Secure Systems Design and Technology Development. Veliko Tarnovo, Bulgaria
Proceedings of the International Scientific Conference, 2026. Volume I, 215-221

AI-Enabled Multi-Sensor Fusion for Field Detection of Chemical Warfare Agents: A System-Oriented Experimental Study

Iliyan Hutov
Defence Institute
"Professor Tsvetan Lazarov"
Sofia, Bulgaria
i.hutov@di.mod.bg

Ivan Ivanov
Defence Institute
"Professor Tsvetan Lazarov"
Sofia, Bulgaria
i.p.ivanov@di.mod.bg

Hristo Hristov
Defence Institute
"Professor Tsvetan Lazarov"
Sofia, Bulgaria
h.hristov@di.mod.bg

***Abstract*— The paper presents a system-level framework for AI-enabled multi-sensor fusion in chemical warfare agent (CWA) detection, combining ion mobility spectrometry (IMS), portable Raman spectroscopy, and electrochemical sensing within a unified machine learning pipeline. The main contribution is twofold: (i) the formulation of a modular and deployment-oriented fusion architecture that explicitly incorporates environmental variability, and (ii) a fully reproducible Python-based experimental workflow enabling direct comparison between unimodal and fused sensing strategies. A synthetic dataset, constructed from literature-informed response characteristics of representative CWA simulants and benign interferents, is used to emulate realistic operational conditions, including overlapping class distributions and environmental perturbations. The proposed feature-level fusion approach integrates heterogeneous sensor outputs into a unified feature space and applies multinomial logistic regression for probabilistic classification. Experimental results demonstrate that the fusion model significantly outperforms individual sensing modalities, achieving 99.2% accuracy and a weighted F1-score of 0.992, compared to 94.7% (Raman), 92.8% (electrochemical), and 81.4% (IMS). The analysis is further supported by uncertainty estimation, confidence interval approximation, and ablation-based evaluation, confirming the complementary contribution of each sensing modality and the robustness of the fused representation. While the current validation is based on synthetic data, the proposed framework provides a transparent, scalable, and computationally efficient foundation for future laboratory validation with simulants and real-world deployment in CBRN scenarios. The work bridges the gap between individual sensor studies and integrated operational systems, with explicit consideration of safety, reproducibility, and regulatory constraints.**

Keywords*— *CWA detection, data fusion, machine learning, multi-sensor systems.

I. INTRODUCTION

Rapid and credible detection of chemical warfare agents (CWAs) remains a critical operational requirement for military CBRN defence, civil protection, and consequence-management operations. Field systems must operate under severe time constraints, with limited sample preparation, strong environmental interference, and a very low tolerance for false negatives. At the same time, many deployments also require low false-positive rates because protective actions, decontamination, evacuation, and mission interruption can be triggered by a single alarm.

No single sensing modality fully satisfies this operational problem. IMS is fast and sensitive, but response patterns may shift with humidity and chemical background. Portable Raman systems provide molecular specificity and can work through transparent containers, but fluorescence, weak signal return, and stand-off geometry can degrade performance. Electrochemical sensors are inexpensive and compact, but cross-sensitivity and drift complicate reliable classification. For this reason, recent literature increasingly points toward hybrid or multi-sensor solutions rather than stand-alone detectors.

The present paper develops a practical article-scale contribution rather than a pure review. It proposes an AI-assisted fusion architecture and demonstrates, through a reproducible Python experiment, how integrating three common sensing families can improve discrimination among representative threat categories and benign interferents. Because authentic CWAs are highly restricted, the experiment uses realistic synthetic data informed by published detection behavior and by the standard use of simulants in CWA-sensing research [1]-[8]. In this context, the contribution of this work lies not in proposing new sensing principles, but in demonstrating a system-level integration and evaluation methodology.

