely changing its representation (Zhao et al. (2026)).

This formulation differs from conventional multi-agent learning, where the population typically exists before learning begins: agents may correspond to robots, vehicles, physical sensors, or controllers (Ito et al. (2013)). Agentification addresses a preceding representational question: *how can computational agents be constructed directly from the structure of an underlying sensing domain?* In the present work, this question is studied through gamma-ray spectroscopy using the finest physically interpretable partition: one spectral channel per agent.

Gamma-ray spectroscopy is particularly relevant to nuclear security because an operationally important challenge is detecting weak radioactive perturbations embedded within naturally fluctuating radiation background (Alamaniotis (2024)). A weak source may contribute only a small number of additional counts over localized portions of a spectrum. The challenge is therefore not only whether the complete spectrum is globally different, but whether localized portions exhibit small yet meaningful departures from expected behavior. This setting provides a physically meaningful test of the hypothesis that local computational responsibility can preserve weak evidence that may be diluted in a monolithic representation.

I transform each 1,024-channel spectrum into 1,024 spectral agents. Each agent learns normal background behavior for one channel using a Gaussian process (GP) (Lee et al. (2024)), independently evaluates an incoming count, and produces bounded local anomaly evidence. The 1,024 local beliefs are then aggregated into a source-agnostic anomaly score. Cs-137 and Co-60 detector response functions are used only to generate controlled evaluation anomalies and are never supplied to the detector; inference therefore identifies departures from learned normality rather than matching predefined source signatures (Ghawaly Jr et al. (2022)).

Because every agent retains responsibility for a known channel, the same representation also preserves the location of evidence supporting the collective decision. The aggregated score answers whether the observation is anomalous, while the distribution of local agent evidence indicates where anomalous behavior arises within the sensing domain. I evaluate this intrinsic localization quantitatively rather than assuming that local responsibility automatically yields perfect hard-region explanations.

The work is motivated more broadly by Agentic Nuclear Sensing Systems (ANSS), in which sensing elements become active computational participants in interpretation and reasoning rather than remaining passive measurement producers (Tsoukalas (2010)). The present study focuses on the agent-formation, local-learning, and collective-inference layers of that broader vision; adaptive sensing, higher-level decision making, and bounded physical action remain future directions.

The paper addresses four questions: (RQ1) Can a structured sensing domain be transformed into locally responsible learning agents? (RQ2) Does agentification improve sensitivity to weak anomalies relative to monolithic representations? (RQ3) Does agent-level evidence provide useful localization of anomalous spectral behavior? (RQ4) How does distributed local responsibility behave when portions of the sensing domain become unavailable, and does explicit inter-agent communication improve inference?

Figure 1 summarizes the proposed transformation from a structured sensing observation into locally responsible agents and collective inference. The contributions are fourfold: (1) a formalization of agentification for structured sensing domains, distinguishing agent formation from feature decomposition and predefined multi-agent systems; (2) a spectral-element realization with 1,024 independently learning GP agents for source-agnostic weak-anomaly detection; (3) quantitative evidence that local agent responses retain channel-level localization information; and (4) a characterization of agentification across anomaly-strength and sensing-degradation regimes, together with an ablation showing that explicit communication is a distinct design dimension rather than a prerequisite for agentification.

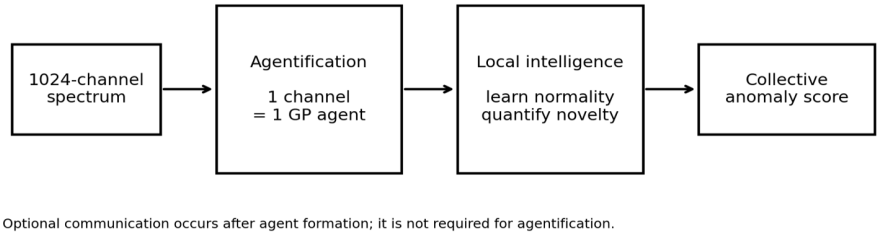

Figure 1: Agentification converts a structured sensing observation into locally responsible learning entities. Local evidence supports collective anomaly detection and spectral localization; communication is optional and follows agent formation.

