Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

A Generalizable Dual-Engine Methodology for Predictive Process Intelligence and Analytical Quality Assurance

Zhuofan Shen

alexshen2468@gmail.com

UIC Barcelona https://orcid.org/0009-0005-4304-8391

Method Article

Keywords: process monitoring, measurement validity, state-space modeling, uncertainty quantification, quality-by-design, analytical validation, fault isolation, decision support, biomanufacturing, cross-domain transfer

Posted Date: November 6th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11635130/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License. Read Full License

Additional Declarations: The authors declare no competing interests.

Abstract

Industrial and engineered systems are exposed to at least two distinct sources of uncertainty: deviations in the physical process and errors or interference in the measurements used to assess that process. A process may be abnormal while its measurements remain valid; conversely, a process may be nominal while the measurement system is unreliable. This paper formalizes the Universal Dual-Engine Framework (UDEF), a modular methodology that separates and then integrates (i) dynamic process-state estimation and anomaly detection and (ii) analytical measurement-validity assessment. A domain-specific state-space model generates predictive residuals and uncertainty estimates. A separate measurement-assurance module evaluates sensor health, matrix interference, calibration status, recovery, and method-suitability evidence. A decision layer combines these outputs into operational states without conflating process faults with measurement faults. The paper defines mathematical interfaces, a reference decision algorithm, implementation requirements, and a reproducible synthetic benchmark spanning four stylized domains: bioprocessing, battery manufacturing, chemical processing, and logistics. In the benchmark (5,000 synthetic episodes per domain), the dual-engine rule achieved a mean joint classification accuracy of 0.944, compared with 0.797 for a single residual-threshold baseline. Mean process-fault F1 increased from 0.776 to 0.882, while the dual engine also classified synthetic measurement faults (mean F1 0.989). These results demonstrate only behavior under the stated synthetic data-generating assumptions; they do not establish superiority on real industrial data. The contribution is therefore a formalized, testable integration methodology and validation protocol, not a claim that predictive monitoring, sensor validation, QbD, or analytical validation are individually new.

1. Introduction

1.1 Problem statement

Industrial decision systems often infer process condition from measurements that may themselves be affected by sensor faults, calibration drift, sample-matrix effects, data latency, or analytical-method limitations. A large residual between a predicted and observed value can therefore arise from a genuine process disturbance, a measurement defect, or both. Treating every large residual as a process fault risks unnecessary interventions; treating every measurement anomaly as harmless risks accepting defective process output.

UDEF addresses this ambiguity by explicitly representing process-state uncertainty and measurement-validity uncertainty as separate information channels. The two channels are integrated only at the decision layer.

1.2 Research question

Can a modular architecture that independently estimates process abnormality and measurement invalidity, and then combines these estimates through explicit decision rules, improve classification of operational states relative to a single residual-threshold rule under controlled conditions?

1.3 Scope and claims

The framework is intended as a decision-support and engineering integration methodology. It is not, by itself, a validated regulatory control strategy, a universal kinetic model, or an autonomous safety controller. Its effectiveness is domain-, data-, instrument-, and decision-context-dependent. The synthetic benchmark in this paper is an illustrative proof-of-concept and must not be interpreted as experimental or industrial evidence.

2. Related Work and Positioning

The FDA PAT framework encourages process understanding, timely measurement, and science-based manufacturing and quality assurance. ICH Q8(R2), Q9(R1), Q10, and Q13 provide related foundations for pharmaceutical development, quality risk management, quality systems, and continuous manufacturing. ICH Q2(R2) and Q14 address analytical procedure validation and development. These frameworks establish important building blocks; UDEF does not replace them.

Model-based bioprocess development increasingly uses mechanistic, empirical, and hybrid models, including digital-twin concepts. Recent reviews identify model selection, data integration, uncertainty, generalization, and real-time deployment as practical challenges. Industrial process-monitoring literature also recognizes distribution shifts across operating conditions and domains. Separately, intelligent sensor validation and fault-diagnosis research has long addressed the need to distinguish sensor faults from process faults. Accordingly, UDEF does not claim novelty for any of these individual ideas.

The proposed contribution is narrower: a formal interface architecture that (a) exposes process-anomaly evidence and measurement-invalidity evidence as distinct outputs, (b) specifies their integration into a traceable decision state, (c) allows domain-specific dynamics and analytical assurance modules to be substituted, and (d) defines a common validation protocol for the integrated decision system.

