



EXPLAINABLE MODELS FOR REGULATORY APPROVAL DELAY USING DEFICIENCY TEXT AND MANUFACTURING READINESS SIGNALS

Gabriel Costa^{1*}, Rafael Mendes¹, Bruno Teixeira², Lucas Ribeiro¹

1. *Department of AI Drug Discovery Systems, Faculty of Pharmacy, Federal University of Rio de Janeiro, Rio de Janeiro, Brazil.*
2. *Department of Computational Pharmaceutical Analytics, Faculty of Engineering, University of Campinas, Campinas, Brazil.*

ARTICLE INFO

Received:

06 February 2026

Received in revised form:

17 June 2026

Accepted:

18 June 2026

Available online:

28 June 2026

Keywords: Explainable artificial intelligence, Regulatory science, Approval delay, Complete response letter, Manufacturing readiness, SHAP

ABSTRACT

Pharmaceutical regulatory approval delays often occur when unresolved chemistry, manufacturing, and controls deficiencies intersect with weak manufacturing readiness, yet agency correspondence detailing these issues is rarely leveraged as predictive evidence before final regulatory action. Currently, sponsors assess submission risk primarily through expert judgment, historical experience, and qualitative escalation processes, leaving an unmet need for an interpretable modeling framework capable of forecasting delays while pinpointing specific, remediable drivers of risk. This article proposes explainable machine learning models that integrate text-derived deficiency features with structured manufacturing readiness indicators to predict regulatory approval delays in a transparent, auditable, and actionable manner for regulatory affairs and compliance teams. By using natural language processing to encode deficiency topics and severity, and employing SHAP explanations to decompose predicted delay risk into individual feature contributions, a conceptual gradient-boosted classifier or survival model could, for example, identify a high-risk submission where delays are driven by aseptic process validation concerns and an unfavorable recent inspection outcome. Linking predictions directly to deficiency language and operational readiness signals, rather than providing an opaque risk score, would enable sponsors to prioritize targeted remediation, allocate resources effectively, and improve overall submission quality.

This is an open-access article distributed under the terms of the [Creative Commons Attribution-Non Commercial-Share Alike 4.0 License](https://creativecommons.org/licenses/by-nc-sa/4.0/), which allows others to remix, and build upon the work non commercially.

To Cite This Article: Costa G, Mendes R, Teixeira B, Ribeiro L. Explainable Models for Regulatory Approval Delay Using Deficiency Text and Manufacturing Readiness Signals. *Pharmacophore*. 2026;17(3):102-11. <https://doi.org/10.51847/MB1GO4VJox>

Introduction

Regulatory approval delays can impose strategic and financial burdens on pharmaceutical sponsors because they postpone market access, disrupt launch planning, and require additional regulatory-response work. Prediction studies in drug development suggest that submission and approval outcomes can be modeled from historical product, sponsor, and development features [1], while oncology approval analyses indicate that approval decisions reflect structured patterns that may be learnable from prior regulatory experience [2]. Manufacturing-related uncertainty is especially important because pharmaceutical quality expectations shape whether a submitted product can be reliably manufactured at commercial scale [3]. Continuous manufacturing submission audits further show that manufacturing strategy and regulatory outcome are intertwined, supporting the idea that delay risk may be partly foreseeable from CMC and facility-readiness evidence [4].

Current regulatory risk assessment remains heavily dependent on expert interpretation, cross-functional review, and sponsor-specific institutional memory. Statistical and machine-learning studies of drug approval have shown that structured variables can support approval prediction [5], yet these approaches rarely represent the full texture of regulatory correspondence or operational readiness evidence. Quantitative benefit–risk modeling has also been proposed as a way to make regulatory product assessments more systematic [6], but it does not by itself explain how specific deficiencies or site-readiness gaps contribute to delay. An explainable regulatory–delay model would therefore complement expert judgment rather than replace it, translating heterogeneous evidence into a transparent risk assessment.

Corresponding Author: Gabriel Costa; Department of AI Drug Discovery Systems, Faculty of Pharmacy, Federal University of Rio de Janeiro, Rio de Janeiro, Brazil. E-mail: gabriel.costa@gmail.com.