2 Related Work

Multi-agent systems and multi-agent reinforcement learning have been applied to distributed sensing, robotics, control, surveillance, and autonomous decision making (Cai et al. (2021)). In most formulations, however, the population is predefined: agents correspond to robots, sensors, controllers, or other existing entities (Mishra et al. (2024),Hou et al. (2025)). Agentification addresses a preceding problem by constructing the computational population from the structure of the observation itself.

Graph neural networks represent relational structure through nodes and edges and propagate information through message passing (Hembert et al. (2024)). Agentification shares the premise that internal structure matters, but gives each computational entity persistent responsibility for a localized sensing subdomain and a local learned model. A graph (Dong et al. (2023)), when used, therefore describes interaction among already formed agents rather than defining agentification itself.

Mixture-of-experts and modular architectures distribute computation among specialized components, typically according to tasks, latent representations, examples, or learned routing (Pan et al. (2025)). In agentification, specialization is anchored to explicit sensing responsibility: an agent owns and reasons about a defined portion of the physical observation domain (Liu et al. (2025)).

Attention enables adaptive information interaction but does not inherently create persistent entities with explicit domain responsibility. Distributed sensor systems physically distribute sensing across devices (Bovenzi et al. (2018)); agentification can instead create many computational agents from one physical detector. Existing approaches therefore address representation, communication, modularity, or multi-agent reasoning individually (Gu & Alamaniotis (2022)). Agentification connects these ideas by making domain decomposition itself the first stage in constructing an agentic sensing system.

3 Agentification of a Sensing Domain

Let a structured sensing observation contain D elements:

$$\mathbf{x} = [x_1, x_2, \ldots, x_D]. \tag{1}$$

A conventional monolithic learner applies a global mapping

$$y = f(\mathbf{x}). \tag{2}$$

Agentification introduces a partition of the observation domain,

$$\mathcal{P} = \{P_1, P_2, \ldots, P_K\}, \qquad \bigcup_{k=1}^{K} P_k = \{1, 2, \ldots, D\}. \tag{3}$$

Partition P_k is assigned to agent a_k. The agent receives only its localized observation and produces local evidence,

$$e_k = f_k(\mathbf{x}_{P_k}). \tag{4}$$

Collective inference is produced by an aggregation operator G,

$$S = G(e_1, e_2, \ldots, e_K). \tag{5}$$

If communication is used, local evidence can be updated before aggregation,

$$e_k^{(t+1)} = \Phi\Big(e_k^{(t)}, \{e_j^{(t)} : j \in \mathcal{N}_k\}\Big). \tag{6}$$

This decomposition separates three design questions: what constitutes an agent, what each agent learns and infers, and whether/how agents communicate. Agentification is defined by the first two; explicit message passing is optional.

4 Spectral Agentification for Gamma-Ray Sensing

4.1 Spectral-Element Agents

A gamma-ray spectrum is $\mathbf{x} = [x_1, \ldots, x_{1024}]$, where x_i is the count in channel i. I adopt the finest fixed partition, $P_i = \{i\}$, yielding $K = D = 1024$. Thus one physical detector produces a population of 1,024 computational agents. This choice isolates the central agentification question without confounding it with a separate learned-boundary algorithm.

4.2 Local Gaussian-Process Learning

For channel agent i, background observations $B_i = \{x_i^{(1)}, \ldots, x_i^{(N)}\}$ define an empirical count distribution. For count q, the local negative log likelihood is

$$\mathcal{L}_i(q) = -\log[p_i(q) + \epsilon]. \tag{7}$$

A Gaussian process learns the mapping $q \mapsto \mathcal{L}_i(q)$ using a squared-exponential kernel Seeger (2004),

$$k(q, q') = \sigma_f^2 \exp\left[-\frac{(q - q')^2}{2\ell^2}\right]. \tag{8}$$

Each agent therefore learns a smooth, channel-specific model of how surprising different measurements are under background conditions.

