UDC 004.8:502.3:656.13
DOI https://doi.org/10.32846/2306-9716/2026.eco.2-65.8

IMPLEMENTATION OF AN AUTONOMOUS SELF-EVOLVING AND SELF-HEALING ALGORITHM FOR AI-ORCHESTRATED HYBRID POLLUTION MODELS IN REAL-TIME ROADSIDE POLLUTION CONTROL SYSTEMS

Prykhodko K.V.
Kharkiv National Automobile and Highway University
Yaroslava Mudrogo St, 25, 61002, Kharkiv, Ukraine
kyrylo.prykhodko1410@gmail.com

This paper presents the full-scale implementation and four-month production operation of the autonomous self-evolving and self-healing mode (activated 1 November 2025) of an algorithm that dynamically orchestrates four hybrid pollution dispersion models (LZA, ADBMS-Light, improved Gaussian with terrain and chemistry, and hybrid Gaussian-Lagrangian with Monte Carlo particle tracking) for real-time roadside ecosystem safety management. Powered by a six-agent multi-agent system on Azure OpenAI and tightly integrated with a central MCP (Modular Computation Platform) Tool Registry Server, the algorithm continuously discovers new tools, autonomously creates and registers entirely new ecological-specific tools, maintains 100% test coverage across specialized tests, and performs zero-downtime blue/green deployments. Through the MCP registry the system automatically recognizes the expected control behavior (alerting drivers or activating active purification means) and integrates both alerting and control tools. The ecological reward functions are built strictly around real roadside metrics: time to return to maximum permissible concentrations (MPC), shrinkage of red danger zones on GIS maps, and impact on seasonal roadside vegetation. Agent «Kiro» (programmer) operates in two modes: reactive (on anomaly or test fail) and proactive (in the last 72 hours of each billing period). The system is designed to support multiple large language models with free quotas. Although today Azure OpenAI remains the most advantageous option due to its stability and quality, the architecture is ready to dynamically switch to any other provider (Gemini, Claude etc.) that currently offers larger remaining free credits. Using only the remaining free credits of the most beneficial provider at the moment – typically 8-15 out of 50 – Agent Kiro creates planned improvements or new tools). Field results on highways M03, E40, and P02 (Kharkiv region, November 2025 – February 2026) demonstrate 99.7% uptime, modest but steady additional 8.8% RMSE reduction purely from self-evolution cycles, 7% faster time to MPC recovery, 9% reduction in red-zone area, and full compliance with Ukrainian and WHO standards. Only three major new tools were created by Agent Kiro during the entire period. The developed system represents the final practical implementation stage before the PhD thesis defense on integration of ecological safety control tools for roadside ecosystems and is ready for nationwide scaling. *Key words:* self-evolving algorithm, multi-agent AI, MCP Tool Registry, roadside pollution dispersion, active control tools, 100% ecological test coverage, blue/green deployment, credit-aware proactive evolution, hybrid Gaussian-Lagrangian model

Впровадження автономного саморозвивального та самовідновлюваного алгоритму для гібридних моделей забруднення, керованих штучним інтелектом, у системах контролю забруднення на узбіччі доріг у режимі реального часу. Приходько К.В.

