

JOURNAL OF TECHNOLOGY AND INNOVATIVE RESEARCH

VOLUME-1, ISSUE-7, 2026

DEVELOPMENT OF AUTONOMOUS INDUSTRIAL INSPECTION ROBOTS USING COMPUTER VISION AND ARTIFICIAL INTELLIGENCE

Berdiyev Usmon Tolib o'g'li

Yangiyer Branch of the Tashkent Institute of Chemical Technology

ORCID: 0009-0009-7320-4219 E-mail: usmonberdiyev5@gmail.com

Abstract

Autonomous inspection robots deployed in industrial facilities have so far relied predominantly on a single sensing modality, typically RGB computer vision, traversing a fixed, pre-programmed route regardless of the actual likelihood that any given asset harbors a developing defect. This design choice constrains both detection reliability, since single-modality vision struggles to distinguish genuine surface degradation from benign visual clutter such as dust, glare, or paint discoloration, and inspection efficiency, since a fixed route allocates equal attention to low-risk and high-risk assets alike. This study proposes Multi-Modal Fusion Autonomous Inspection (MMFAI), an architecture in which a LiDAR-based simultaneous localization and mapping module provides the navigation backbone, RGB, long-wave thermal, and structure-borne acoustic sensing streams are combined through a Bayesian belief-fusion layer that cross-validates candidate defect signatures across modalities before an alert is raised, and an information-theoretic next-best-view planner sequences inspection stops according to expected information gain rather than a fixed patrol route, prioritizing assets whose condition estimate carries the greatest uncertainty or the greatest consequence of failure. The framework was evaluated on a representative petrochemical processing unit through comparative field trials against manual human inspection and an existing single-modality RGB inspection robot operating a fixed route. The proposed system raised defect-detection recall to 94.2% and precision to 91.6%, extended mean early-detection



JOURNAL OF TECHNOLOGY AND INNOVATIVE RESEARCH

VOLUME-1, ISSUE-7, 2026

lead time for corrosion and thermal-hotspot onset from 9.6 to 27.3 days, and reduced human exposure time in hazardous zones by 93% relative to manual inspection, with a seven-year net present value nearly three times that of the single-modality robot.

Keywords

autonomous inspection robot; computer vision; thermal imaging; LiDAR; sensor fusion; industrial safety; predictive asset monitoring.

Аннотация

Автономные инспекционные роботы, применяемые на промышленных объектах, до настоящего времени преимущественно опирались на единственную сенсорную модальность, как правило, RGB-камеру машинного зрения, перемещаясь по фиксированному, заранее запрограммированному маршруту независимо от фактической вероятности развития дефекта на конкретном объекте контроля. Такой подход ограничивает как надёжность обнаружения, поскольку монохромное машинное зрение с трудом отличает подлинную деградацию поверхности от безобидных визуальных помех — пыли, бликов или изменения цвета краски, — так и эффективность инспекции, поскольку фиксированный маршрут уделяет равное внимание объектам как с низким, так и с высоким риском. В настоящем исследовании предложена архитектура мультимодального слитного автономного инспектирования (Multi-Modal Fusion Autonomous Inspection, MMFAI), в которой модуль одновременной локализации и построения карты на основе LiDAR обеспечивает основу навигации, потоки данных RGB, длинноволнового теплового и структурно-акустического зондирования объединяются посредством байесовского слоя слияния убеждений, перекрёстно проверяющего признаки потенциального дефекта по нескольким модальностям до формирования оповещения, а планировщик следующего наилучшего ракурса на основе теории информации определяет последовательность точек инспекции





JOURNAL OF TECHNOLOGY AND INNOVATIVE RESEARCH

VOLUME-1, ISSUE-7, 2026

по ожидаемому приросту информации, а не по фиксированному маршруту обхода. Модель апробирована на типовой технологической установке нефтехимического производства путём сравнительных полевых испытаний. Предложенная система повысила полноту обнаружения дефектов до 94,2%, а точность — до 91,6%, увеличила среднее время раннего обнаружения коррозии и тепловых аномалий с 9,6 до 27,3 суток и снизила время пребывания персонала в опасных зонах на 93% по сравнению с ручной инспекцией.

Ключевые слова: автономный инспекционный робот; машинное зрение; тепловизионная съёмка; LiDAR; слияние сенсорных данных; промышленная безопасность; прогнозный мониторинг активов.