4.3 Local Anomaly Evidence

For an incoming count, agent i produces GP response r_i. Background responses provide robust calibration through median m_i and scale $s_i = 1.4826\,\mathrm{MAD}_i$. I retain only excess surprise relative to normal background,

$$z_i = \max\left(\frac{r_i - m_i}{s_i}, 0\right), \tag{9}$$

$$e_i = 1 - \exp\left(-\frac{z_i^2}{2}\right), \qquad 0 \le e_i \le 1. \tag{10}$$

The quantity e_i is both a local anomaly belief and an attribution signal: because agent i corresponds exactly to channel i, high e_i identifies a spectral location contributing to the collective decision.

4.4 Collective Source-Agnostic Inference

A weak source may affect only part of the spectrum; averaging all 1,024 agents can dilute localized evidence. I therefore combine population-wide evidence with evidence from the strongest 10% of agents,

$$S_{\mathrm{global}} = \frac{1}{1024} \sum_{i=1}^{1024} e_i, \tag{11}$$

$$S_{\mathrm{tail}} = \frac{1}{|\mathcal{T}|} \sum_{i \in \mathcal{T}} e_i, \tag{12}$$

where $\mathcal{T}$ contains the top 10% of agents ranked by evidence. The final score is

$$S_{\mathrm{agent}} = 0.5 S_{\mathrm{global}} + 0.5 S_{\mathrm{tail}}. \tag{13}$$

No source signature appears in this detector. The inference path is therefore spectrum $\rightarrow$ local GP evaluations $\rightarrow$ local anomaly evidence $\rightarrow$ collective source-agnostic score.

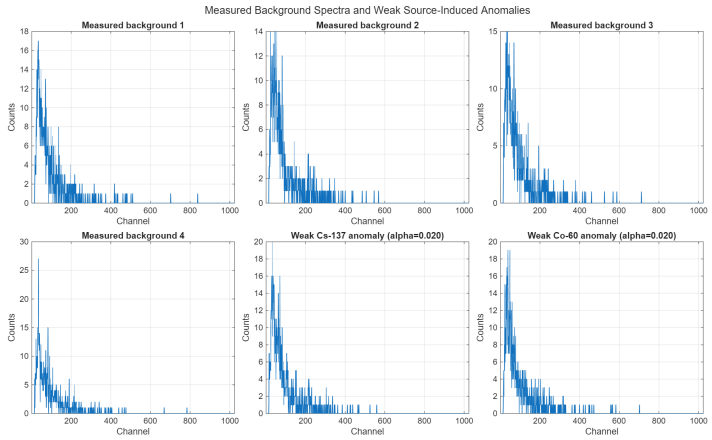

Figure 2: Four representative measured background spectra and two weak source-induced anomalous spectra (Cs-137 and Co-60, $\alpha = 0.02$). The weak anomalies visually illustrate the nuclear-security setting in which source-induced counts can be masked by natural background variability.

4.5 Optional Communication

For the communication ablation, each agent exchanges evidence with channels within radius two for three rounds, retaining 60% self-weight. This intentionally simple mechanism tests whether communication itself explains the performance of the agentified representation.

5 Experimental Methodology

Figure 2 shows representative measured backgrounds together with weak Cs-137 and Co-60 source-injected spectra at $\alpha = 0.02$, illustrating how source contributions can remain visually masked by background variability.

The dataset contains 100 measured background spectra, each with 1,024 channels (i.e., overall 102,400 values). Total counts range from 759 to 23,929. Ten repeated random 70/30 holdouts are used; each repetition trains on 70 background spectra and evaluates on 30 held-out spectra. Identical partitions and source realizations are used across competing methods, enabling paired comparison Ştefănescu et al. (2019).

Controlled anomalies are generated from normalized Cs-137 and Co-60 detector response functions (DRFs). The DRFs are used only for source injection and retrospective localization validation, never for inference. Let $\mathbf{d}_s$ denote the normalized DRF of source s and $\bar{B}$ the mean total background counts in the training set. Expected source contribution is

$$\boldsymbol{\lambda}_s = \alpha \bar{B} \frac{\mathbf{d}_s}{\sum_i d_{s,i}}, \tag{14}$$

$$\mathbf{y}_s \sim \mathrm{Poisson}(\boldsymbol{\lambda}_s), \tag{15}$$

$$\mathbf{x}_{\mathrm{anomaly}} = \mathbf{x}_{\mathrm{background}} + \mathbf{y}_s. \tag{16}$$

The measured held-out background is preserved and is not re-Poissonized. Source levels are $\alpha \in \{0.005, 0.01, 0.02, 0.05, 0.10, 0.25, 0.50\}$. Thus α directly denotes expected added source counts relative to the mean background total. The low-α regime is especially relevant to nuclear security because weak source contributions can be masked by stochastic background (Bandstra et al. (2023)).