Natural language processing creates a practical route for converting regulatory correspondence into analyzable features. Text-mining studies of clinical-trial termination demonstrate that free-text documents can be transformed into structured signals for prediction [7], and topic-modeling approaches show how latent themes in regulatory or development text may reveal patterns associated with outcomes [8]. Explainable artificial intelligence provides methods for connecting such extracted features to model outputs, with SHAP offering local and global attribution mechanisms for tree-based models. In this setting, the same logic could be applied to deficiency text so that a predicted delay is explained through specific CMC topics, severity language, and manufacturing readiness indicators.

The central thesis of this article is that an explainable model could integrate deficiency text with manufacturing readiness signals to predict regulatory approval delay in a way that is auditable by sponsors and understandable to regulatory professionals. Prior drug-approval modeling challenges have shown the value of combining diverse product and development features for regulatory outcome prediction [1], while broader reviews of machine learning in drug development highlight the importance of transparent and fit-for-purpose applications. Pharmaceutical manufacturing research emphasizes that digital and quality-system transformation should support reliable regulatory decision-making rather than produce isolated analytics [9]. A model designed for explainability from the outset would therefore provide not only a risk estimate, but also a structured account of the evidence behind that estimate.

Background

The Pharmaceutical Regulatory Approval Process and Sources of Delay

The pharmaceutical approval process requires agencies to determine whether submitted evidence supports safety, efficacy, and quality, while sponsors must respond to unresolved questions that arise during review. Public-disclosure studies of FDA filing decisions and refuse-to-file correspondence show that regulatory communication can contain substantive information about deficiencies and review barriers [10, 11]. Comparative regulatory analyses across agencies also indicate that approval timing and decision patterns vary by product type, therapeutic area, and review system [12]. For delay modeling, Complete Response Letters and deficiency correspondence should therefore be treated as structured regulatory evidence sources whose language may reveal unresolved CMC, clinical, or administrative issues.

Manufacturing Readiness and Pre-Approval Inspections

Manufacturing readiness reflects whether the proposed facilities, processes, controls, and quality systems are prepared to support commercial supply under GMP expectations. FDA pharmaceutical quality perspectives emphasize that regulatory confidence depends on robust quality systems and reliable manufacturing performance [3], while manufacturing-benchmarking studies show that quality practices vary across global production sites [13]. Data-integrity challenges further demonstrate how compliance weaknesses can create regulatory concern even when product-specific development data appear otherwise mature [14]. A delay-prediction model should therefore encode inspection history, Warning Letters, OAI classifications, and related compliance signals as readiness indicators that may influence review outcomes.

Table 1 shows the key manufacturing readiness dimensions, associated regulatory/compliance indicators, and their relevance for predicting regulatory review outcomes in pharmaceutical development.

Table 1. Manufacturing Readiness Indicators and Compliance Signals for Regulatory Assessment Models

Readiness Dimension	Key Indicators / Signals	Typical Data Sources	Relevance to Regulatory Assessment
Facility Preparedness	GMP certification status, facility qualification, inspection outcomes	FDA/EMA inspection records, internal QA audits	Indicates whether sites can consistently operate under GMP expectations
Quality Management System (QMS)	Deviation rates, CAPA effectiveness, audit findings	Quality system databases, audit reports	Reflects robustness of quality oversight and risk control
Regulatory Compliance History	FDA Warning Letters, OAI classification, 483 observations	FDA inspection database, regulatory correspondence	Strong predictor of regulatory concern and approval risk
Data Integrity Controls	ALCOA+ compliance, audit trail completeness, electronic record integrity	QA audits, IT validation reports	Identifies risk of unreliable or non-traceable manufacturing data
Process Performance & Control	Process capability (Cp/Cpk), batch failure rates, variability metrics	Manufacturing execution systems (MES), QC systems	Demonstrates stability and reproducibility of commercial-scale production
Global Manufacturing Consistency	Site-to-site variability, tech transfer success rates	Global supply chain records, tech transfer reports	Highlights variability across manufacturing networks
Inspection Risk Signals	Frequency of inspections, repeat observations, escalation patterns	Regulatory inspection history	Aggregated indicator of regulatory scrutiny and risk exposure