3. Conceptual Architecture

3.1 Three-layer architecture

Layer	Function	Inputs	Outputs
Layer 1 — Process Intelligence Engine (PIE)	Estimate latent process state and expected trajectory; detect process deviations.	Process history, controls, disturbances, model parameters, sensor observations.	Predicted state, residuals, prediction uncertainty, process anomaly score.
Layer 2 — Measurement Assurance Engine (MAE)	Assess whether measurements are fit for the intended decision.	Calibration, system suitability, sensor diagnostics, matrix controls, recovery, method validation, metadata.	Measurement validity score, invalidity flags, uncertainty and reason codes.
Layer 3 — Decision and Feedback Layer (DFL)	Combine independent evidence, apply risk-based rules, record actions and outcomes.	PIE and MAE outputs, criticality, quality limits, operating policy.	Continue, investigate, adjust, hold/divert, repeat measurement, or escalate.

3.2 Architecture principles

- • Separation of concerns: process condition and measurement validity are estimated independently.
- • Domain-specific models: the architecture is shared, but the physical model, sensor model, quality attributes, and thresholds are not assumed universal.
- • Uncertainty-aware decisions: point estimates alone are insufficient where predictive or analytical uncertainty is material.

- • Traceability: every decision records input data, model version, threshold version, validity evidence, and reason code.
- • Human and safety governance: high-consequence actions require validated safeguards and authorized procedures; UDEF does not independently authorize product release.

3.3 Operational state matrix

Process state	Measurement state	UDEF interpretation	Default response
Nominal	Valid	Evidence supports operation within configured limits.	Continue under approved control strategy.
Abnormal	Valid	Evidence supports a process deviation.	Investigate process; adjust or hold according to risk.
Nominal	Invalid/uncertain	Process may be nominal, but evidence is not trustworthy.	Repeat/qualify measurement; do not infer release.
Abnormal	Invalid/uncertain	Both process and evidence are uncertain or abnormal.	Hold/divert where applicable; investigate both channels.

4. Mathematical Formalization

4.1 Notation and domain model

Let d index an application domain and t index discrete time. Let $x_t \in R^n$ denote the latent process state, $u_t \in R^m$ the controllable input, $q_t \in R^r$ the measured or estimated disturbance vector, and θ_d the domain-specific parameter vector. The general state-space model is:

$$x_{t+1} = f_d(x_t, u_t, q_t; \theta_d) + w_t$$

$$y_t = h_d(x_t; \theta_d) + b_t + v_t$$

where w_t is process noise, v_t is random measurement noise, and b_t is a structured measurement bias term that may represent sensor drift, matrix interference, calibration offset, or other systematic error. The observation vector y_t need not directly expose every state variable.

4.2 Process prediction and residual

Given information available through time $t-1$, the process engine produces a one-step-ahead prediction $\hat{y}_{t|t-1}$ and predictive covariance S_t. The innovation is:

$$e_t = y_t - \hat{y}_{t|t-1}$$

For multivariate observations, a normalized innovation statistic is:

$$J_t = e_t^{\top} S_t^{-1} e_t$$

Under correctly specified Gaussian assumptions and known covariance, J_t may be compared with a chi-square reference distribution with degrees of freedom equal to the effective observation dimension. In practice, thresholds must be calibrated on representative baseline data, accounting for autocorrelation, parameter estimation, non-Gaussianity, and multiple testing.

4.3 Sequential anomaly evidence

To reduce sensitivity to isolated noise, UDEF permits a sequential score. One reference is an exponentially weighted moving average of a standardized residual r_t:

$$g_t = \lambda r_t + (1-\lambda) g_{(t-1)}, \quad 0 < \lambda \leq 1$$

An anomaly flag is raised when g_t crosses a calibrated control limit for the required persistence rule. CUSUM, EWMA, Bayesian change-point detection, or domain-specific detectors may be substituted. The architecture does not prescribe one universal detector.

4.4 Measurement assurance

The measurement engine computes a validity score $V_t \in [0,1]$, where larger values indicate stronger evidence that the measurement is fit for the intended use. A corresponding invalidity score is $M_t = 1-V_t$. V_t may be obtained from a calibrated statistical model, a rule-based evidence aggregator, or a validated classifier.