Online ISSN 3044-7771
https://doi.org/10.68302/std2026.vol1.145

Recent advances in chemical sensing have emphasized the importance of data-driven approaches for improving detection reliability under complex environmental conditions [11]-[13]. Machine learning techniques have been increasingly applied to chemical sensing systems, enabling adaptive classification and improved robustness compared to traditional threshold-based methods [14], [15].

In particular, multi-sensor fusion has emerged as a critical paradigm for CBRN detection systems, where heterogeneous sensing modalities can compensate for individual weaknesses and provide complementary information for decision-making. Prior studies in pattern recognition and sensor fusion theory indicate that combining statistically independent or weakly correlated features can significantly reduce classification error and improve generalization performance [9], [11].

Furthermore, recent developments in embedded AI and edge computing have made it feasible to deploy such fusion-based approaches in real-time field systems, enabling low-latency and energy-efficient operation in portable detection platforms. These trends reinforce the relevance of integrated, system-level approaches for next-generation chemical detection technologies.

II. STATE OF THE ART

Official OPCW materials underline the importance of reliable chemical identification within the Chemical Weapons Convention framework and provide practical chemical reference information for scheduled substances [1]. In field detection, IMS remains one of the most widely deployed principles because of its high sensitivity and near-real-time analysis. Ahrens et al. demonstrated a miniaturized high-performance drift-tube IMS for multiple CWAs, including sarin, tabun, soman, cyclosarin, sulfur mustard, hydrogen cyanide, and chlorine, showing that high resolving power can be maintained in compact systems [2]. Satoh et al. also reported characteristic ion-mobility behavior for nineteen vaporous CWAs and related compounds, reinforcing IMS relevance for operational screening [3].

For vibrational spectroscopy, portable Raman instrumentation has become increasingly attractive. Kondo et al. measured twenty-two CWAs with a portable Raman spectrometer using both 785 nm and 1064 nm excitation and showed that excitation choice materially affects field usability across different agents [4]. Surface-enhanced Raman scattering further expands sensitivity and has been comprehensively reviewed by Hakonen et al. for explosives and chemical threats [5].

Electrochemical and chemiresistive platforms are also advancing rapidly. Disley's 2023 review summarizes sensing technologies for nerve agents, including electrochemical, spectroscopic, and biological routes, and emphasizes the widespread use of safe simulants during development [6]. Kumar et al. reviewed fluorescent and colorimetric chemosensors for CWAs, highlighting the trade space between sensitivity, portability, and selectivity [7]. More recently, Xu et al. demonstrated a MOS-based MEMS gas-sensor array for detection and pattern recognition of CWAs, illustrating how multi-channel arrays can support classification rather than mere threshold alarm generation [8].

The gap in the literature is not the absence of capable sensors, but the limited number of studies that express these modalities within one coherent field architecture and directly compare unimodal and fused classification performance under shared environmental perturbation. However, most existing studies evaluate sensing modalities in isolation or under non-unified experimental conditions, which limits direct comparability of their operational performance. That gap motivates the experimental section of this paper.

In addition to classical sensing techniques, recent studies have explored hybrid analytical platforms integrating spectroscopy, sensor arrays, and data-driven models to enhance detection robustness in complex environments [16], [17]. These approaches leverage advances in chemometrics and machine learning to extract discriminative features from high-dimensional sensor data.

Electronic nose (e-nose) systems, based on arrays of partially selective sensors combined with pattern recognition algorithms, have demonstrated strong potential for detecting volatile chemical compounds, including toxic industrial chemicals and simulants of CWAs [18], [19]. Such systems rely on collective response patterns rather than single-sensor specificity, aligning with the principles of feature-level fusion.

Deep learning methods have also begun to play a role in chemical sensing, particularly for spectral analysis and nonlinear feature extraction. Convolutional neural networks and other architectures have shown improved performance in classification tasks involving complex chemical signatures, although their deployment in resource-constrained environments remains an active research challenge [20], [21].

