

ORIGINAL RESEARCH

Open access

From Quality-by-Design to Quality-by-Intelligence: Governance of AI-Assisted Pharmaceutical Manufacturing Systems

Michael Turner¹, Sophia Nguyen^{2*}, David Clark¹, Emma Wilson³

Abstract

Artificial intelligence is rapidly entering pharmaceutical manufacturing through predictive analytics, process analytical technology, continuous manufacturing platforms, and data-driven quality control. These developments promise earlier detection of process deviation, adaptive optimization of critical process parameters, and more responsive assurance of critical quality attributes. Yet the dominant regulatory and quality vocabulary remains anchored in Quality-by-Design, a paradigm built around pre-defined knowledge, structured risk assessment, and validated control strategies. The central problem is that AI-assisted manufacturing systems do not behave like conventional pharmaceutical processes governed solely through fixed design spaces. Machine learning models may change performance as data distributions shift, as sensors age, as materials vary, or as feedback loops alter the operating environment. This creates a governance gap between the static logic of pre-validation and the dynamic reality of algorithmic decision-making. This critical conceptual review examines whether Quality-by-Design remains sufficient for AI-assisted pharmaceutical manufacturing. It argues that QbD is still necessary but no longer sufficient because it was designed for processes whose boundaries, models, and control logic can be defined before routine operation. AI-assisted systems require an additional governance logic capable of supervising data, models, explanations, adaptation, and accountability throughout the product lifecycle. The review develops the concept of Quality-by-Intelligence as a forward-looking governance framework. QbI does not discard QbD; rather, it extends it by embedding data stewardship, model lifecycle management, explainable decision-making, adaptive risk control, and regulatory translation into the pharmaceutical quality system. Its purpose is to govern not only the manufacturing process, but also the intelligence layer that increasingly mediates quality decisions. Quality-by-Intelligence reframes pharmaceutical quality as an adaptive, evidence-generating, and accountable system rather than a one-time design achievement. The transition will require new regulatory expectations, new validation evidence, new operator competencies, and stronger collaboration between industry, regulators, and academic experts. Without this shift, AI may be deployed into manufacturing faster than the quality systems needed to govern it.

Keywords Quality-by-design, Quality-by-intelligence, AI-assisted manufacturing, Pharmaceutical governance, Explainable AI, Risk-based control

*Correspondence:

Sophia Nguyen
sophia.nguyen@gmail.com

¹ Department of Pharmaceutical Technology and Drug Delivery Systems, Faculty of Pharmacy, School of Pharmaceutical Sciences, University of Glasgow, Glasgow, United Kingdom

² Department of Applied Pharmaceutical Engineering and Therapeutic Systems, Faculty of Pharmacy, National University of Singapore, Singapore, Singapore

³ Department of Pharmaceutical Manufacturing Technologies, Faculty of Medicine and Health, University of Sydney, Sydney, Australia

Introduction

Problem

Quality-by-Design has provided a powerful organising logic for pharmaceutical development by linking product understanding, process understanding, design space, and control strategy. Its success lies in replacing end-product testing as the dominant assurance mechanism with prospective scientific understanding and risk-based process control [1]. However, the rise of AI-assisted manufacturing places pressure on this model because the sources of process knowledge are no longer only mechanistic, pre-specified, or fixed before commercial operation.

The problem is not that QbD is obsolete, but that its assumptions are incomplete for systems where learning continues after deployment. Reviews of contemporary QbD practice show that the paradigm remains centred on defining critical quality attributes, critical process parameters, and design spaces through structured experimentation and prior knowledge [2, 3]. AI-assisted manufacturing challenges this logic because the model itself becomes a quality-relevant component whose performance can evolve as data, materials, equipment, and operating contexts change.

Continuous manufacturing illustrates this tension particularly clearly. Modern lines integrate process analytical technology, automated monitoring, and real-time feedback to maintain quality under conditions that are more data-rich and time-sensitive than traditional batch production [4, 5]. When artificial neural networks or other machine learning models are added to these systems, the manufacturing process becomes partly governed by statistical representations that may be accurate yet difficult to interpret, audit, or contain within a static design space [6].