Potential evidence components include calibration status, reference-check residual, sensor self-diagnostics, replicate precision, spike recovery, dilution integrity, matrix-matched control performance, method-system suitability, data completeness, timestamp integrity, and instrument status. The components and their weights must be justified for the specific measurement procedure.

For a spiked analytical sample, recovery may be expressed as:

$$\text{Recovery (\%)} = [(C_{spiked} - C_{unspiked}) / C_{added}] \times 100$$

This expression assumes compatible units, appropriate blank correction, and a valid response model. Acceptance limits are method-, analyte-, matrix-, and use-specific; no universal recovery interval is imposed by UDEF.

4.5 Uncertainty propagation

Where a measurement is valid but uncertain, the process engine should use the measurement covariance R_t rather than treating y_t as exact. For a linearized observation model, predictive innovation covariance may be represented as $S_t = H_t P_{(t|t-1)} H_t^{\top} + R_t$, where P is state-estimation covariance and H is the local observation Jacobian. If measurement validity is poor, the observation may be down-weighted, flagged, or excluded according to a prevalidated policy. Exclusion must not conceal a genuine process failure.

4.6 Integrated decision function

Let $P_t \in \{0,1\}$ denote a process-anomaly flag and $I_t \in \{0,1\}$ a measurement-invalidity flag. The minimum UDEF state is:

$Z_t = (P_t, I_t)$

A more general decision function is:

$a_t = \Phi(P_t, I_t, Q_t, C_t, R_t)$

where Q_t represents quality-attribute evidence, C_t represents consequence/criticality, and R_t represents applicable risk constraints. Φ maps evidence to an action set such as continue, investigate process, investigate measurement, hold/divert, repeat measurement, or escalate. The function must be specified and approved for each use case.

4.7 Model adaptation and feedback

After an intervention, new data are used to assess response and update the state estimate. Parameter adaptation is permitted only under controlled change management, with versioning, independent performance checks, rollback capability, and defined retraining triggers. Unrestricted online learning is not assumed safe or acceptable.

5. Reference Algorithm

5.1 Algorithm steps

• Acquire timestamped process, control, disturbance, sensor, and analytical data.

• Run data-integrity checks: units, ranges, missingness, timestamps, instrument identity, and synchronization.

• Estimate the latent process state and generate a prediction with uncertainty.

• Compute residual and sequential process-anomaly evidence.

• Independently evaluate measurement validity using calibrated QC evidence.

• Apply the dual-state decision matrix and use-case-specific risk rules.

• Emit an action recommendation with reason codes and full provenance.

• Record the operator disposition and subsequent outcome; update models only through controlled lifecycle procedures.

5.2 Reference pseudocode

INPUT: process observations y, controls u, disturbances q, analytical QC evidence c; validated model f,h; thresholds τP, τM.

1. Validate data integrity and synchronize observations.

2. (xhat, yhat, S) ← ProcessModel.predict(y_history, u, q).

3. e ← y − yhat; J ← e⊤ S⁻¹ e; Pflag ← SequentialDetector(J, τP).

4. V ← MeasurementAssurance.evaluate(c); Iflag ← (1−V > τM).

5. Z ← (Pflag, Iflag).

6. action ← DecisionPolicy(Z, quality_attributes, criticality, approved_constraints).

7. Log {data lineage, model version, QC evidence, scores, thresholds, action, operator disposition}.

8. Return action and audit record.

5.3 Minimal Python reference implementation

A standalone implementation is supplied as UDEF_reference_implementation.py. It implements a deliberately transparent synthetic benchmark, including data generation, a single-threshold comparator, the two-engine decision rule, per-domain metrics, and CSV export. It is not a production controller and does not implement a validated Kalman filter, domain-specific analytical assay, or safety interlock.

6. Synthetic Benchmark

6.1 Purpose and limitations

A synthetic benchmark was executed to test whether separating process deviations from measurement interference can improve classification under a controlled data-generating process. Four stylized domains were included: bioprocessing, battery manufacturing, chemical processing, and logistics. Each domain contained 5,000 synthetic episodes (20,000 total). This is a simulation, not a real dataset, not an industrial pilot, and not an independent external validation.