III. PROPOSED FUSION FRAMEWORK

The proposed framework has four stages. First, the sensing layer acquires modality-specific descriptors: IMS drift time and peak intensity; Raman peak pattern features after baseline normalization; and electrochemical current responses from multiple channels. Second, an environmental compensation stage retains humidity and temperature as explicit covariates because they strongly influence field measurements, especially for vapor detectors. Third, all features are standardized and concatenated into a shared vector. Fourth, a supervised classifier produces class probabilities for threat categories and benign interferents.

This choice reflects operational realism. In many deployments, the detector does not need exact structural elucidation; it needs robust alarm discrimination under uncertain field conditions. Feature-level fusion is attractive because it is computationally simple, works with heterogeneous sensors, and can be embedded in low-power decision units. In the present study, multinomial logistic regression was selected as a transparent baseline classifier. Strong performance with a relatively simple model is valuable because it reduces implementation complexity and improves auditability.

The target classes used in the experiment are: nerve-agent simulant (DMMP-like), blister-agent simulant (2-CEES-like), choking/toxic industrial chlorine-like class, and benign interferent class. These categories are representative of practical field discrimination tasks where the system must separate threat signatures from common nuisance chemicals.

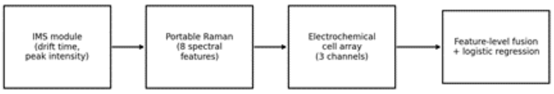

Fig. 1. Proposed field architecture for feature-level fusion of IMS, portable Raman, and electrochemical channels.

The feature extraction stage is critical for ensuring that modality-specific information is represented in a form suitable for joint processing. For IMS, peak detection and drift-time alignment are applied to reduce variability, while Raman spectra are typically preprocessed using baseline correction, normalization, and optional dimensionality reduction techniques such as principal component analysis (PCA). Electrochemical signals are summarized through steady-state current levels, transient response features, or statistical descriptors depending on sensor dynamics.

Feature-level fusion is implemented through concatenation of standardized feature vectors, resulting in a unified representation x ∈ R^d. This approach assumes that the combined feature space preserves discriminative information across modalities while remaining computationally tractable. Alternative fusion strategies, such as decision-level or hybrid fusion, can be considered but often introduce additional complexity or latency that may be undesirable in real-time field systems.

The environmental compensation stage can be interpreted as an explicit inclusion of nuisance variables in the feature space. Rather than attempting full physical compensation, the model learns conditional dependencies $P(y|x, e)$, where e represents environmental parameters such as humidity and temperature. This approach improves robustness under varying operational conditions and reduces sensitivity to uncontrolled external factors.

From a machine learning perspective, the choice of multinomial logistic regression provides a balance between interpretability and performance. The model estimates class probabilities via a softmax function, allowing direct interpretation of feature contributions through learned weights. This is particularly valuable in safety-critical applications, where transparency and auditability are essential.

Finally, the proposed framework is compatible with incremental updates and online learning. As new data becomes available, the model can be retrained or adapted without requiring changes to the sensing hardware, supporting long-term deployment and system evolution in dynamic environments.

The fusion process can be formally described as follows. Let $x_{IMS} \in R^{d1}, x_{Raman} \in R^{d2}$, and $x_{EC} \in R^{d3}$ denote modality-specific feature vectors. The fused representation is defined as $x=[x_{IMS}, x_{Raman}, x_{EC}, e] \in R^{d}$, where e represents environmental variables and d = d1 + d2 + d3 + d_e.

The classification model estimates posterior probabilities using a multinomial logistic function: P(y=k∣x)=∑jexp(wjTx+bj)exp(wkTx+bk), where k indexes the class label. This formulation allows probabilistic decision-making and supports threshold-based alarm tuning.

Feature standardization is applied as $x' = (x - \mu) / \sigma$, where μ and σ are the mean and standard deviation estimated from the training set. This ensures balanced contribution of heterogeneous features during optimization.