Реалізація автономного самоадаптивного та самовідновлювального алгоритму для автоматизованого оркестрування гібридних моделей розсіювання забруднення в системах управління екологічною безпекою узбіччя автомобільних доріг у реальному часі. У цій статті представлено повномасштабну реалізацію та чотиримісячну виробничу експлуатацію (з 1 листопада 2025 року) автономного режиму саморозвитку та самовідновлення алгоритму, який динамічно оркеструє чотири гібридні моделі розсіювання забруднення (LZA, ADBMS-Light, удосконалена гауссівська модель з урахуванням рельєфу та хімічних процесів, а також гібридна гауссівсько-лагранжева модель з відстеженням частинок методом Монте-Карло) для управління безпекою екосистем узбіччя автомобільних доріг у реальному часі. Алгоритм, що працює на базі шестиагентної багатоагентної системи Azure OpenAI та тісно інтегрований з центральним сервером реєстру інструментів MCP (Modular Computation Platform), постійно виявляє нові інструменти, автономно створює та реєструє повністю нові екологічно-специфічні інструменти, забезпечує 100 % тестове покриття спеціалізованих тестів і виконує розгортання blue/green без простою. Завдяки реєстру MCP система автоматично розпізнає очікувану поведінку керування (оповіщення водіїв або активацію засобів активного очищення повітря) та інтегрує як інструменти оповіщення, так і інструменти керування. Екологічні функції винагороди побудовані виключно на реальних дорожніх метриках: часі повернення концентрацій до максимально допустимих концентрацій (МДК), зменшенні площі червоних зон небезпеки на картах ГІС та впливу на сезонну рослинність узбіччя. Агент «Kiro» (програміст) працює у двох режимах: реактивному (при виявленні аномалії або невдачі тесту) та проактивному (лише в останні 72 години кожного розрахункового періоду). Система спроектована для підтримки кількох великих мовних моделей з безкоштовними квотами. Хоча сьогодні Azure OpenAI залишається найбільш вигідним варіантом завдяки стабільності та якості, архітектура готова динамічно перемикатися на будь-якого іншого провайдера (Gemini, Claude тощо), який наразі пропонує більший залишок безкоштовних кредитів. Використовуючи лише залишок безкоштовних кредитів найбільш вигідного провайдера (зазвичай 8-15 із 50), агент Kiro створює заплановані вдосконалення або нові інструменти. Результати польових випробувань на автомагістралях M03, E40 та P02 (Харківська область, листопад 2025 – лютий 2026 рр.) демонструють 99,7

% безперервної роботи, скромне, але стабільне додаткове зниження RMSE на 8,8 % виключно завдяки циклам саморозвитку, на 7 % швидше повернення до МДК, зменшення площі червоних зон на 9 % та повну відповідність українським нормативам і стандартам ВООЗ. За весь період агент Kiro створив лише три основні нові інструменти. Розроблена система є завершальним практичним етапом підготовки до захисту дисертації на здобуття ступеня доктора філософії з інтеграції інструментів контролю екологічної безпеки екосистем узбіччя доріг і повністю готова до масштабування на національному рівні. *Ключові слова:* алгоритм саморозвитку, багатоагентний ШІ, реєстр інструментів MCP, дисперсія забруднення узбіччя доріг, інструменти активного керування, 100 % екологічне тестове покриття, розгортання blue/green, проактивна еволюція з урахуванням кредитів, гібридна гауссівсько-лагранжева модель

Problem statement. Roadside ecosystems along major highways are exposed to constant and multifaceted pressures from vehicle emissions of fine particulate matter (PM2.5), nitrogen dioxide (NO_2), and carbon monoxide (CO). These pressures are compounded by highly variable meteorological conditions, complex terrain features, seasonal changes in plant cover, and specific wartime operational constraints, including civilian drone bans that force reliance on manual terrain surveys and occasional instability in IoT data streams. Previous static orchestration of pollution dispersion models proved insufficient for maintaining long-term accuracy and reliability under rapidly changing ecological conditions. Manual updates to model parameters or code were too slow, labor-intensive, and costly, creating a critical need for an autonomous system capable of continuous adaptation without human intervention.

Relevance of the study. Real-time roadside ecosystem safety management is crucial for compliance with Ukrainian national standards and WHO air quality guidelines. The development of adaptive systems is particularly relevant given the instability of infrastructure and the need for precise ecological monitoring under wartime conditions. Transforming passive monitoring platforms into closed-loop ecological safety systems capable of active intervention (alerting drivers, activating purification means) addresses an urgent practical necessity. Furthermore, ensuring that such advanced AI systems operate without additional financial costs through credit-aware strategies enhances their viability for widespread deployment in resource-constrained environments.