Three monolithic baselines are evaluated. The PCA baseline standardizes the complete 1,024-dimensional spectrum and learns a background PCA representation using up to 20 components subject to a 95% explained-variance target. The residual baseline uses a single standardized global residual-energy statistic,

$$S_{\mathrm{residual}} = \sqrt{\frac{1}{D} \sum_{i=1}^{D} \left(\frac{x_i - \mu_i}{\sigma_i} \right)^2}. \tag{17}$$

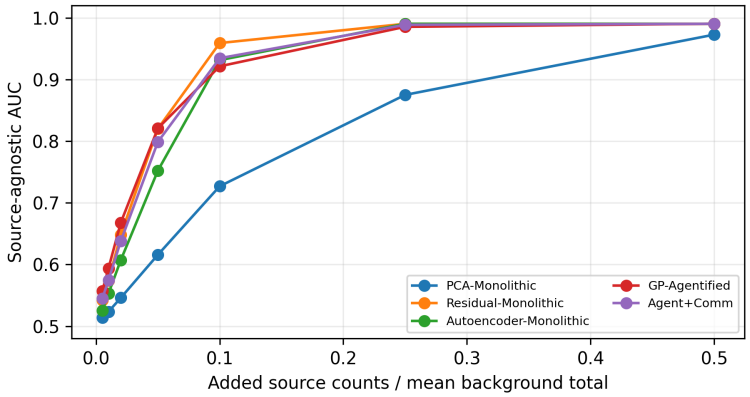

Figure 3: Source-agnostic anomaly detection for the three monolithic baselines and the agentified methods. The largest relative benefit of agentification occurs in the weak-source regime.

Table 1: Source-agnostic anomaly-detection AUC (mean ± SD over ten repeated holdouts).

α	PCA	Residual	Autoencoder	GP-Agent	Agent+Comm
0.005	.5141±.0113	.5421±.0125	.5253±.0091	**.5569±.0181**	.5449±.0123
0.010	.5234±.0149	.5740±.0170	.5528±.0194	**.5937±.0277**	.5743±.0216
0.020	.5460±.0253	.6484±.0351	.6073±.0299	**.6680±.0388**	.6382±.0411
0.050	.6158±.0540	.8203±.0608	.7523±.0548	**.8209±.0562**	.7989±.0697
0.100	.7269±.0818	**.9591±.0415**	.9314±.0478	.9214±.0414	.9345±.0504
0.250	.8750±.0753	**.9903±.0156**	.9903±.0156	.9854±.0221	.9887±.0184
0.500	.9727±.0308	.9903±.0156	.9903±.0156	**.9904±.0155**	.9903±.0156

A nonlinear autoencoder provides a modern machine-learning baseline. It is trained only on the 70 background training spectra in each split after channel-wise standardization and uses a compact $1024 \to 64 \to 8 \to 64 \to 1024$ architecture. Mean standardized reconstruction error is used as its anomaly score. The compact architecture limits model capacity given the small measured-background training set and preserves the same background-only anomaly-detection setting as the agentified method.

Detection is evaluated using ROC AUC for Cs-137 vs. background, Co-60 vs. background, and the primary source-agnostic task (Cs-137 or Co-60) vs. background. Localization is evaluated independently at the channel level. Ground truth is defined retrospectively by whether a channel received at least one injected source count. This ground truth is not available to the detector. I report channel-level localization AUC, precision among the highest-evidence agents, and the fraction of injected source-count mass captured by those agents.