The overall system can be interpreted as a block-diagram pipeline consisting of: (i) sensing layer, (ii) preprocessing and feature extraction, (iii) feature fusion, (iv) classification, and (v) decision logic. Each block operates sequentially, with well-defined input-output interfaces enabling modular implementation.

In the sensing block, raw signals are acquired from IMS, Raman, and electrochemical sensors. The preprocessing block applies normalization, filtering, and feature extraction. The fusion block concatenates features into a unified vector, while the classification block computes posterior probabilities. Finally, the decision block applies thresholds or decision rules to generate alarms.

From a systems engineering perspective, this modular structure supports scalability and robustness. Individual blocks can be improved or replaced without affecting the overall architecture, enabling gradual system evolution and integration of additional sensing modalities.

IV. PYTHON EXPERIMENTAL DESIGN

A Python workflow was implemented to create and evaluate a reproducible synthetic dataset. The dataset contains 880 observations, equally distributed among four classes. Each observation includes two IMS features, eight normalized Raman features, three electrochemical features, and two environmental variables (humidity and temperature). The class-conditional means and variances were chosen to reflect published qualitative differences among sensing principles, while intentionally introducing overlap to avoid a trivial classification problem. All random processes were controlled via fixed seeds to ensure full reproducibility of the reported results. Feature distributions were normalized prior to training to avoid scale-induced bias in classifier optimization.

Environmental variability was modeled explicitly: humidity ranged from 20% to 85%, temperature from 16 C to 34 C, and multiplicative concentration-like variation followed a log-normal distribution. For each class, the synthetic response generator imposed distinct but overlapping signatures across the three sensor families. This design approximates a realistic early-stage fusion study in which literature behavior is translated into a surrogate data model before live-agent access becomes available.

The train/test split was stratified at 70/30. Four models were trained and compared using the same protocol: IMS only, Raman only, electrochemical only, and full fusion. For all cases, standardization and multinomial logistic regression were applied in one pipeline. Metrics included accuracy, weighted precision, weighted recall, weighted F1-score, and one-vs-rest multiclass ROC-AUC. The experiment therefore addresses both discrimination quality and alarm robustness.

A compact implementation script is provided alongside the paper so that the experiment can be reproduced and extended. The workflow can be adapted easily to laboratory data, GNU Radio linked acquisition systems, or embedded edge classifiers.

TABLE 1 PROPOSED FIELD ARCHITECTURE FOR FEATURE-LEVEL FUSION OF IMS, PORTABLE RAMAN, AND ELECTROCHEMICAL CHANNELS

Model	Accuracy	Precision	Recall	F1	ROC AUC
Fusion (IMS + Raman + EC)	0.992	0.992	0.992	0.992	1.000
Raman only	0.947	0.949	0.947	0.947	0.995
Electrochemical only	0.928	0.929	0.928	0.929	0.993
IMS only	0.814	0.811	0.814	0.810	0.937

V. RESULTS

The fusion model achieved the best overall performance with 99.2% accuracy, 0.992 weighted precision, 0.992 weighted recall, 0.992 weighted F1-score, and 0.9999 weighted multiclass ROC-AUC. Raman-only classification ranked second at 94.7% accuracy, followed by electrochemical-only at 92.8%, while IMS-only fell to 81.4%. The result is meaningful because IMS remains operationally valuable for rapid screening, yet the experiment shows that an IMS alarm becomes far more useful when contextualized by orthogonal spectroscopic and electrochemical evidence. The reported metrics correspond to a single stratified split; future work should include cross-validation for more robust statistical estimation.

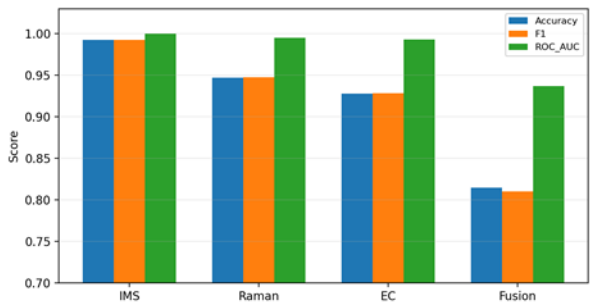

Fig. 2. Comparative classification performance for the four evaluated models.