Natural Language Processing of Regulatory Text

Natural language processing can convert unstructured regulatory correspondence into structured variables such as deficiency categories, cited controls, severity indicators, and recurrence patterns. Text-mining approaches have already been used to

quantify risk from narrative development documents [7], while latent topic models have shown how recurring themes can be extracted from clinical-trial text and associated with termination risk [8]. In regulatory deficiency analysis, similar methods could identify recurring CMC topics such as sterility assurance, stability, analytical validation, process control, or data integrity. These extracted features would allow the model to learn from the content and tone of agency concerns rather than relying only on binary submission outcomes.

Machine Learning in Regulatory Science

Machine learning has been applied to drug approval prediction, clinical-trial outcome forecasting, and early-development transition analysis. Studies using multimodal artificial intelligence have shown that target choice, trial design, and related development features can support prediction of clinical-trial outcomes [15], while chemical and development features have been used to estimate drug approval likelihood [16]. Oncology-specific work has similarly explored the prediction of approval after early-phase trials [17], and phase-transition studies have used machine-learning frameworks to identify indicators associated with development progression [18]. These studies support the plausibility of regulatory predictive analytics, but they also reveal the need for models that incorporate regulatory text and provide explanations suitable for high-stakes compliance use.

Explainable AI for High-Stakes Decision-Making

Explainability is essential when model outputs may influence regulatory strategy, remediation priorities, or executive governance. Reviews of explainable artificial intelligence emphasize that transparency, interpretability, and accountability are central requirements in high-stakes domains, while interpretable machine-learning scholarship distinguishes between post hoc explanations, transparent models, and explanation fidelity. SHAP-based methods are particularly relevant because they can provide both submission-level explanations and global summaries of feature influence for tree-based models. In regulatory delay prediction, such explanations would be expected to show whether the forecast is driven by deficiency language, manufacturing history, product complexity, or combinations of these factors.

Model Development Overview

High-Level Predictive Pipeline

A high-level predictive pipeline would represent each regulatory submission as a combined feature vector containing text-derived deficiency signals, structured manufacturing readiness indicators, and product-level descriptors. Prior approval-prediction studies demonstrate that diverse features can be assembled into models of regulatory outcome [1, 5], while multimodal clinical-development prediction supports the principle that heterogeneous evidence can improve decision support [15]. The output would be a predicted probability of a significant approval delay and a conceptual estimate of additional review time, expressed with appropriate caution rather than as a deterministic forecast. The explanation layer would then attribute the prediction to specific input features so that regulatory teams can understand why the submission is considered high risk.

Table 2 shows the key components of a predictive regulatory submission pipeline, including text-derived deficiency signals, manufacturing readiness indicators, product-level descriptors, and the corresponding probabilistic and interpretable outputs used to estimate approval delay risk.

Table 2. Predictive features and outputs in a regulatory submission delay prediction pipeline

Component	Feature Type	Description	Example Inputs	Role in Model
Text-derived deficiency signals	Unstructured text features	Signals extracted from regulatory review narratives and deficiency letters	Reviewer comments, inspection findings, prior approval critiques	Capture qualitative risk indicators linked to historical regulatory concerns
Manufacturing readiness indicators	Structured operational features	Measures of GMP readiness, process robustness, and site compliance	Inspection history, deviation rates, quality system maturity scores	Quantifies manufacturing-related delay risk factors
Product-level descriptors	Structured product features	Intrinsic attributes of the drug product and modality	Dosage form, complexity, modality type, stability profile	Accounts for inherent product-driven approval complexity
Integrated submission vector	Multimodal feature fusion	Combined representation of all feature groups	Concatenated text + structured variables	Serves as unified model input for prediction
Prediction output	Probabilistic + temporal estimate	Risk and delay estimation outputs	Probability of significant delay; estimated review time extension	Supports decision-making with uncertainty-aware forecasting
Explanation layer	Feature attribution outputs	Model interpretability module highlighting key drivers	Feature importance scores, attention weights, SHAP values	Provides transparency for regulatory review justification