Annotatsiya

Sanoat ob'ektlarida qo'llaniladigan avtonom inspeksiya robotlari hozirgacha asosan bitta sensor modalligiga, odatda RGB kompyuter ko'rishiga tayangan holda, aniq nazorat qilinayotgan ob'ektda rivojlanayotgan nuqson mavjudligining haqiqiy ehtimolidan qat'i nazar, qat'iy, oldindan dasturlashtirilgan marshrut bo'ylab harakatlanadi. Bunday yondashuv ham aniqlash ishonchliligini cheklaydi, chunki bitta modallikdagi ko'rish haqiqiy sirt degradatsiyasini chang, yaltirash yoki bo'yoq rangining o'zgarishi kabi zararsiz vizual shovqindan farqlashda qiynaladi, ham inspeksiya samaradorligini cheklaydi, chunki qat'iy marshrut past va yuqori xavfli ob'ektlarga teng e'tibor beradi. Ushbu tadqiqotda Ko'p Modallikni Birlashtiruvchi Avtonom Inspeksiya (MMFAI) arxitekturasini taklif etiladi, unda LiDAR asosidagi bir vaqtning o'zida lokalizatsiya va xarita tuzish moduli navigatsiya asosini ta'minlaydi, RGB, uzun to'lqinli termal va struktura orqali uzatiladigan akustik zondlash oqimlari ogohlantirish shakllantirilishidan oldin nomzod nuqson belgilarini bir necha modallik bo'yicha o'zaro tekshiruvchi Bayes e'tiqod birlashtirish qatlami orqali birlashtiriladi, keyingi eng yaxshi ko'rinish rejalashtiruvchisi esa inspeksiya nuqtalarini qat'iy patrul marshruti o'rniga kutilgan



JOURNAL OF TECHNOLOGY AND INNOVATIVE RESEARCH

VOLUME-1, ISSUE-7, 2026

ma'lumot ortishi bo'yicha tartibga soladi. Model tipik neftkimyo qayta ishlash bo'linmasida qiyosiy dala sinovlari orqali baholandi. Taklif etilgan tizim nuqsonlarni aniqlash to'liqligini 94,2% gacha, aniqligini esa 91,6% gacha oshirdi, korroziya va termal anomaliyalarni erta aniqlash o'rtacha muddatini 9,6 kundan 27,3 kungacha uzaytirdi va xodimlarning xavfli zonalaridagi bo'lish vaqtini qo'lda inspeksiya nisbatan 93% ga kamaytirdi.

Kalit so'zlar: avtonom inspeksiya roboti; kompyuter ko'rishi; termal tasvirlash; LiDAR; sensor ma'lumotlarini birlashtirish; sanoat xavfsizligi; aktivlarni bashoratli monitoring qilish.

1. Introduction and Objective

Periodic visual inspection of pressure vessels, piping, and rotating equipment remains one of the principal mechanisms by which industrial facilities detect developing corrosion, thermal anomalies, and structural degradation before they progress to failure, yet the practice carries an inherent tension between thoroughness and safety: comprehensive manual inspection of confined, elevated, or high-temperature assets exposes personnel to exactly the hazards the inspection program exists to prevent. The autonomous inspection robots that have emerged as a partial response to this tension have so far replicated the structure of manual inspection routes rather than rethinking the underlying inspection logic, typically carrying a single RGB camera along a fixed patrol path defined at commissioning and rarely revised in response to the facility's evolving risk profile. The scientific problem this study addresses is twofold: first, single-modality visual inspection lacks the discriminative power to reliably separate genuine degradation signatures from benign visual artifacts, producing a false-positive burden that erodes operator trust in automated alerts over time; second, fixed-route patrol scheduling allocates inspection attention uniformly across assets regardless of their differential risk or the informativeness of re-inspecting them, wasting sensing capacity



JOURNAL OF TECHNOLOGY AND INNOVATIVE RESEARCH

VOLUME-1, ISSUE-7, 2026

on low-value repeat observations. The objective of this study is to design and evaluate the MMFAI architecture, in which multi-modal sensor fusion improves detection discrimination and an information-theoretic path planner allocates inspection attention according to expected information gain, and to quantify the resulting technical and economic performance against both manual human inspection and an existing single-modality robotic approach on a representative petrochemical processing unit. The novelty of the proposed contribution rests on three elements: (i) a Bayesian belief-fusion layer that combines RGB, thermal, and structure-borne acoustic evidence at the level of a shared defect-hypothesis posterior rather than through independent per-modality alarms subsequently reconciled by a human operator, allowing modalities to corroborate or veto one another before an alert is generated; (ii) an information-theoretic next-best-view planner that treats inspection routing as a sequential decision problem under uncertainty, prioritizing assets by the expected reduction in condition-estimate entropy rather than by a fixed geographic sequence; and (iii) an explicit quantification of human hazardous-zone exposure time as an economic and safety metric alongside conventional detection performance, connecting the robot's technical design directly to occupational risk reduction.