Forcing AI into the existing QbD mould risks creating a false sense of control. A pre-approved design space may appear to define the boundaries of acceptable operation, while the AI layer may be making predictions or recommendations based on changing correlations that were not fully visible during validation [7]. The central governance challenge is therefore to retain the discipline of QbD while adding a lifecycle framework for intelligent, adaptive, and explainable quality decision-making.

Limits of Quality-by-Design in AI-Assisted Manufacturing

The design space is one of the most valuable but also most strained concepts in QbD. It assumes that the relationship between material attributes, process parameters, and quality outcomes can be sufficiently characterised to define a region of acceptable operation [2]. In AI-assisted manufacturing, however, the relevant operating space includes not only process variables but also data pipelines, model features, algorithmic assumptions, retraining rules, and prediction confidence, which are not easily represented as conventional multidimensional design spaces.

Traditional control strategy is similarly challenged because AI can shift control from rule-based intervention to probabilistic recommendation. Continuous manufacturing studies demonstrate that process analytical technology can support rapid detection and control of quality variation [8, 9]. Yet when AI models interpret PAT signals or recommend parameter adjustments, the control strategy must also specify how model uncertainty, model drift, sensor bias, and operator override will be governed.

Validation is the third point of friction. Conventional validation is strongest when process models and control rules are fixed, but machine learning systems may require periodic recalibration, transfer learning, or monitored adaptation to remain valid in changing environments [6, 10]. This means the validation question shifts from “Was the model validated?” to “Is the model still valid, under what conditions, and according to which evidence thresholds?”

AI-assisted manufacturing therefore exposes a mismatch between QbD’s static governance architecture and the dynamic character of algorithmic quality systems. **Table 1** contrasts the assumptions of QbD with the realities of AI-assisted manufacturing.

Table 1. Quality-by-Design versus AI-Assisted Manufacturing: Mismatched Assumptions, Gaps, and Emerging Risks

QbD assumption	AI-assisted manufacturing reality	Governance gap	Emerging quality risk
Process knowledge can be prospectively structured into a stable design space	Knowledge is partly generated through evolving data streams and model outputs	Design space does not fully include data, feature, and model behaviour	Undetected drift outside the validated knowledge boundary
Control strategy can be pre-defined around known parameters and limits	AI may recommend actions based on probabilistic predictions and changing correlations	Control logic becomes partly algorithmic and less transparent	Inappropriate parameter adjustment or delayed human intervention
Validation confirms readiness for routine operation	Model performance may change after deployment as data distributions shift	Validation becomes a lifecycle activity rather than a one-time event	Use of stale or degraded models in GMP decisions
Operators can understand process behaviour through established scientific rationale	AI outputs may be accurate but difficult to explain operationally	Human oversight may become symbolic rather than substantive	Reduced accountability for algorithm-driven decisions
Regulatory evidence is centred on process and product understanding	Evidence must also cover data integrity, model governance, explainability, and monitoring	Existing QbD submissions may under-specify the intelligence layer	Incomplete regulatory assessment of AI-enabled control

The limits of QbD are especially visible when intelligence is treated merely as a tool rather than as a governed quality subsystem. Pharma 4.0 and smart factory discussions emphasise connectivity, automation, and analytics, but these capabilities do not automatically translate into auditable quality assurance [7, 11]. A system may be technologically advanced while remaining weakly governed if its data lineage, model assumptions, retraining triggers, and decision rights are not formally embedded in the pharmaceutical quality system.

Quality-by-Intelligence Logic

Quality-by-Intelligence begins from the premise that pharmaceutical quality in AI-assisted manufacturing is maintained through governed learning, not only through pre-defined design. It preserves QbD's commitment to scientific understanding and risk management, but extends the object of governance to include the intelligence layer that predicts, recommends, detects, and sometimes controls [1, 7]. In this sense, QbI is not a replacement for QbD but a necessary evolution of it.

The core shift is from fixed assurance to adaptive assurance. In conventional QbD, quality is protected by operating within an approved design space and control strategy [2]. In QbI, quality is protected by ensuring that AI-enabled decisions remain traceable, explainable, monitored, and bounded by risk-based guardrails even as the system learns from new data [6, 12].