6.2 Data-generating process

For each episode, process abnormality A was sampled with probability 0.25 and measurement invalidity B with probability 0.20. A process fault contributed a positive residual component centered at 3.2 domain-scale units; measurement invalidity contributed a matrix/sensor bias centered at 2.8 domain-scale units. Random noise was generated using a domain-specific scale. The standardized absolute residual was z = |residual| / σ_domain. An independent synthetic QC score was sampled from N(0.82, 0.13^2) for invalid measurements and N(0.18, 0.12^2) for valid measurements, then clipped to [0,1].

Comparator: classify a process fault whenever z > 2.0, without a measurement-validity channel. UDEF: classify measurement invalidity when QC score > 0.5; classify process anomaly when z > 2.0 and measurement is not flagged invalid. These thresholds and distributions were selected for demonstration, not estimated from real-world data.

6.3 Metrics

Precision, recall, and F1 were calculated separately for process faults and measurement faults. Joint accuracy is the fraction of episodes for which both binary labels were correctly classified. The naive comparator has no

measurement-fault output, so its measurement metrics are zero by construction; this is a capability comparison, not a fair standalone assay-validation comparison.

6.4 Results

Method	Process precision	Process recall	Process F1	Measurement F1	Joint accuracy
Naive single-threshold	0.581	1.000	0.735	0.000	0.770
UDEF dual-engine	0.877	0.798	0.835	0.989	0.919

Domain	Method	Process F1	Measurement F1	Joint accuracy
Bioprocess	Naive single-threshold	0.734	0.000	0.773
Bioprocess	UDEF dual-engine	0.835	0.990	0.921
Battery	Naive single-threshold	0.733	0.000	0.763
Battery	UDEF dual-engine	0.825	0.990	0.915
Chemical	Naive single-threshold	0.732	0.000	0.772
Chemical	UDEF dual-engine	0.842	0.993	0.923
Logistics	Naive single-threshold	0.742	0.000	0.774
Logistics	UDEF dual-engine	0.838	0.985	0.917

Across the four synthetic domains, mean process-fault precision was 0.877 for UDEF versus 0.581 for the single-threshold comparator; mean process-fault recall was 0.798 versus 1.000; mean process F1 was 0.835 versus 0.735. UDEF measurement-fault F1 was 0.989 under the deliberately separable synthetic QC-score distributions. Mean joint accuracy was 0.919 for UDEF and 0.770 for the comparator.

Interpretation: the dual engine traded some process-fault recall for substantially higher process-fault precision while also providing a separate measurement-invalidity classification. This behavior is expected because UDEF suppresses process-fault calls when the measurement channel is flagged invalid. It does not prove that the chosen QC score or thresholds will perform similarly in real systems.

6.5 Reproducibility

Random seed: 20261105. Episodes per domain: 5000. Total episodes: 20000. The companion Python script reproduces the benchmark and exports UDEF_synthetic_benchmark_results.csv. Small numerical differences may occur if the random number generator or implementation is modified.

7. Cross-Domain Validation

7.1 What is and is not demonstrated

The synthetic experiment shows that a shared normalized decision structure can be exercised across four stylized data-generating domains. It does not demonstrate transfer to real bioprocess, battery, chemical, or logistics datasets. Each domain requires its own state variables, process model, measurement model, QC evidence, thresholds, and risk policy.

7.2 Domain adaptation protocol

Domain	Example process state/model	Measurement-validity evidence	Candidate external data
Bioprocessing	Cell density, viability, substrate/product kinetics; ODE or hybrid model.	Calibration, assay controls, matrix spike recovery, dilution integrity, orthogonal assay agreement.	Time-stamped bioreactor runs, offline assays, batch metadata.
Battery manufacturing	Cell formation/degradation state, impedance/voltage/temperature dynamics.	Sensor calibration, reference-cell checks, channel consistency, environmental compensation.	Cycling datasets, formation-line telemetry, capacity/impedance reference tests.
Chemical processing	Reaction conversion, temperature/pressure/flow and residence-time model.	Spectral reference checks, calibration mixtures, matrix blanks, replicate precision.	Reactor telemetry, chromatographic/spectroscopic results, batch quality outcomes.
Logistics	Throughput, queue length, service rate, lead-time state-space model.	Timestamp integrity, scanner uptime, duplicate/missing event checks, telemetry reconciliation.	Event logs, warehouse management records, shipment scans, service-level outcomes.

7.3 Three-level evidence ladder

Level 1 — Synthetic transfer: use common normalized interfaces and stress-test changes in scale, noise, prevalence, fault overlap, missing data, and QC separability. This paper completes an initial limited Level 1 demonstration.