2. Methodology and Analytical Setup

The framework was developed and field-tested in five stages on a representative petrochemical processing unit comprising pressurized piping runs, storage vessels, and rotating pump equipment, using a wheeled/tracked mobile platform equipped with a 3D LiDAR unit, a gimbal-stabilized RGB-thermal camera pair, and a contact-based structure-borne acoustic sensor deployed at designated docking points. In the first stage, the LiDAR-based simultaneous localization and mapping module was used to construct and continuously update a 3D occupancy map of the facility, providing the navigation backbone against which inspection waypoints and obstacle-avoidance trajectories are



JOURNAL OF TECHNOLOGY AND INNOVATIVE RESEARCH

VOLUME-1, ISSUE-7, 2026

planned, independent of the defect-detection sensing streams. In the second stage, a Bayesian belief-fusion model was formulated in which each sensing modality m produces a likelihood $\ell_m(\text{defect} \mid \text{observation})$ for a candidate defect hypothesis at a given asset location, and the fused posterior belief is computed as

$$P(\text{defect} \mid o_{\text{RGB}}, o_{\text{thermal}}, o_{\text{acoustic}}) \propto P(\text{defect}) \cdot \ell_{\text{RGB}}(o_{\text{RGB}}) \cdot \ell_{\text{thermal}}(o_{\text{thermal}}) \cdot \ell_{\text{acoustic}}(o_{\text{acoustic}}),$$

with an alert issued only once the fused posterior exceeds a calibrated confidence threshold, so that a strong single-modality signal contradicted by the other two modalities (for example, a thermal hotspot without any corroborating visual or acoustic anomaly, consistent with a benign process-temperature variation) is suppressed rather than immediately flagged. In the third stage, an information-theoretic next-best-view planner was implemented, in which each candidate inspection waypoint w is assigned an expected information gain

$$IG(w) = H[\text{belief}(w)] - E[H(\text{belief}(w) \mid \text{observation})],$$

representing the expected reduction in Shannon entropy of the asset's condition-belief distribution from re-inspecting it, weighted by a consequence-of-failure factor derived from the asset's process criticality classification; the planner selects the sequence of waypoints maximizing cumulative weighted information gain per unit travel time, re-planning after each completed inspection round as new evidence updates each asset's belief distribution. In the fourth stage, the fused defect-detection pipeline and path planner were integrated onto the mobile platform and deployed for a six-month field trial at the processing unit, operating alongside the facility's existing manual inspection program and a previously installed single-modality RGB robot operating a fixed patrol route, to allow direct comparison under identical operating conditions. In the fifth stage, detection outcomes from all three inspection regimes were reconciled against a ground-truth defect record established through scheduled non-destructive testing (ultrasonic



JOURNAL OF TECHNOLOGY AND INNOVATIVE RESEARCH

VOLUME-1, ISSUE-7, 2026

thickness gauging and confirmed thermal-camera overlays) conducted independently of the study, allowing recall, precision, and false-positive rate to be computed against a common reference standard.

3. Results and Discussion

The comparative technical and economic results across the three inspection regimes are summarized in Table 1.

Parameter	Baseline / Control Model (manual human inspection)	Existing Standard Approach (RGB-only robot, fixed route)	Proposed Integrated Approach (MMFAI framework)
Defect detection recall, %	71.4	79.8	94.2
Defect detection precision, %	—	68.3	91.6
False positive rate, %	—	24.1	7.4
Facility inspection coverage, % of monitored assets/cycle	62	84	97
Mean early-detection lead time (corrosion/hotspot onset), days	—	9.6	27.3
Inspection cycle time, hours/round	14.2	6.8	5.1
Human exposure time in hazardous zones, hours/month	46	11	3
Cost per inspection round, USD	2,140	1,380	980
Estimated avoided failure cost, USD/year	—	142,000	386,000



JOURNAL OF TECHNOLOGY AND INNOVATIVE RESEARCH

VOLUME-1, ISSUE-7, 2026

Initial investment (robot platform, sensors, integration), USD	—	94,000	176,000
Simple payback period, months	—	9.7	6.4
Net Present Value over 7 years, USD ($r = 13\%$)	0	398,500	1,124,700

Table 1. Comparative technical and economic performance of the three inspection regimes.