6 Weak-Anomaly Detection Results

Figure 3 and Table 1 summarize source-agnostic anomaly detection across source strengths and all five methods. The principal result is concentrated in the weak-anomaly regime. At $\alpha = 0.005$, 0.010, and 0.020, GP agentification improves AUC over the strongest monolithic baseline by +0.0148, +0.0197, and +0.0196, respectively. Relative to the nonlinear autoencoder, the paired mean improvements are +0.0316, +0.0409, and +0.0607 at these same levels, and +0.0687 at $\alpha = 0.05$. At $\alpha = 0.05$ the GP agent and global residual detector are essentially tied; at $\alpha = 0.10$ the global residual representation is stronger, and at $\alpha \geq 0.25$ the methods approach saturation. Agentification should therefore not be interpreted as universally superior; its advantage is concentrated where anomalous evidence is weak enough to benefit from localized normality models.

Source morphology also matters. For Co-60, agentification yields 0.7344 vs. 0.6802 AUC at $\alpha = 0.02$ and 0.9168 vs. 0.8560 at $\alpha = 0.05$. For Cs-137, the residual baseline becomes stronger at moderate levels. This suggests that local responsibility is particularly useful when the anomalous structure remains locally distinctive rather than simply producing a large global deviation.

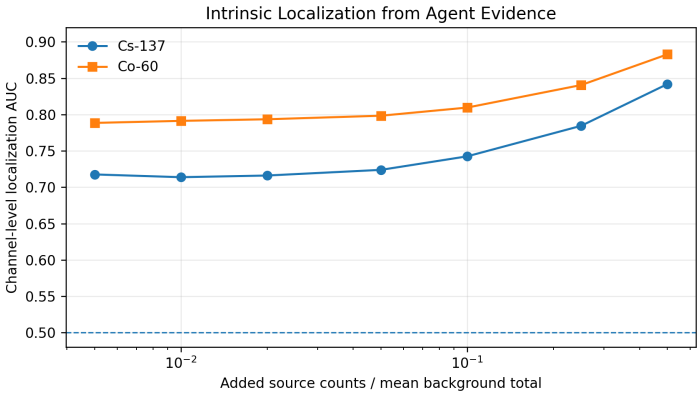

Figure 4: Channel-level localization AUC from GP-agent evidence. Localization is evaluated against injected source counts only after inference; DRFs are not used by the detector.

Table 2: Channel-level localization AUC. Values are means over ten repeated holdouts.

α	Cs-137	Co-60
0.005	0.7176	0.7885
0.010	0.7138	0.7913
0.020	0.7161	0.7935
0.050	0.7239	0.7984
0.100	0.7425	0.8096
0.250	0.7845	0.8406
0.500	0.8419	0.8828

7 Intrinsic Spectral Localization and Decision Attribution

Figure 4 and Table 2 quantify the channel-level localization retained by the agent evidence. A central consequence of agentification is that the collective decision retains a direct map back to the sensing domain. Because agent i owns channel i, the vector $\mathbf{E} = [e_1, \ldots, e_{1024}]$ is an intrinsic channel-level attribution map. No post-hoc explainer is required to establish which sensing elements contributed evidence.

Localization is already substantially above chance at the weakest source level: 0.7176 for Cs-137 and 0.7885 for Co-60 at $\alpha = 0.005$. The values remain approximately 0.71–0.80 through the low-source regime and increase to 0.8419 and 0.8828 at $\alpha = 0.50$. This indicates that even when collective detection is difficult, the distribution of local agent evidence contains information about where source-induced perturbations occur.

The strongest-agent analysis provides a complementary view. At the weakest levels, precise hard region extraction remains difficult because only a small number of injected counts exist and background fluctuations can also produce saturated local evidence. Consequently, I distinguish channel-level localization from exact hard-boundary explanation. At stronger source levels, the reported regions become increasingly aligned with actual injected counts. For example, representative Cs-137 explanations at $\alpha = 0.25$ identify channels 218–223 and 225–226 with substantial injected counts, while at $\alpha = 0.50$ the region 225–230 contains a large source contribution.

This motivates a careful interpretability claim: agentification does not guarantee perfect localization of every weak event, but it preserves localized evidence by construction. The same local beliefs that generate the collective decision can be inspected to determine which spectral regions support that decision. This is qualitatively different from a monolithic scalar detector whose output does not inherently retain responsibility at the sensing-element level.