The connection of the author's contribution with important scientific and practical tasks. The developed system represents the final practical implementation stage before the PhD thesis defense on the integration of ecological safety control tools for roadside ecosystems. The research directly supports the task of nationwide scaling across the entire Ukrainian road network. By integrating ecological safety control tools with autonomous self-evolution mechanisms, the authors address the scientific task of creating living computational ecosystems that grow together with the environment they protect. Practically, the system enables immediate deployment on highways to reduce health risks associated with pollution spikes near schools and residential areas.

Analysis of recent studies and publications. Previous research conducted by the authors successfully developed and field-validated four hybrid dispersion models of progressively increasing complexity: Linear Zone Approximation (LZA), ADBMS-Light, improved Gaussian, and hybrid Gaussian-Lagrangian models. Contemporary developments in dispersion modeling include adaptive and hybrid optimization approaches consistent with real-time air quality systems. Recent studies have explored hybrid quantum-classical and advanced neural network approaches for air quality prediction, including data-driven time-series prediction based on deep learning architectures such as BLSTM. Traditional static modeling platforms such as SimScale or AERMOD provide measurably less long-term accuracy and adaptability compared to autonomous solutions. Comparative studies of dispersion models for roadside ecosystems highlight the benefits of hybrid particle–puff formulations validated in roadside conditions.

Identification of previously unresolved parts of the overall problem addressed in this article. Despite previous advancements, several critical gaps remained unresolved. Static orchestration was insufficient for maintaining accuracy under rapidly changing ecological conditions typical of Ukrainian highways. There was a lack of autonomous self-healing modes capable of detecting ecological anomalies and improving orchestration logic without human input. Additionally, existing systems lacked tight integration with a central tool registry that enables automatic discovery and autonomous creation of ecological-specific tools. Finally, the high cost of AI-driven self-improvement remained a barrier, as previous architectures did not utilize a credit-aware strategy to ensure cost-free evolution using remaining free quotas of LLM providers.

The novelty of the study lies. The core innovations distinguishing this work are fourfold. First, a six-agent multi-agent chat system operates on Azure OpenAI with function calling, guided by carefully designed ecological-specific reward functions. Second, tight integration with a central MCP (Modular Computation Platform) Tool Registry Server enables automatic discovery of externally provided tools as well as autonomous creation and registration of entirely new ecological tools by the system itself. Third, a comprehensive testing framework consisting of specialized tests guarantees 100% code coverage and 100% ecological coverage of all relevant roadside scenarios. Fourth, a credit-aware proactive evolution mode ensures that all self-improvements remain completely cost-free by utilizing only the remaining free Azure OpenAI credits in the final 72 hours of each monthly billing period.

Methodological or general scientific significance. The study introduces six ecological-specific self-evolution algorithms: reinforcement learning using Proximal Policy Optimization (PPO) with weighted ecological indicators; Bayesian optimization with a Gaussian Process surrogate model for tuning dispersion parameters; evolutionary strategies using a (μ + λ) evolution strategy for tool combination discovery; an adaptive Monte Carlo self-correction loop for turbulence changes; physics-informed neural network (PINN) meta-learning for residual correction; and concept drift detection using the Page-Hinkley statistical test. Methodologically, the implementation of Continuous Deployment (CD) with Blue/Green and Canary strategies for ecological models ensures zero-downtime releases. The comprehensive testing framework covers unit, integration, regression, load, chaos, ecological compliance, seasonal, active control, security, credit-aware, mutation, and end-to-end testing, setting a new standard for reliability in safety-critical ecological applications.

Presentation of the main material. Roadside The platform builds upon four hybrid pollution dispersion models previously developed and field-validated by the authors along highways M03, E40, and P02: the Linear Zone Approximation (LZA), ADBMS-Light, the improved Gaussian model with terrain and chemistry corrections, and the hybrid Gaussian-Lagrangian model with Monte Carlo particle tracking. These models were initially orchestrated statically. To overcome the limitations of static orchestration under highly variable meteorological conditions, complex terrain, seasonal vegetation dynamics, and wartime operational constraints (including the civilian drone ban that forced reliance on manual Leica TS07 surveys), the fully autonomous self-evolving and self-healing mode was activated on 1 November 2025 [1, 2].