Figure 1 shows how Quality-by-Intelligence extends the static design-space logic of Quality-by-Design into a lifecycle governance architecture for AI-assisted pharmaceutical manufacturing.




Figure 1. From Quality-by-Design to Quality-by-Intelligence: Expanding Pharmaceutical Quality Governance for AI-Assisted Manufacturing

This logic is particularly relevant to continuous manufacturing, where high-frequency data streams and real-time control create conditions in which AI may offer genuine value. Case studies of PAT-enabled continuous manufacturing show that quality assurance can move closer to real-time process understanding [5, 8, 9]. QbI builds on this trajectory by requiring that real-time intelligence be governed with the same seriousness as equipment qualification, analytical method validation, and process control.

QbI also changes the meaning of pharmaceutical knowledge. Instead of treating knowledge as something largely captured before approval, it treats knowledge as a lifecycle asset that must be curated, challenged, updated, and explained [13]. This is a policy-relevant shift because it requires regulators and firms to define acceptable mechanisms for controlled learning without allowing uncontrolled algorithmic change.

Data and Model Governance

Data governance is the foundation of QbI because AI systems cannot generate reliable quality decisions from unreliable, biased, incomplete, or poorly contextualised data. Pharmaceutical AI models depend on sensor streams, batch records, material attributes, environmental conditions, and laboratory results, all of which must preserve provenance and integrity [6]. Without this foundation, even sophisticated models may produce outputs that are statistically impressive but GMP-fragile.

Model governance is equally important because the model becomes part of the manufacturing control environment. Lifecycle governance must define model development, feature selection, training data boundaries, performance metrics, version control, change management, retraining triggers, and retirement criteria [10]. The pharmaceutical quality system must therefore treat models as controlled assets whose status, limitations, and decision authority are continuously documented.

A further governance issue is representativeness. Machine learning models can perform well on historical development data while failing under new raw material lots, scale-up conditions, equipment states, or rare disturbances [12, 14]. This is why QbI requires monitoring for concept drift, data drift, and performance drift, not merely initial accuracy during development.

Digital twins and advanced analytics may support QbI by linking mechanistic process understanding with data-driven prediction. However, their governance value depends on whether they remain transparent enough to support investigation, deviation management, and regulatory review [13, 15]. **Table 2** summarises the key data and model governance requirements for QbI.

Table 2. Data and Model Governance for Quality-by-Intelligence: Lifecycle Management, Integrity, and Control

Governance domain	Qbi requirement	Practical implementation mechanism	Quality-system implication
Data provenance	Every data element used for training, validation, monitoring, or release support must be traceable	Metadata capture, audit trails, source-system mapping, and controlled data dictionaries	Data become regulated quality assets rather than passive records
Data integrity	AI inputs must remain accurate, complete, contemporaneous, attributable, and protected from unauthorised alteration	Validated pipelines, access controls, checksum monitoring, and periodic data-quality review	Data integrity failures become potential model and product-quality failures
Data representativeness	Training and monitoring datasets must reflect expected operating variability and known risk scenarios	Stratified sampling, raw material variability coverage, disturbance cases, and periodic dataset review	Model performance is linked to process relevance, not only statistical fit
Model lifecycle control	Models must be developed, approved, versioned, monitored, updated, and retired under formal governance	Model registry, change-control workflows, predefined retraining triggers, and independent review	Models are managed like GMP-relevant controlled components
Drift monitoring	Model performance must be assessed against changing process, material, and sensor conditions	Control charts for prediction error, population stability metrics, and alarm thresholds	Continued validity becomes an operational requirement
Human accountability	AI decisions must remain attributable to defined roles and reviewable by qualified personnel	Decision logs, operator interfaces, override rules, and escalation procedures	Human oversight becomes evidence-based rather than ceremonial

Explainable Pharmaceutical Decision-Making

Explainability is a governance requirement, not a cosmetic add-on, when AI contributes to GMP-relevant manufacturing decisions. Pharmaceutical operators must be able to understand why a model flags a deviation, predicts a critical quality attribute failure, or recommends an adjustment to a process parameter [16]. Without an interpretable rationale, AI may improve prediction while weakening accountability, because responsibility becomes distributed between model developers, quality units, operators, and automated systems.