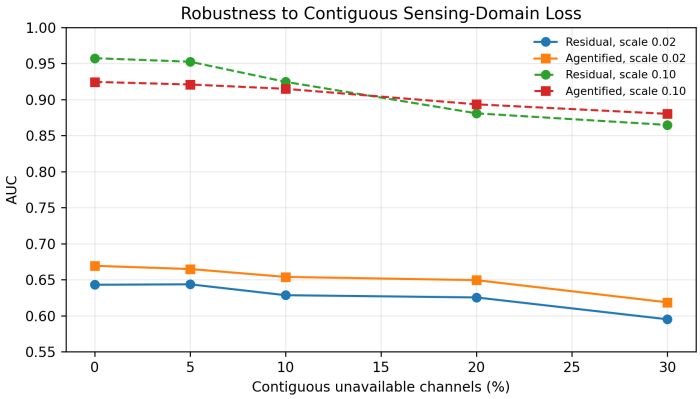

Figure 5: Robustness to contiguous sensing-domain loss. Agentification retains its weak-source advantage and can reverse the intact-domain ordering under substantial structured loss.

Table 3: Robustness summary for intact sensing and 30% channel loss.

α	Loss	Residual 0%	Agent 0%	Residual 30%	Agent 30%
0.010	Random	0.5782	0.5964	0.5682	0.5832
0.020	Random	0.6429	0.6692	0.6304	0.6531
0.050	Random	0.8257	0.8222	0.7786	0.7929
0.100	Random	0.9572	0.9243	0.9342	0.9043
0.010	Block	0.5782	0.5964	0.5605	0.5705
0.020	Block	0.6429	0.6692	0.5950	0.6185
0.050	Block	0.8257	0.8222	0.7588	0.7843
0.100	Block	0.9572	0.9243	0.8646	0.8800

8 Robustness to Partial Sensing-Domain Failure

Figure 5 and Table 3 summarize performance under random and contiguous sensing-domain loss. I next test whether distributed local responsibility changes how inference degrades when portions of the sensing domain become unavailable. Random channel loss and contiguous block loss are evaluated at 0%, 5%, 10%, 20%, and 30%. Failed agents abstain; the monolithic residual baseline is likewise computed only over available channels.

At $\alpha = 0.02$, agentification retains an advantage at every tested random and block-loss level. More revealingly, at $\alpha = 0.10$ the residual baseline is clearly stronger with intact sensing (0.9572 vs. 0.9243), but the ordering reverses under 20% contiguous loss (0.8807 vs. 0.8932) and 30% contiguous loss (0.8646 vs. 0.8800). The result suggests that global representations can be advantageous when the complete observation is available, while distributed local responsibility becomes comparatively valuable when observation integrity deteriorates.

9 Communication Ablation

The Agent+Comm results in Table 1 isolate the effect of explicit message passing. Explicit adjacent-agent communication is not consistently beneficial. In the source-agnostic task, Agent+Communication minus Agentification equals -0.0120, -0.0194, -0.0298, and -0.0220 at $\alpha = 0.005$, 0.010, 0.020, and 0.050, respectively. A small positive difference appears at $\alpha = 0.10$, while stronger regimes are essentially saturated. The simplest explanation is that indiscriminate local averaging can dilute the very weak, localized evidence that agentification is designed to preserve.

This negative result is conceptually useful: agentification and communication should not be conflated. Agentification establishes localized computational responsibility. Communication is a separate design choice whose topology and content must preserve task-relevant information. Future work can investigate confidence-aware, learned, or attention-based communication rather than assuming that message passing is intrinsically beneficial.

10 DISCUSSION

The weakest source levels are not merely a difficult benchmark selected to separate algorithms; they represent the nuclear-security regime motivating this work, where source-induced counts can be sparse relative to background and masked by stochastic fluctuations. A monolithic representation asks whether the complete spectrum is sufficiently unusual, whereas agentification first asks 1,024 localized questions about channel-specific normality and then forms a collective conclusion. The low-α advantage in Figure 3 and Table 1 is therefore aligned with the physical structure of the problem rather than indicating universal superiority.