Core Input Features

Core deficiency features would include the presence of sterility, stability, analytical-method, process-validation, labeling, and data-integrity concerns, as well as topic proportions and severity language extracted from correspondence. Manufacturing readiness features would encode recent Warning Letters, OAI classifications, inspection outcomes, time since last inspection, and facility-level compliance history, reflecting the importance of quality-system maturity in pharmaceutical manufacturing [13, 14]. Product-specific features could include dosage form, novelty, therapeutic area, orphan designation, and whether the submission concerns an original or supplemental application, consistent with evidence that product and development attributes influence approval likelihood [16, 19]. These feature groups would allow the model to distinguish between a delay driven mainly by textual deficiency content and one driven by manufacturing readiness or product complexity.

Design Principles

The model should be explainable by default, auditable by regulatory affairs teams, and robust to missing or incomplete data. Explainable AI taxonomies stress that interpretability must be aligned with the user's decision context, and SHAP-based tree explanations provide a practical mechanism for translating complex models into local feature contributions. The system should also support agency-specific adaptation because comparative regulatory studies show that decision patterns and approval timelines can differ across FDA, EMA, Swissmedic, and other review environments [12, 20]. A well-designed framework would therefore preserve traceability from raw evidence to engineered features, prediction, explanation, and subsequent human action.

Data Sources and Feature Engineering

Construction of a Submission-Outcome Dataset

A submission-outcome dataset would conceptually link original and supplemental NDA or BLA records to review timelines, regulatory actions, and eventual approval outcomes. Public data sources such as FDA approval records and filing disclosures have been used to study regulatory application transparency and agency decision communication [10, 11], while cross-agency comparisons demonstrate that approval timing can be studied using curated regulatory records [20]. Proprietary sponsor data could add internal milestones, deficiency correspondence, response-cycle dates, and manufacturing site mappings that are not fully available in public sources. Dataset construction would therefore require careful provenance tracking so that each modeled feature can be traced to a regulatory record, internal document, or facility-compliance source.

NLP Feature Extraction from Deficiency Correspondence

NLP feature extraction would parse deficiency correspondence into structured concepts, including deficiency category, severity language, affected module, referenced manufacturing process, and recurrence across review cycles. Topic-modeling and text-mining studies show that narrative documents can be transformed into predictive representations without requiring manual coding of every document [7, 8]. For regulatory text, named entity recognition could identify references to stability protocols, validation batches, analytical procedures, sterility controls, and data-integrity terms. The resulting feature set would preserve interpretable links between model inputs and the original deficiency language, allowing explanations to point back to specific text segments rather than abstract embeddings alone.

Encoding Manufacturing Readiness and Compliance Signals

Manufacturing readiness features would map each proposed facility to recent inspection outcomes, Warning Letters, data-integrity findings, quality metrics, and the time elapsed since prior regulatory review. Pharmaceutical quality literature emphasizes that manufacturing capability and quality-system performance are central to regulatory confidence [3], while work on Industry 4.0 in pharmaceutical manufacturing argues that digital readiness and process monitoring can support smarter quality oversight [9]. Regulatory audits of continuous manufacturing submissions further illustrate how manufacturing strategy can be linked to regulatory outcomes and review considerations [4]. These signals should be engineered as transparent variables, such as recent unfavorable inspection status or unresolved compliance action, so that both predictions and explanations remain operationally meaningful.

Table 3 maps each regulatory and manufacturing evidence source to interpretable engineered features, explanatory value, and remediation relevance within the proposed delay-prediction framework.