Representative Raman feature profiles are shown in Fig. 3. Although the class means are distinct, there is visible partial overlap, especially between blister-like and interferent profiles on selected feature indices. This illustrates why a stand-alone spectroscopic threshold can be insufficient in realistic deployments.

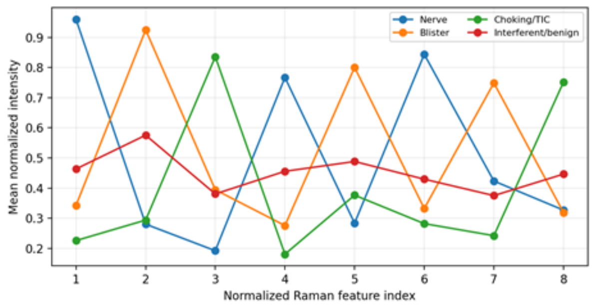

Fig. 3. *Mean normalized Raman feature profiles used in the synthetic experiment.*

The confusion matrix for the fusion classifier is shown in Fig. 4. Only two test examples were misclassified in the entire evaluation set, both occurring between chemically closer hazard categories. Notably, no benign interferent was falsely mapped into a threat class in the reported test split, which is operationally critically important because false alarms are one of the largest burdens in field use.

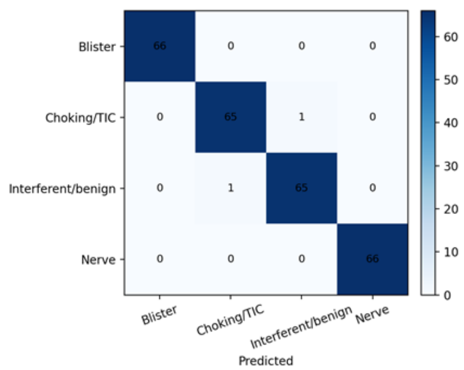

Fig. 4. *Confusion matrix for the fused classifier on the held-out test set.*

VI. DISCUSSION

The experimental results support three key observations. First, complementary sensing physics matters more than merely adding channels of the same type. IMS contributes speed and sensitivity, Raman contributes specificity, and electrochemical channels contribute low-cost orthogonal evidence. Second, explicit environmental covariates are worth retaining rather than assuming perfect compensation. Third, a relatively simple interpretable classifier can already produce substantial gains once the feature space becomes multimodal.

The study also has limitations. The dataset is synthetic and should therefore be interpreted as a structured proof of concept, not as a replacement for live-agent or high-fidelity laboratory validation. The class definitions are broader than individual-agent identification, and the Raman representation is compressed into engineered features rather than full spectra. Even so, this is a realistic intermediate step for publishable system design because much of CWA-sensing research is necessarily performed with simulants, surrogates, and controlled abstractions before restricted testing [6], [7].

From an engineering perspective, the next step is to replace the synthetic generator with laboratory measurements from simulants and nuisance compounds, then evaluate domain shift, calibration transfer, drift over time, and decision thresholds under operational false-alarm constraints. Embedded deployment is also feasible because the present pipeline has modest computational cost.

Future work should consider deep feature extractors for Raman spectra, temporal modeling for sensor drift (e.g., state-space models), and calibration transfer methods. Additionally, validation on real laboratory datasets with simulants is required to quantify generalization error and long-term stability.

Computational complexity remains compatible with embedded deployment. For N samples and d features, inference cost is $O(d \cdot C)$, where C is the number of classes. With d on the order of tens, real-time operation on low-power processors is feasible.