The core of continuous improvement is a six-agent multi-agent system powered by Azure OpenAI with function calling, tightly integrated with the central MCP (Modular Computation Platform) Tool Registry Server [3]. The MCP registry provides automatic discovery of external tools, secure self-registration of newly created ecological-specific tools, full versioning with rollback, sandbox execution, and seamless provision of both alerting tools (driver mobile notifications and SCADA alarms) and active control tools (Roadside Air Purifiers and Roadvent curb-integrated systems). This bidirectional integration transforms the platform from passive forecasting into a closed-loop ecological safety system capable of active pollution mitigation.

Self-evolution is driven by six complementary mechanisms specifically tailored to the dynamic nature of roadside ecosystems (terrain effects, traffic-induced turbulence, seasonal vegetation buffers, and data-stream instability under wartime conditions). These mechanisms operate under strictly ecological reward functions based on real roadside metrics. The primary reward function used within Proximal Policy Optimization (PPO) is defined as:

$$R = \omega_1\left(1-\frac{RMSE}{RMSE_0}\right)+\omega_1\left(1-\frac{T_{MPC}}{T_{MPC,0}}\right)+ \\ +\omega_1\left(1-\frac{A_{red}}{A_{red,0}}\right)+b_1B_{season}+b_2B_{tool} \qquad (1)$$

where RMSE is the root mean square error of pollution predictions, T_{MPC} is the time required for concentrations to return to maximum permissible levels after a spike, A_{red} is the area of red danger zones on GIS maps, subscript 0 denotes baseline (pre-improvement) values, B_{season} and B_{tool} are bonus terms rewarding improved seasonal performance and successful integration of new MCP tools, and ω_i, b_j are weighting coefficients [4].

The six mechanisms are:

1. PPO-based reinforcement learning – optimizes model orchestration and tool selection using the multi-component ecological reward.
2. Bayesian optimization with Gaussian Process surrogate – efficiently tunes dispersion parameters (σ, κ, λ) against live sensor data.
3. ($\mu+\lambda$) evolutionary strategies – discovers optimal combinations of MCP tools and model modules through population-based fitness evaluation.
4. Adaptive Monte Carlo self-correction loop – dynamically adjusts turbulence modeling in response to detected changes in wind or traffic-induced fluctuations.
5. Physics-informed neural network (PINN) meta-learning – learns and corrects systematic residuals of the hybrid models from recent roadside observations.
6. Page-Hinkley concept drift detection – identifies seasonal or traffic-induced shifts and triggers creation of specialized model variants [5].

All proposed changes are executed only after passing a comprehensive multi-layered test suite that guarantees 100% code coverage and 100% ecological coverage across relevant roadside scenarios, with no regression in key performance metrics.

A typical self-evolution cycle is triggered every 15 minutes or immediately upon anomaly detection. Agent Monitor ingests real-time IoT and GIS data, Agent Analyzer computes the ecological reward, Agent Ecologist provides roadside-specific interpretation, and Agent Model Optimizer applies the appropriate evolution mechanisms. After validation by Agent Tester, Agent Kiro generates code or new tools. Approved updates are deployed using zero-downtime Blue/Green strategy, with Canary rollout (initially 10% of traffic) applied for major modifications.

The system implements a strictly credit-aware proactive evolution mode. Major improvements and creation of new tools by Agent Kiro occur exclusively during the final 72 hours of each billing period, using only the remaining free credits (typically 8-15 out of 50) of the currently most advantageous LLM provider (primarily Azure OpenAI, with architecture ready to dynamically switch to Gemini, Claude or other providers offering larger remaining quotas) [6]. This approach ensures

Table 1

Performance before and after self-evolution (average values, November 2025 – February 2026)

Metric	Before (static orchestration)	After self-evolution	Improvement
RMSE (PM2.5)	1.02 μg/m³	0.93 μg/m³	8.8%
Time to MPC recovery	47 min	43.7 min	7%
Red danger zone area	12.4 ha	11.3 ha	9%
Computation time	2.1 min	1.95 min	7%
Uptime	97.2%	99.7%	+2.5%
Test coverage	89%	100%	+11%
New MCP tools created by Kiro	1	3	+300%

completely cost-free self-evolution while preserving high-quality code generation for ecological applications.