The spectral agents are intentionally simple and should not be interpreted as general-purpose autonomous AI systems. Their intelligence is localized and functional: each owns a sensing subdomain, learns its expected behavior, independently evaluates novelty, and contributes local evidence to collective inference. Persistent localized responsibility and local inference—not model complexity—distinguish agentification from feature partitioning. This organization also produces detection and attribution from the same representation: aggregate evidence indicates whether a spectrum is anomalous, while local evidence identifies where support originates. As Figure 4 and Table 2 show, channel-level localization remains informative even for weak events, although exact contiguous-region explanations are noisy at the weakest levels.

Within Agentic Nuclear Sensing Systems, agentification converts passive measurements into computational entities participating in interpretation. The present implementation addresses the sensing and reasoning side of the broader sense–reason–act loop and does not claim mission-level autonomy. The principle is also substrate independent: images, time series, power-system measurements, and multimodal observations could be transformed into spatial, temporal, electrical-region, or hierarchical agent populations.

The main limitations are the 100 measured background spectra and controlled Poisson source additions rather than a large independent source-present campaign. Repeated holdouts reduce dependence on one partition but do not replace additional data. The current one-channel-per-agent partition is fixed rather than learned, communication is intentionally simple, and localization is strongest as a channel-ranking signal rather than an exact weak-source boundary estimator. These limitations motivate independent source-present validation, learned agent boundaries, and task-aware communication.

11 CONCLUSION

I introduced agentification, a learning paradigm in which a structured sensing domain is transformed into locally responsible computational agents. In gamma-ray spectroscopy, 1,024 GP agents independently learn channel-level background behavior and contribute evidence to collective inference. Agentification improves source-agnostic detection in the weakest tested regimes, provides intrinsic channel-level localization, and exhibits favorable behavior under partial sensing-domain degradation. Conversely, strong intact-domain anomalies can favor a global residual representation, and simple neighboring-agent communication is not inherently beneficial.

The value of agentification therefore lies not in universally replacing monolithic learning, but in reorganizing responsibility across structured observations. For nuclear security, this is particularly relevant when weak source-induced counts must be distinguished from a stronger fluctuating background. More broadly, agentification provides a computational foundation for sensing architectures in which local learning, intrinsic attribution, and collective inference are built into the organization of the observation domain, including Agentic Nuclear Sensing Systems.

AI USE STATEMENT

Generative AI tools were used to assist with language editing, manuscript organization, and code development/debugging. The author reviewed the AI-assisted material, verified the experimental code and numerical results, and takes responsibility for the final content, claims, and artifacts in this work.

Reproducibility statement

The manuscript specifies the data organization, repeated-holdout protocol, source-injection procedure, GP-agent construction, anomaly-evidence computation, collective aggregation, monolithic baselines, localization evaluation, communication ablation, and sensing-degradation experiments. The corresponding implementation and additional experimental outputs can be provided as supplementary material to support reproduction of the reported results.

References

Deepak Bhaskar Acharya, Karthigeyan Kuppan, and B Divya. Agentic ai: Autonomous intelligence for complex goals—a comprehensive survey. *IEEE Access*, 13:18912–18936, 2025.

Miltiadis Alamaniotis. Preventing a nuclear september 11th: Challenges, solutions, and concerns in ai-empowered analysis software in sensors. In *International Conference on Artificial Intelligence-empowered Software Engineering*, pp. 150–160. Springer, 2024.

MS Bandstra, N Abgrall, RJ Cooper, D Hellfeld, THY Joshi, V Negut, BJ Quiter, M Salathe, R Sankaran, Y Kim, et al. Background and anomaly learning methods for static gamma-ray detectors. *IEEE Transactions on Nuclear Science*, 70(10):2352–2363, 2023.

Giampaolo Bovenzi, Domenico Ciuonzo, Valerio Persico, Antonio Pescapè, and Pierluigi Salvo Rossi. Iot-enabled distributed detection of a nuclear radioactive source via generalized score tests. In *International symposium on signal processing and intelligent recognition systems*, pp. 77–91. Springer, 2018.