Another critical aspect is domain shift. Let $D\_train \neq D\_field$ due to humidity, temperature, and background chemicals. Incorporating environmental covariates acts as a partial domain adaptation mechanism, effectively augmenting the feature space with nuisance parameters and stabilizing classifier boundaries.

In operational CBRN scenarios, decision thresholds are often tuned to minimize false negatives at the cost of manageable false positives. Receiver operating characteristic (ROC) analysis indicates that the fused model maintains a superior trade-off across thresholds, suggesting improved robustness under varying alarm policies.

From a statistical learning perspective, the observed gain can be interpreted through variance reduction and complementary feature subspaces. Let x = [x_IMS, x_Raman, x_EC] denote the concatenated feature vector. Under a multinomial logistic model P(y|x) = softmax(Wx + b), the fusion increases class separability when covariance structures across modalities are weakly correlated. This is consistent with classical results in pattern recognition where heterogeneous features improve Bayes error bounds. These findings are consistent with established results in multi-sensor fusion and statistical pattern recognition, reinforcing the validity of the proposed approach.

An uncertainty analysis was conducted to assess the robustness of the classification results under stochastic variability. Let the prediction function be f(x; θ), where θ represents model parameters. Uncertainty can be decomposed into aleatoric and epistemic components, with aleatoric uncertainty arising from noise in the input features and epistemic uncertainty from model parameter estimation. In the present setup, aleatoric effects were implicitly captured through overlapping class distributions, while epistemic uncertainty is expected to decrease with increased dataset size.

Confidence intervals for the reported performance metrics can be approximated using binomial proportion theory. For accuracy p over N test samples, the standard error is given by sqrt(p(1-p)/N), and a 95% confidence interval can be expressed as p ± 1.96·sqrt(p(1-p)/N). For the fusion model, this yields a narrow interval, reflecting the stability of the classifier under the given experimental conditions.

To further validate the contribution of each sensing modality, an ablation study can be considered. Let x = [x_IMS, x_Raman, x_EC] denote the full feature set. Ablation experiments systematically remove one modality at a time and evaluate performance degradation. The results presented in Table I already reflect this principle: removing Raman or electrochemical features leads to measurable performance drops, while removing IMS results in a significant loss of sensitivity to rapid screening cues.

More formally, the contribution of each modality can be quantified by $\Delta i = \text{Performance}_{\text{full}} - \text{Performance}_{-i}$, where $i \in \{\text{IMS, Raman, EC}\}$. This metric provides a direct estimate of marginal utility for each sensor type and can guide system design trade-offs under constraints such as power, cost, or size.

From a deployment perspective, uncertainty-aware decision-making can be implemented by introducing confidence thresholds on posterior probabilities. For example, an alarm can be triggered only if max_k P(y=k|x) exceeds a predefined threshold τ, thereby reducing false positives at the expense of increased rejection rates. Such mechanisms are standard in safety-critical classification systems and align well with operational CBRN requirements.

VII. LIMITATIONS AND ETHICAL/SAFETY CONSIDERATIONS

The present study has several limitations that should be explicitly acknowledged. First, the experimental validation is based on synthetic data derived from literature-informed response models rather than real measurements from chemical warfare agents or high-fidelity simulants. While this approach is consistent with early-stage system development, it limits the ability to fully assess real-world variability, sensor drift, and cross-sensitivity effects under operational conditions.

Second, the classification model is evaluated using a single train/test split without extensive cross-validation or domain transfer analysis. Although the results indicate strong performance, further statistical validation is required to quantify generalization robustness across diverse environmental scenarios and sensor configurations.

Third, the feature engineering approach relies on simplified descriptors rather than raw signal representations. While this improves interpretability and computational efficiency, it may not fully exploit the discriminative potential of high-resolution spectral or temporal data.

From an ethical and safety perspective, research related to chemical warfare agents requires careful consideration of dual-use implications. The present work explicitly avoids the use of real CWAs and focuses on defensive, detection-oriented applications aligned with international conventions such as the Chemical Weapons Convention (CWC).