The broader explainable AI literature distinguishes between transparent models, post hoc explanation methods, and human-centred interpretability [17]. In pharmaceutical manufacturing, this distinction matters because an explanation that satisfies a data scientist may not satisfy a production operator, quality investigator, or regulator. Qbi therefore requires explanations that are technically faithful, operationally usable, and linked to quality risk.

Methods such as feature-importance analysis, SHAP-style attribution, local surrogate explanations, rule extraction, and attention-based visualisation can help translate model behaviour into quality-relevant reasoning [18]. Yet these methods should not be accepted uncritically, because post hoc explanations may simplify or distort the real internal logic of complex models [19]. Table 3 outlines the explainability methods applicable to pharmaceutical AI and their regulatory acceptability.

Table 3. Explainable AI Methods for Pharmaceutical Manufacturing: Transparency Requirements and Decision-Making Integration

Explainability method	Pharmaceutical use case	Transparency contribution	Regulatory acceptability consideration
Feature-importance ranking	Identifying material attributes, sensor signals, or process variables driving prediction	Shows which variables most influence model output globally	Useful for model review, but insufficient for explaining individual batch decisions
Local explanation methods	Explaining why a specific batch, time point, or process state triggered an alarm	Connects a single prediction to local input patterns	More useful for deviation investigation when linked to batch records
Rule extraction	Translating complex model behaviour into simplified decision rules	Supports operator understanding and procedural alignment	Acceptability depends on whether extracted rules are faithful to the original model
Surrogate models	Using simpler models to approximate black-box decisions	Improves communication of model behaviour	Requires evidence that approximation error does not mislead quality decisions
Attention or saliency methods	Highlighting influential time-series regions, spectra, or image areas	Helps operators inspect high-dimensional inputs	Must be validated against domain knowledge to avoid visually persuasive but weak explanations
Counterfactual explanation	Showing what changes would alter a prediction or classification	Supports corrective action and process understanding	Strong potential for quality decisions if constrained by feasible manufacturing actions
Human-in-the-loop review	Combining model output with expert assessment and documented override	Preserves accountability and contextual judgement	Most acceptable when roles, escalation, and documentation are predefined

Explainability must also be embedded into operator interfaces rather than confined to validation reports. Healthcare AI scholarship shows that explainability depends on the needs of different users and decision contexts, not simply on making algorithms more transparent in an abstract sense [20]. In Qbl, the relevant question is whether an operator or quality reviewer can use the explanation to decide whether to trust, challenge, override, or escalate the AI-supported decision.

Risk-Based Control Strategy

Risk-based control under Qbl reframes AI as both an opportunity and a hazard. AI can predict critical quality attribute trajectories, identify early warning signals, and support real-time release decisions, but it can also introduce new failure modes through biased data, unstable correlations, overfitting, and automation bias [21]. A Qbl control strategy must therefore govern the risk of the model as well as the risk of the process.

In practical terms, this means that AI predictions should be linked to predefined risk categories and response pathways. For low-risk predictions, the model may support operator awareness; for medium-risk predictions, it may trigger intensified monitoring; for high-risk predictions, it may require quality review or process intervention [13]. This graded approach avoids the false binary in which AI is either fully trusted or completely excluded from GMP decision-making.

Adaptive control limits are central to QbI, but they must not become uncontrolled moving targets. Data analytics in process industries show that prediction, diagnosis, and prognosis can be integrated into increasingly sophisticated monitoring systems [14]. In pharmaceutical manufacturing, however, adaptive limits must be bounded by approved risk logic, documented model performance, and clear evidence that adaptation improves rather than erodes quality assurance.

Human oversight remains a risk-control layer, but QbI rejects purely symbolic oversight. Studies of responsible machine learning in healthcare warn that human review can fail when users do not understand model limitations or when workflows encourage passive acceptance of algorithmic outputs [21, 22]. In pharmaceutical manufacturing, meaningful oversight requires trained personnel, accessible explanations, documented override authority, and quality-system procedures that make human accountability operational.