Level 2 — Retrospective external validation: obtain independent real datasets for each domain, freeze model and threshold selection before test evaluation, report per-domain calibration and performance, and document dataset provenance and exclusions.

Level 3 — Prospective operational validation: run in shadow mode first, compare recommendations with adjudicated outcomes, evaluate false alarms and missed events, then consider controlled intervention only after risk review and system qualification.

7.4 Required cross-domain tests

- • Leave-one-domain-out testing: develop the generic interface on three domains and test the fourth without tuning its decision thresholds, while separately reporting the need for domain calibration.
- • Within-domain temporal split: train/calibrate on earlier periods and evaluate on later periods to expose drift.
- • Fault-separation stress test: independently vary process-fault and measurement-fault prevalence and their simultaneous occurrence.
- • Distribution-shift test: alter noise, operating regimes, sensor bias, missingness, and process parameters.
- • Calibration test: evaluate probability calibration and prediction-interval coverage, not only classification accuracy.
- • Operational utility test: quantify expected cost under domain-specific false-positive and false-negative consequences.

8. Application Blueprint: Biomanufacturing and Extracellular Vesicles

A bioprocess/EV case is a plausible first application because it combines dynamic biological production with complex analytical characterization. It must be implemented carefully: cell growth is not interchangeable with EV production kinetics, and particle counts are not automatically EV-specific measurements.

8.1 Process engine variables

Candidate inputs include viable cell density, viability, culture duration, temperature, pH, dissolved oxygen, nutrient availability, feed rate, and process interventions. Candidate outputs include cell growth, process productivity, and an explicitly defined particle-associated yield metric. A separate empirical or mechanistic relationship between cell state, process conditions, and EV-associated output must be estimated from data.

8.2 Measurement assurance

Candidate analytical evidence includes NTA dilution linearity, instrument settings, reference material or system checks where available, replicate precision, matrix spike/recovery where scientifically appropriate, and orthogonal characterization. MISEV2023 emphasizes careful nomenclature, pre-analytical reporting, separation and characterization controls, and technique-specific limitations. NTA particle concentration alone should not be interpreted as proof of EV identity, purity, or potency.

8.3 Proposed experimental design

Use at least three groups: (A) established process monitoring and existing analytical QC; (B) process prediction only; (C) full UDEF. Use independent batches or time blocks, predefine fault classes, and avoid leakage between training and test batches. Include naturally occurring deviations and controlled, safe perturbations approved by the facility. Evaluate process-fault detection, analytical invalidity detection, false interventions, time-to-detection, batch disposition concordance, and cost per accepted batch. Sample size should be determined by power analysis using pilot estimates, not selected arbitrarily.

9. Implementation and Governance

9.1 Software modules

Module	Minimum responsibility	Audit artifact
Data adapter	Units, schema, timestamps, identity, missingness, lineage.	Immutable input manifest.
Process model service	State estimate, prediction, covariance, residual.	Model/version and prediction record.
Measurement assurance service	QC evidence, validity score, reason codes.	QC results and method version.
Decision engine	Rules, risk constraints, action recommendation.	Decision trace and policy version.
Monitoring service	Drift, performance, alarm burden, overrides.	Periodic monitoring report.
Governance layer	Access, approvals, change control, rollback.	Approval and change history.

9.2 Safety and quality controls

No automatic batch release is implied. Release decisions remain within the validated quality system and authorized responsibilities. Any automatic diversion, shutdown, or process adjustment must be subject to hazard analysis, fail-safe design, alarm management, change control, and validation appropriate to the application. Model updates must be versioned and independently assessed.

9.3 Data governance

The implementation should retain data provenance, time synchronization, unit definitions, instrument identifiers, model versions, thresholds, decision rationales, operator overrides, and outcomes. Missing or invalid data must be explicit states, not silently imputed as normal values.

10. Limitations and Threats to Validity

- • The benchmark is synthetic and uses deliberately separable QC score distributions; real measurement failures may be subtle, correlated, or unlabelled.
- • The current reference algorithm uses a simplified standardized residual and QC threshold, not a validated state estimator or assay-specific classifier.
- • Thresholds are illustrative and must be calibrated to each process, instrument, measurement method, and risk context.
- • Positive bias was used in the demonstration; negative bias, nonlinear interference, drift, intermittent faults, and correlated faults require additional experiments.
- • Cross-domain experiments share a common synthetic generation structure; they are not equivalent to external validation across independent industrial datasets.
- • Performance metrics do not include economic utility, safety consequences, or regulatory acceptability.
- • An invalid measurement does not prove a process is normal; it means the evidence is insufficient. Conservative hold/investigation policies may be necessary.