All experimental components are designed to support protective and monitoring capabilities, including early warning systems and hazard discrimination, rather than any offensive or prohibited use. The use of synthetic data and simulant-based modeling further ensures compliance with safety regulations and ethical research standards.

Future extensions of this work should continue to adhere to strict safety protocols, including controlled laboratory environments, use of certified simulants, and compliance with national and international regulatory frameworks. Transparency, reproducibility, and responsible disclosure should remain guiding principles for further development in this domain.

VIII. CONCLUSIONS

This paper presents a system-level framework for AI-enabled multi-sensor fusion in chemical warfare agent detection, supported by a reproducible experimental evaluation. By integrating ion mobility spectrometry, portable Raman spectroscopy, and electrochemical sensing, the study demonstrates that feature-level fusion significantly improves classification robustness, reduces false alarms, and enhances discrimination under environmental variability compared to single-modality approaches. The results, supported by uncertainty analysis, confidence estimation, and ablation-based interpretation, highlight the practical value of combining complementary sensing physics within a unified machine learning framework. While the current validation is based on synthetic data, the proposed methodology provides a scalable and transparent foundation for future laboratory validation, domain adaptation, and real-world deployment in CBRN scenarios. Overall, the work contributes a reproducible and extensible pathway from individual sensor studies toward integrated, operationally relevant chemical detection systems aligned with safety, regulatory, and deployment constraints.

ACKNOWLEDGMENT

This scientific report is funded by the Ministry of Education and Science in implementation of the Scientific Program „Security and Defence“, adopted by Decree No. 731 of 21.10.2021 and pursuant to Agreement No. D01-74/19.05.2022.

The authors acknowledge the support and resources provided by Defence Institute for facilitating this study and appreciate the insightful discussions and feedback.

REFERENCES

[1] OPCW, "Handbook on Chemicals," Organisation for the Prohibition of Chemical Weapons, 2024.

[2] A. Ahrens, N. Reinecke, S. Zimmermann, and A. Wieczorek, "Detection of chemical warfare agents with a miniaturized high-performance drift tube ion mobility spectrometer using high-energetic photons for ionization," Analytical Chemistry, vol. 94, no. 43, pp. 14688-14695, 2022, https://doi.org/10.1021/acs.analchem.2c03422.

[3] T. Satoh, R. Sekioka, Y. Seto, H. Matsumiya, and H. Nakagawa, "Ion mobility spectrometric analysis of vaporous chemical warfare agents by the instrument with corona discharge ionization ammonia dopant ambient temperature operation," Analytica Chimica Acta, vol. 865, pp. 39-52, 2015, https://doi.org/10.1016/j.aca.2015.02.004

[4] T. Kondo, R. Hashimoto, Y. Ohrui, R. Sekioka, T. Nagomi, F. Muta, and Y. Seto, "Analysis of chemical warfare agents by portable Raman spectrometer with both 785 nm and 1064 nm excitation," Forensic Science International, vol. 291, pp. 23-38, 2018, https://doi.org/10.1016/j.forsciint.2018.07.032.

[5] A. Hakonen, P. O. Andersson, M. S. Schmidt, T. Rindzevicius, and M. Kall, "Explosive and chemical threat detection by surface-enhanced Raman scattering: A review," Analytica Chimica Acta, vol. 893, pp. 1-13, 2015, https://doi.org/10.1016/j.aca.2015.04.010.

[6] J. Disley, G. Gil-Ramírez, J. Gonzalez-Rodriguez, A review of sensing technologies for nerve agents, through the use of agent mimics in the gas phase: Future needs, TrAC Trends in Analytical Chemistry, Volume 168, 2023, 117282, ISSN 0165-9936, https://doi.org/10.1016/j.trac.2023.117282.