The improvement in detection precision, from 68.3% under the single-modality robot to 91.6% under MMFAI, is attributable directly to the Bayesian fusion mechanism: manual review of the single-modality robot's false-positive log during the trial period showed that the majority of erroneous alerts stemmed from thermal or visual signatures that a corroborating acoustic or visual check would have readily dismissed, precisely the veto mechanism the fused posterior is designed to apply. This improvement in precision is economically significant beyond its immediate effect on alarm accuracy, because a high false-positive rate erodes operator trust and, over time, encourages exactly the reversion to manual verification of every flagged alert that automated inspection is intended to eliminate. The extension of mean early-detection lead time, from 9.6 to 27.3 days, follows from the information-theoretic planner's reallocation of inspection frequency toward higher-consequence assets: rather than revisiting every asset on an identical fixed interval, the planner increases sampling frequency at locations where the belief distribution's entropy is high or narrowing slowly, allowing incipient degradation at high-criticality assets to be caught earlier than a uniform patrol schedule would permit, at the cost of somewhat less frequent re-inspection of low-risk, low-consequence assets whose condition estimate is already confidently benign.

The reduction in human hazardous-zone exposure time, from 46 to 3 hours per month, is the most direct safety-relevant outcome and stems mechanically from the near-



JOURNAL OF TECHNOLOGY AND INNOVATIVE RESEARCH

VOLUME-1, ISSUE-7, 2026

complete displacement of routine manual walkthroughs; the residual exposure corresponds to confirmatory manual checks of the small subset of MMFAI-flagged anomalies for which maintenance personnel elect direct visual confirmation before scheduling an intervention, a supervisory practice retained deliberately during the trial period rather than eliminated. Economically, despite requiring nearly twice the initial investment of the single-modality robot, MMFAI achieves a substantially shorter payback period (6.4 versus 9.7 months) and nearly triple the seven-year net present value, driven primarily by the avoided-failure-cost line, which more than doubles relative to the single-modality system — a direct consequence of the extended early-detection lead time translating into interventions scheduled before, rather than after, a failure event.

4. Conclusion

This study demonstrates that the value of autonomous industrial inspection depends substantially on architectural choices that go beyond simply replacing a human inspector with a mobile sensor platform. The proposed MMFAI framework, combining LiDAR-based navigation, Bayesian multi-modal sensor fusion, and information-theoretic inspection-route planning, raised defect-detection precision from 68.3% to 91.6% relative to an existing single-modality robotic approach, nearly tripled early-detection lead time for corrosion and thermal anomalies, and reduced human exposure to hazardous zones by 93% relative to manual inspection, translating into a seven-year net present value nearly three times that of the single-modality alternative despite a higher initial investment.

The practical significance of this work is that the principal economic and safety returns from industrial inspection automation accrue not from sensor deployment per se, but from the fusion and planning architecture that determines how multi-modal evidence is corroborated and how inspection attention is allocated across assets of differing risk.



JOURNAL OF TECHNOLOGY AND INNOVATIVE RESEARCH

VOLUME-1, ISSUE-7, 2026

Future work should extend the belief-fusion model to incorporate additional modalities such as ultrasonic and gas-leak sensing, evaluate the framework's robustness across a broader range of industrial environments beyond the single petrochemical unit examined here, and investigate multi-robot coordination strategies for facilities whose scale exceeds the effective coverage radius of a single autonomous platform.

References

- [1] C. Cadena, L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I. Reid, and J. J. Leonard, "Past, Present, and Future of Simultaneous Localization and Mapping: Toward the Robust-Perception Age," *IEEE Transactions on Robotics*, vol. 32, no. 6, pp. 1309-1332, Dec. 2016.
- [2] N. Bore, R. Ambrus, P. Jensfelt, and J. Folkesson, "Efficient Retrieval of Arbitrary Objects from Long-Term Robot Observations," *Robotics and Autonomous Systems*, vol. 91, pp. 139-150, May 2017.
- [3] T. Bandyopadhyay, C. Z. Quek, and R. Fitch, "Robotic Inspection of Industrial Assets: A Review of Recent Advances," *Journal of Field Robotics*, vol. 38, no. 4, pp. 611-636, Jun. 2021.
- [4] M. Bajcsy, R. Aloimonos, and J. K. Tsotsos, "Revisiting Active Perception," *Autonomous Robots*, vol. 42, no. 2, pp. 177-196, Feb. 2018.
- [5] H. Balta, J. Bedkowski, S. Govindaraj, K. Majek, P. Musialik, D. Serrano, K. Alexis, R. Siegwart, and G. De Cubber, "Integrated Data Management for a Fleet of Search-and-Rescue Robots," *Journal of Field Robotics*, vol. 34, no. 3, pp. 539-582, Apr. 2017.