Table 3. Regulatory Evidence-to-Feature Mapping for Explainable Approval-Delay Prediction

Evidence domain	Raw evidence examples	Engineered feature representation	Regulatory meaning	Expected explanatory role	Practical remediation use
Deficiency topic content	Complete Response Letter language; agency deficiency correspondence; CMC review comments	Binary or proportional indicators for sterility, stability, analytical validation, process validation, labeling, and data-integrity topics	Identifies the substantive regulatory issues most closely associated with unresolved review concerns	Shows which deficiency themes increase or reduce predicted delay risk	Directs technical teams toward the CMC issues requiring additional evidence, clarification, or corrective action
Deficiency severity language	Terms indicating major concern, inadequate	Ordinal severity score; weighted language	Distinguishes minor clarification requests	Explains whether risk is driven by the	Helps prioritize severe or repeated deficiencies

	evidence, unresolved issue, repeat observation, or insufficient validation	intensity index; recurrence marker across review cycles	from deficiencies likely to delay approval	presence of a topic or by the strength and recurrence of agency concern	over lower-impact wording
Regulatory module location	References to Module 3, stability sections, analytical methods, validation reports, facility information, or labeling documents	Module-specific feature flags; affected-document category; section-level traceability label	Locates the deficiency within the submission structure	Links model attribution to specific submission components rather than abstract text categories	Supports targeted document revision, response preparation, and ownership assignment
Manufacturing-site compliance history	Prior inspection outcomes; Warning Letters; OAI classification; unresolved compliance actions	Recent unfavorable inspection flag; time since inspection; unresolved action indicator; compliance-history score	Represents whether the proposed manufacturing site is likely to raise review confidence concerns	Explains delay risk arising from facility readiness rather than product-specific text alone	Triggers site-readiness review, quality-system escalation, or alternative facility planning
Quality-system maturity	Data-integrity observations; CAPA history; deviation trends; process-monitoring maturity; quality metrics	Quality-system maturity index; data-integrity risk flag; CAPA closure status	Captures broader operational reliability and evidence credibility	Shows whether technical deficiencies are compounded by quality-system weakness	Supports cross-functional remediation beyond the regulatory response document
Product and review complexity	Dosage form; biologic versus small molecule; novel modality; orphan status; original versus supplemental submission	Product-type category; review-pathway variable; formulation or modality complexity indicator	Accounts for baseline differences in regulatory review difficulty	Separates intrinsic product complexity from remediable submission or manufacturing weaknesses	Prevents over-attributing delay risk to deficiencies when complexity is partly structural
Timing and missingness signals	Absence of correspondence by a milestone; incomplete facility mapping; pending inspection status; missing internal response dates	Missingness indicators; review-stage flags; observation-window variables	Prevents missing data from being treated as neutral or meaningless	Clarifies whether uncertainty is caused by unavailable evidence, early review status, or incomplete internal records	Encourages data-completeness checks before interpreting risk predictions
Outcome and review-cycle history	Approval timing; response-cycle duration; resubmission date; final regulatory action	Delay label; time-to-event endpoint; response-cycle outcome variable	Defines the prediction target and supports time-aware validation	Provides the historical outcome structure against which features are learned	Enables prospective benchmarking and post-action calibration

Explainable Model Architecture

Model Choice – Gradient-Boosted Trees with SHAP

Gradient-boosted tree models would be suitable for this conceptual application because they can handle mixed categorical, ordinal, binary, and continuous features while supporting post hoc SHAP explanations. SHAP methods provide local explanations that allocate the difference between a baseline prediction and a submission-specific prediction across individual features. Interpretable machine-learning literature also emphasizes that model choice should balance predictive flexibility with explanation fidelity and governance needs. In a prospective regulatory setting, the model should be trained and evaluated using time-aware splits so that future submissions are not implicitly informed by later regulatory outcomes.

Input Feature Vector and Pre-processing

The input feature vector would combine normalized text-derived deficiency variables, structured manufacturing readiness indicators, and product-specific descriptors. Drug-approval prediction work using molecular, clinical, patent, and development features supports the broader principle of integrating multiple data domains into a unified predictive representation [16], while multimodal approval and trial-outcome studies show that heterogeneous features can capture complementary sources of regulatory risk [1, 15]. Missing deficiency data should not be treated as a simple blank field, because the absence of a letter before a given milestone may itself carry timing information. Instead, missingness indicators should be encoded explicitly so that the model can learn whether unavailable correspondence reflects low risk, incomplete observation, or early review status.