Regulatory Translation Pathway

The regulatory pathway from QbD to QbI should be evolutionary rather than disruptive. Existing QbD concepts such as design space, control strategy, lifecycle management, and pharmaceutical quality systems remain essential, but they must be extended to include data governance, model governance, explainability, and post-deployment monitoring [1, 23]. The regulatory question is not whether AI can fit inside QbD unchanged, but what additional evidence is needed when the control logic includes intelligent systems.

Drug discovery and development AI literature shows that regulatory confidence depends on transparency, validation, reproducibility, and evidence of clinical or operational value [24-27]. Although manufacturing AI differs from clinical AI, the underlying translation challenge is similar: regulators must evaluate systems whose performance depends on data, context, model architecture, and lifecycle maintenance. QbI can make this evaluation easier by defining the evidence package for intelligence-enabled quality assurance.

Medical AI regulatory experience also offers cautionary lessons. Analyses of approved AI-based medical devices show that regulatory pathways have often struggled with model updates, real-world monitoring, and cross-jurisdictional consistency [28, 29]. Pharmaceutical manufacturing should not repeat this pattern by approving AI-enabled control systems without clear expectations for change management, drift surveillance, explanation, and periodic review.

Regulatory translation should therefore proceed through staged acceptance, beginning with advisory models and progressing toward release-supporting or control-supporting models only when evidence is sufficient.

Figure 2 maps the staged lifecycle evidence pathway through which AI-assisted manufacturing can progress from advisory analytics to governed Quality-by-Intelligence acceptance.




Figure 2. Lifecycle Evidence Pathway for Quality-by-Intelligence Acceptance in AI-Assisted Pharmaceutical Manufacturing

Table 4 maps the regulatory translation pathway from QbD to Qbl.

Table 4. Regulatory Translation Pathway: Steps, Milestones, and Evidence Needed for Quality-by-Intelligence Acceptance

Translation step	Milestone	Evidence needed	Regulatory purpose
Define AI role in the control strategy	AI classified as advisory, monitoring, release-supporting, or control-supporting	Intended use statement, risk classification, decision authority, and human oversight plan	Clarifies the regulatory significance of the model
Establish data-governance readiness	Data sources and pipelines qualified for GMP-relevant AI use	Data provenance, integrity controls, representativeness assessment, and audit trails	Demonstrates that AI inputs are reliable and traceable
Validate model performance	Model shown to perform acceptably within defined operating contexts	Accuracy, robustness, uncertainty, stress testing, and failure-mode analysis	Shows that the model is fit for its intended quality role
Demonstrate explainability	Model outputs can be interpreted by operators and quality reviewers	Explanation method validation, interface design, and investigation examples	Supports accountability and deviation management
Implement lifecycle monitoring	Continued model validity assessed during routine operation	Drift metrics, alarm thresholds, periodic review, and retraining criteria	Ensures that approval is not limited to initial deployment
Control model change	Updates managed through predefined change-control pathways	Version history, comparability evidence, retraining records, and approval workflow	Prevents uncontrolled algorithmic evolution
Scale toward Qbl acceptance	AI-enabled quality system shown to deliver equal or superior assurance	Process capability, deviation trends, intervention records, and quality outcomes	Builds confidence in intelligence-enabled quality governance

Regulatory sandboxes and pilot programmes would be especially valuable for Qbl because they allow regulators and manufacturers to learn from controlled implementation before expectations harden into formal norms. Pharma 4.0 readiness studies show that technical adoption often outpaces organisational and regulatory maturity [11, 30, 31]. A staged translation pathway would help prevent both excessive caution, which could suppress useful innovation, and premature deployment, which could undermine quality assurance.

Implementation Barriers

The first implementation barrier is cultural. Pharmaceutical manufacturing has strong reasons to value predictability, documentation, and procedural control, yet AI introduces probabilistic reasoning, model uncertainty, and adaptive learning into environments that have historically privileged fixed instructions [7]. This can generate resistance from quality units, production personnel, and senior management unless Qbl is presented not as deregulation, but as a stronger governance model for more complex systems.

The second barrier is capability. Studies of Industry 4.0 deployment in pharmaceutical and related manufacturing contexts show that successful adoption depends on data infrastructure, digital maturity, workforce competence, and cross-functional collaboration [30, 31]. QbI requires professionals who can bridge pharmaceutical science, process engineering, data science, validation, quality assurance, and regulatory strategy.