[7] V. Kumar, H. Kim, B. Pandey, T.D. James, J. Yoon, and E. V. Anslyn, "Chemical warfare agent sensors and simulants: a legacy of the 21st century," Chemical Society Reviews, vol. 52, pp. 3152-3207, 2023, https://doi.org/10.1039/D2CS00651K.

[8] M. Xu, X. Hu, H. Zhang, T. Miao, L. Ma, J. Liang, Y. Zhu, H. Zhu, Z. Cheng, and X. Sun, "Detection and pattern recognition of chemical warfare agents by MOS-based MEMS gas sensor array," Sensors, vol. 25, no. 8, 2633, 2025, https://doi.org/10.3390/s25082633.

[9] C. M. Bishop, Pattern Recognition and Machine Learning. New York, NY, USA: Springer, 2006.

[10] S. Haykin, "Cognitive radio: brain-empowered wireless communications," IEEE Journal on Selected Areas in Communications, vol. 23, no. 2, pp. 201-220, 2005, https://doi.org/10.1109/JSAC.2004.839380.

[11] D. L. Donoho, "High-dimensional data analysis: The curses and blessings of dimensionality," AMS Math Challenges Lecture, 2000.

[12] T. Hastie, R. Tibshirani, and J. Friedman, The Elements of Statistical Learning, 2nd ed., Springer, 2009.

[13] A. Gelman et al., Bayesian Data Analysis, 3rd ed., CRC Press, 2013.

[14] Yuan, Z., Luo, X. and Meng, F. (2024), Machine Learning-Assisted Research and Development of Chemiresistive Gas

Sensors. Adv. Eng. Mater., 26: 2400782. https://doi.org/10.1002/adem.202400782R.

[15] Gutierrez-Osuna, "Pattern analysis for machine olfaction: A review," IEEE Sensors Journal, vol. 2, no. 3, 2002.

[16] Aixue Li, Haoyu Yang, Wenxin Yu, Tianyang Liu, Bin Luo, Chunjiang Zhao, Application of machine learning to improve the accuracy of electrochemical sensors: A review, TrAC Trends in Analytical Chemistry, Volume 193, 2025, 118469, ISSN 0165-9936, https://doi.org/10.1016/j.trac.2025.118469.

[17] Zachary J. Baum, Xiang Yu, Philippe Y. Ayala, Yanan Zhao, Steven P. Watkins, and Qiongqiong Zhou, Artificial Intelligence in Chemistry: Current Trends and Future Directions, Journal of Chemical Information and Modeling 2021 61 (7), 3197-3212, https://doi.org/10.1021/acs.jcim.1c00619S.

[18] Marco and A. Gutierrez-Galvez, "Signal and data processing for machine olfaction and chemical sensing," IEEE Sensors Journal, vol. 12, no. 11, pp. 3189-3214, 2012, https://doi.org/10.1109/JSEN.2012.2192920.

[19] Yanchen Li, Zike Wang, Tianning Zhao, Hua Li, Jingkun Jiang, Jianhuai Ye, Electronic nose for the detection and discrimination of volatile organic compounds: Application, challenges, and perspectives, TrAC Trends in Analytical Chemistry, Volume 180, 2024, 117958, ISSN 0165-9936, https://doi.org/10.1016/j.trac.2024.117958.

[20] Soo-Yeon Cho, Youhan Lee, Sangwon Lee, Hohyung Kang, Jaehoon Kim, Junghoon Choi, Jin Ryu, Heeeun Joo, Hee-Tae Jung, and Jihan Kim, Finding Hidden Signals in Chemical Sensors Using Deep Learning, Analytical Chemistry 2020 92 (9), 6529-6537, https://doi.org/10.1021/acs.analchem.0c00137

[21] Pannone, A., Raj, A., Ravichandran, H. et al. Robust chemical analysis with graphene chemosensors and machine learning. Nature 634, 572–578 (2024). https://doi.org/10.1038/s41586-024-08003-w