Over four months of uninterrupted production operation (November 2025 – February 2026) on a total of 47 km of highways M03, E40, and P02, the system completed 127 self-evolution cycles. Agent Kiro autonomously created and registered exactly three new ecological tools: the winter salt deposition adjuster, the vegetation buffer optimizer, and the traffic spike predictor. These enhancements, combined with the self-evolution mechanisms, resulted in the quantitative improvements summarized in Table 1.

The system maintained full compliance with Ukrainian national standards and WHO air quality guidelines [7].

Main conclusions. The implemented autonomous self-evolving and self-healing algorithm, augmented with seamless MCP Tool Registry integration (including automatic recognition of expected control behaviour and support for both alerting and active purification tools), a comprehensive test suite achieving 100% coverage, and a strictly credit-aware proactive evolution strategy, successfully transforms a previously static AI-orchestrated modelling platform into a living, continuously improving system for real-time roadside ecosystem safety management. Four months of uninterrupted operation of the self-evolving mode on three major Ukrainian highways have confirmed exceptional reliability (99.7% uptime), modest but consistent accuracy gains (8.8% RMSE reduction attributable solely to self-evolution), and zero additional operational cost. Only three major ecological tools were created by Agent Kiro, each delivering targeted improvements validated by real sensor data. The system is fully prepared for immediate nationwide scaling across the entire Ukrainian road network and provides the complete practical foundation.

References

1. Warfare Ecology: A New Synthesis for Peace and Security / G. E. Machlis, T. Hanson, Z. Spiric, J. Mckendry. Springer, 2011. (NATO Science for Peace and Security Series C: Environmental Security). DOI: 10.1007/978-94-007-1214-0
2. The effects of modern war and military activities on biodiversity and the environment / M. J. Lawrence et al. Environmental Reviews. 2015. Vol. 23, no. 4. P. 443–460. DOI: 10.1139/er-2015-0039
3. Teliura N., Reshetchenko A, Tsapko, N., Vorobiov, O., Kurakova, N. Environmental Safety Management on the Basis of Sustainable Development: Decision-Making Model for Technogenic Loads. Smart Technologies in Urban Engineering. STUE 2023. Lecture Notes in Networks and Systems, vol 808. 2023. P. 167–176. Springer, Cham. https://doi.org/10.1007/978-3-031-46877-3_15
4. Li Y., Hou K., Chang Y., Yuan B., Li X. A methodological study on the identification of ecological security change processes and zoning control strategies – Based on the perspective of sustainable development. Science of The Total Environment. 2024. Vol. 946. Art. 174190. https://doi.org/10.1016/j.scitotenv.2024.174190
5. He, C., Mao, N., Cheng, L., & Du, G. (2025). An Intelligent Choquet Fuzzy Integral-Based Framework for Risk Assessment in Seismic Acquisition Processes. Processes, 13(11), 3558. https://doi.org/10.3390/pr13113558
6. European Commission. The European Green Deal: Communication from the Commission, COM (2019) 640 final, 11.12.2019 URL: https://ec.europa.eu/info/strategy/priorities-2019-2024/european-green-deal (дата звернення: 01.03.2026).
7. Directive 2008/50/EC of the European Parliament and of the Council of 21 May 2008 on ambient air quality and cleaner air for Europe. 2008.

Дата першого надходження статті до видання: 30.03.2026
Дата прийняття статті до друку після рецензування: 30.04.2026
Дата публікації (оприлюднення) статті: 29.05.2026