The third barrier is economic and regulatory uncertainty. Building validated data pipelines, model registries, monitoring systems, explainability interfaces, and lifecycle governance processes is costly, especially when expectations for AI in GMP manufacturing remain unsettled [15, 28]. Firms may therefore adopt AI at the margins for efficiency while avoiding deeper integration into quality decisions, producing a gap between technological possibility and governed transformation.

Conclusion

Quality-by-Design remains one of the most important achievements in modern pharmaceutical quality thinking, but it was not designed to govern intelligent systems that learn from data, mediate decisions, and may evolve after deployment. AI-assisted manufacturing exposes the limits of static design spaces, fixed validation assumptions, and control strategies that do not explicitly govern data and models. The central conclusion of this review is that QbD is necessary but insufficient for the AI era.

Quality-by-Intelligence offers a conceptual framework for closing this gap. It places data stewardship, model lifecycle control, explainability, adaptive risk management, and regulatory translation at the centre of pharmaceutical quality governance. Its purpose is not to celebrate AI, but to discipline it through accountable, auditable, and risk-based quality systems.

The transition to QbI will require collaboration across industry, regulators, standards bodies, technology providers, and academic researchers. Pilot programmes, shared terminology, model-governance expectations, and evidence standards will be needed before AI can responsibly move from peripheral analytics into core quality decision-making. The future of pharmaceutical quality should not be defined by forcing intelligence into yesterday's paradigm, but by building a governance model capable of making intelligence trustworthy.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Rights and permissions

Open Access The author(s) retain copyright. This article is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. It may be shared and adapted for non-commercial purposes with appropriate attribution, an indication of changes, and distribution of adaptations under the same license. Third-party material may be subject to separate terms identified in its credit line. View the license at <https://creativecommons.org/licenses/by-nc-sa/4.0/>.

References

- Lawrence XY, Kopcha M. The future of pharmaceutical quality and the path to get there. *Int J Pharm*. 2017;528(1-2):354-9.
- Grangeia HB, Silva C, Simoes SP, Reis MS. Quality by design in pharmaceutical manufacturing: A systematic review of current status, challenges and future perspectives. *Eur J Pharm Biopharm*. 2020;147:19-37.
- Mishra V, Thakur S, Patil A, Shukla A. Quality by design (QbD) approaches in current pharmaceutical set-up. *Expert Opin Drug Deliv*. 2018;15(8):737-58.
- Vanhoorne V, Vervaeke C. Recent progress in continuous manufacturing of oral solid dosage forms. *Int J Pharm*. 2020;579:119194.
- Vargas JM, Nielsen S, Cárdenas V, Gonzalez A, Aymat EY, Almodovar E, et al. Process analytical technology in continuous manufacturing of a commercial pharmaceutical product. *Int J Pharm*. 2018;538(1-2):167-78.
- Nagy B, Galata DL, Farkas A, Nagy ZK. Application of artificial neural networks in the process analytical technology of pharmaceutical manufacturing—a review. *AAPS J*. 2022;24(4):74.
- Arden NS, Fisher AC, Tyner K, Yu LX, Lee SL, Kopcha M. Industry 4.0 for pharmaceutical manufacturing: Preparing for the smart factories of the future. *Int J Pharm*. 2021;602:120554.
- Pauli V, Roggo Y, Pellegatti L, Trung NQ, Elbaz F, Ensslin S, et al. Process analytical technology for continuous manufacturing tableting processing: A case study. *J Pharm Biomed Anal*. 2019;162:101-11.
- Roggo Y, Pauli V, Jelsch M, Pellegatti L, Elbaz F, Ensslin S, et al. Continuous manufacturing process monitoring of pharmaceutical solid dosage form: A case study. *J Pharm Biomed Anal*. 2020;179:112971.
- Nagy B, Szabados-Nacsá Á, Fülöp G, Nagyne AT, Galata DL, Farkas A, et al. Interpretable artificial neural networks for retrospective QbD of pharmaceutical tablet manufacturing based on a pilot-scale developmental dataset. *Int J Pharm*. 2023;633:122620.
- Malheiro V, Duarte J, Veiga F, Mascarenhas-Melo F. Exploiting pharma 4.0 technologies in the non-biological complex drugs manufacturing: innovations and implications. *Pharmaceutics*. 2023;15(11):2545.
- Shang C, You F. Data analytics and machine learning for smart process manufacturing: Recent advances and perspectives in the big data era. *Engineering*. 2019;5(6):1010-6.
- Reis MS, Gins G. Industrial process monitoring in the big data/industry 4.0 era: From detection, to diagnosis, to prognosis. *Processes*. 2017;5(3):35.
- Ge Z, Song Z, Ding SX, Huang B. Data mining and analytics in the process industry: The role of machine learning. *IEEE Access*. 2017;5:20590-616.
- Domokos A, Nagy B, Szilágyi B, Marosi G, Nagy ZK. Integrated continuous pharmaceutical technologies—a review. *Org Process Res Dev*. 2021;25(4):721-39.
- Jiménez-Luna J, Grisoni F, Schneider G. Drug discovery with explainable artificial intelligence. *Nat Mach Intell*. 2020;2(10):573-84.