Ting Cai, Zhihua Yang, Yufei Chen, Wuhui Chen, Zibin Zheng, Yang Yu, and Hong-Ning Dai. Cooperative data sensing and computation offloading in uav-assisted crowdsensing with multi-agent deep reinforcement learning. *IEEE Transactions on Network Science and Engineering*, 9 (5):3197–3211, 2021.

Guimin Dong, Mingyue Tang, Zhiyuan Wang, Jiechao Gao, Sikun Guo, Lihua Cai, Robert Gutierrez, Bradford Campbel, Laura E Barnes, and Mehdi Boukhechba. Graph neural networks in iot: A survey. *ACM Transactions on Sensor Networks*, 19(2):1–50, 2023.

James M Ghawaly Jr, Andrew D Nicholson, Daniel E Archer, Michael J Willis, Irakli Garishvili, Brandon Longmire, Andrew J Rowe, Ian R Stewart, and Matthew T Cook. Characterization of the autoencoder radiation anomaly detection (arad) model. *Engineering Applications of Artificial Intelligence*, 111:104761, 2022.

Siyao Gu and Miltiadis Alamaniotis. Radiation sensor placement using reinforcement learning in nuclear security applications. In *2022 13th International Conference on Information, Intelligence, Systems & Applications (IISA)*, pp. 1–6. IEEE, 2022.

Pierre Hembert, Chady Ghnatios, Julien Cotton, and Francisco Chinesta. Assessing sensor integrity for nuclear waste monitoring using graph neural networks. *Sensors*, 24(5):1580, 2024.

Xiangwang Hou, Jingjing Wang, Jun Du, Chunxiao Jiang, and Yong Ren. Distributed machine learning for autonomous agent swarm: A survey. *IEEE Communications Surveys & Tutorials*, 2025.

Taishi Ito, Yasuhiro Kurita, Hideyuki Takahashi, Kazuto Sasai, Gen Kitagata, and Tetsuo Kinoshita. A cooperative multi-agent system infrastructure for heterogeneous sensor environment. *Int. J. of Energy, Information and Communications*, 4(2):55–70, 2013.

Patrik Kamencay, Peter Hockicko, and Robert Hudec. Sensors data processing using machine learning, 2024.

Jaewon Lee, Tenzing H Joshi, Mark S Bandstra, Donald L Gunter, Brian J Quiter, Reynold J Cooper, and Kai Vetter. Radiation image reconstruction and uncertainty quantification using a gaussian process prior. *Scientific Reports*, 14(1):22958, 2024.

Yimei Liu, Changyuan Wu, Minglei Guan, and Jingzhe Wang. Mixturers: A mixture of expert network based remote sensing land classification. *Remote Sensing*, 17(14):2494, 2025.

Manav Mishra, Prithvi Poddar, Rajat Agrawal, Jingxi Chen, Pratap Tokekar, and PB Sujit. Multi-agent deep reinforcement learning for persistent monitoring with sensing, communication, and localization constraints. *IEEE Transactions on Automation Science and Engineering*, 22:2831–2843, 2024.

Licheng Pan, Hao Wang, Zhichao Chen, Yuxin Huang, Yunlong Niu, Zhaoran Liu, Qunshan He, and Xinggao Liu. Controllable mixture-of-experts for multivariate soft sensors. *IEEE Transactions on Automation Science and Engineering*, 2025.

Matthias Seeger. Gaussian processes for machine learning. *International journal of neural systems*, 14(02):69–106, 2004.

Răzvan Ştefănescu, Jason Hite, Jared Cook, Ralph C Smith, and John Mattingly. Surrogate-based robust design for a non-smooth radiation source detection problem. *Algorithms*, 12(6):113, 2019.

Lefteri Tsoukalas. Intelligent component monitoring for nuclear power plants. Technical report, Purdue University, 2010.

Changyuan Zhao, Ruichen Zhang, Jiacheng Wang, Dusit Niyato, Geng Sun, Xianbin Wang, Shiwen Mao, and Abbas Jamalipour. From agentification to self-evolving agentic ai for wireless networks: Concepts, approaches, and future research directions. *IEEE Communications Magazine*, 2026.