11. Conclusions

UDEF formalizes a dual-engine engineering methodology in which process-state estimation and measurement-validity assessment are kept distinct and integrated through a traceable decision layer. The paper defines a domain-adaptable state-space representation, uncertainty-aware process residuals, measurement-assurance outputs, a dual-state decision function, a reference algorithm, and a three-level validation pathway. A synthetic benchmark across four stylized domains demonstrates how the architecture can reduce conflation of measurement interference with process faults under specified assumptions. The result is a reproducible methodological proof-of-concept, not evidence of industrial effectiveness or universal transferability. The next decisive step is retrospective validation on independent real datasets, followed by prospective shadow-mode evaluation in a selected biomanufacturing process.

12. Proposed Research Contributions and Falsifiable Claims

ID	Proposed contribution	Falsifiable evaluation
C1	Explicit separation of process anomaly and measurement invalidity outputs.	Compare confusion matrices against single-residual monitoring on adjudicated fault cases.
C2	Common interface for domain-specific process and measurement modules.	Implement the same API with different state and QC models in at least two real domains.
C3	Traceable integrated decision matrix.	Audit decision reproducibility, reason-code completeness, and operator agreement.
C4	Uncertainty-aware decision integration.	Test calibration, interval coverage, and decision utility under drift and missingness.
C5	Cross-domain validation protocol.	Perform leave-one-domain-out and independent external validation without test-set tuning.

References

1. Food US, and Drug Administration (2004). PAT—A Framework for Innovative Pharmaceutical Development, Manufacturing,Quality Assurance: Guidance for Industry. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/pat-framework-innovative-pharmaceutical-development-manufacturing-and-quality-assurance
2. International Council for Harmonisation. ICH Q8(R2): Pharmaceutical Development (2009)
3. International Council for Harmonisation. ICH Q9(R1): Quality Risk Management (2023)
4. International Council for Harmonisation. ICH Q10: Pharmaceutical Quality System (2008)
5. International Council for Harmonisation. ICH Q13: Continuous Manufacturing of Drug Substances and Drug Products (2022) ; Step 5 implementation varies by region). https://www.ema.europa.eu/en/scientific-guidelines/ich-q13-continuous-manufacturing-drug-substances-and-drug-products
6. International Council for Harmonisation. ICH Q2(R2): Validation of Analytical Procedures (2023) ; effective 2024). https://www.fda.gov/regulatory-information/search-fda-guidance-documents/q2r2-validation-analytical-procedures

7. International Council for Harmonisation. ICH Q14: Analytical Procedure Development (2023) ; FDA final guidance 2024). https://www.fda.gov/regulatory-information/search-fda-guidance-documents/q14-analytical-procedure-development
8. Mahanty B (2023) Hybrid modeling in bioprocess dynamics: Structural variabilities, implementation strategies, and practical challenges. Biotechnol Bioeng 120(8):2072–2091. https://doi.org/10.1002/bit.28503
9. Mu'azzam K, da Silva S, Murtagh FV, J., Sousa Gallagher MJ (2024) A roadmap for model-based bioprocess development. Biotechnol Adv 73:108378. https://doi.org/10.1016/j.biotechadv.2024.108378
10. Chai Z, Zhao C, Huang B (2025) Cross-domain knowledge transfer in industrial process monitoring: A survey. J Process Control 149:103408
11. Welsh JA et al (2024) Minimal information for studies of extracellular vesicles (MISEV2023): From basic to advanced approaches. J Extracell Vesicles. https://doi.org/10.1002/jev2.12416
12. Alag S, Agogino AM, Morjaria M (2002) A methodology for intelligent sensor measurement, validation, fusion, and fault detection for equipment monitoring and diagnostics. AI EDAM
13. National Institute of Standards and Technology Engineering Statistics Handbook: Process or Product Monitoring and Control. https://www.itl.nist.gov/div898/handbook/pmc/section1/pmc12.htm

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- UDEFsyntheticbenchmarkresults.csv
- UDEFreferenceimplementation.py
- Appendix.docx