Barredo Arrieta A, Díaz-Rodríguez N, Del Ser J, Bennetot A, Tabik S, Barbado A, et al. Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Inf Fusion*. 2020;58:82-115.

Adadi A, Berrada M. Peeking inside the black-box: a survey on explainable artificial intelligence (XAI). *IEEE Access*. 2018;6:52138-60.

Holzinger A, Langs G, Denk H, Zatloukal K, Müller H. Causability and explainability of artificial intelligence in medicine. *WIREs Data Min Knowl Discov*. 2019;9(4):e1312.

Amann J, Blasimme A, Vayena E, Frey D, Madai VI, Precise4Q Consortium. Explainability for artificial intelligence in healthcare: a multidisciplinary perspective. *BMC Med Inform Decis Mak*. 2020;20(1):310.

Wiens J, Saria S, Sendak M, Ghassemi M, Liu VX, Doshi-Velez F, et al. Do no harm: a roadmap for responsible machine learning for health care. *Nat Med*. 2019;25(9):1337-40.

Kelly CJ, Karthikesalingam A, Suleyman M, Corrado G, King D. Key challenges for delivering clinical impact with artificial intelligence. *BMC Med*. 2019;17(1):195.

Lee SH, Kim JK, Jee JP, Jang DJ, Park YJ, Kim JE. Quality by Design (QbD) application for the pharmaceutical development process. *J Pharm Investig*. 2022;52(6):649-82.

Vamathevan J, Clark D, Czodrowski P, Dunham I, Ferran E, Lee G, et al. Applications of machine learning in drug discovery and development. *Nat Rev Drug Discov*. 2019;18(6):463-77.

Mak KK, Pichika MR. Artificial intelligence in drug development: present status and future prospects. *Drug Discov Today*. 2019;24(3):773-80.

Mak KK, Wong YH, Pichika MR. Artificial intelligence in drug discovery and development. In: Hock FJ, Gralinski MR, editors. *Drug discovery and evaluation: safety and pharmacokinetic assays*. 2nd ed. Cham: Springer; 2024. p. 1461-98.

Dara S, Dhamercherla S, Jadav SS, Babu CM, Ahsan MJ. Machine learning in drug discovery: a review. *Artif Intell Rev*. 2022;55(3):1947-99.

Benjamens S, Dhunoo P, Meskó B. The state of artificial intelligence-based FDA-approved medical devices and algorithms: an online database. *NPJ Digit Med*. 2020;3(1):118.

Muehlematter UJ, Daniore P, Vokinger KN. Approval of artificial intelligence and machine learning-based medical devices in the USA and Europe (2015–20): a comparative analysis. *Lancet Digit Health*. 2021;3(3):e195-203.

Sharma D, Patel P, Shah M. A comprehensive study on Industry 4.0 in the pharmaceutical industry for sustainable development. *Environ Sci Pollut Res*. 2023;30(39):90088-98.

McDermott O, Wojcik AM, Trubetskaya A, Sony M, Antony J, Kharub M. Pharma industry 4.0 deployment and readiness: a case study within a manufacturer. *TQM J*. 2024;36(9):456-76.