Output: Delay Probability and Explanation

The model output would be a predicted probability that a submission experiences a meaningful regulatory delay, accompanied by an explanation that decomposes the prediction into deficiency and readiness contributions. SHAP waterfall explanations could show, for example, that a stability-program concern increases predicted delay while a clean recent inspection lowers it, following the general logic of local feature attribution. Regulatory product-assessment frameworks highlight the importance

of making decision evidence explicit and interpretable for stakeholders [6], and explainable AI reviews emphasize that explanations should support human oversight rather than merely decorate model outputs. The architecture would therefore present the prediction together with traceable evidence links, enabling regulatory teams to inspect the underlying deficiency text and manufacturing signals behind the forecast.

Figure 1 presents the proposed explainable regulatory-delay prediction architecture linking deficiency correspondence, manufacturing readiness signals, SHAP-based attribution, and human-governed remediation actions.

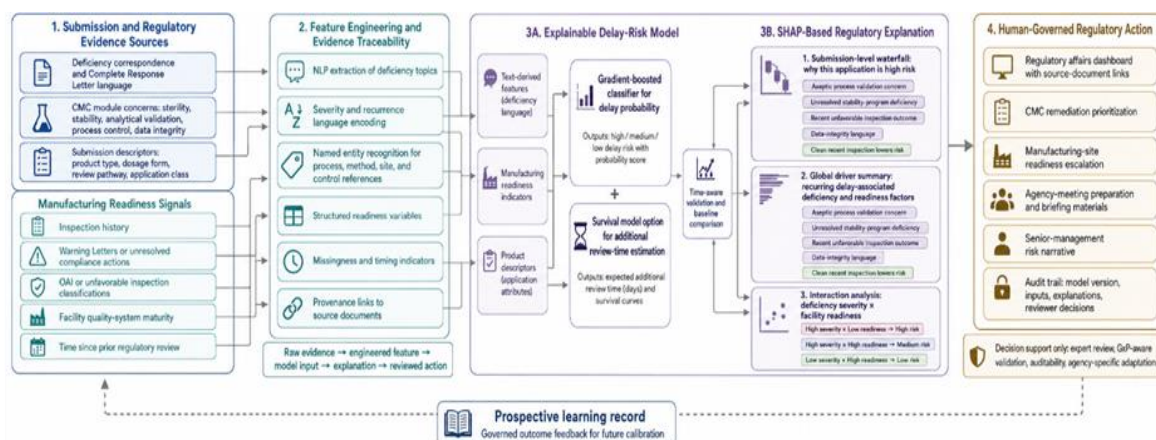


Figure 1. Explainable Regulatory Delay Prediction Architecture Integrating Deficiency Text and Manufacturing Readiness Signals.

Identifying Key Drivers of Regulatory Delay

Global Explanation of Regulatory Delay Factors

Global explanation would summarize which deficiency topics and manufacturing readiness signals are most consistently associated with predicted approval delay across submissions. SHAP summary views could conceptually identify sterility assurance, stability-program concerns, data-integrity language, and recent unfavorable inspection outcomes as recurring high-influence features, using the same attribution logic that supports global interpretation in tree-based explainable AI. Pharmaceutical quality work suggests that such drivers should be interpreted in the context of quality-system maturity and manufacturing reliability rather than as isolated textual labels [3]. A dynamic view of GMP compliance would further support the idea that prior facility behavior and current deficiency content should be analyzed together when assessing regulatory readiness [21].

Submission-Level Explanation

At the submission level, the model would provide a local explanation showing why one application is expected to face higher delay risk than another. For example, a SHAP-style waterfall could attribute elevated risk to a major stability-program deficiency, unresolved manufacturing-site compliance concerns, and a product profile that has historically required complex review, while also showing any features that reduce risk. Drug approval prediction studies using product and development features demonstrate that application-specific characteristics can contribute meaningfully to outcome forecasts [16, 22]. The explanation should therefore be expressed as a traceable regulatory narrative, linking the forecast to the precise deficiency topics and facility-readiness signals present in the submission record.

Interaction between Deficiency Severity and Manufacturing Readiness

The model should also be designed to capture interactions between the severity of deficiency language and the strength of manufacturing readiness evidence. A deficiency involving process validation may carry different delay implications when the manufacturing site has a recent unfavorable compliance history than when the same topic appears at a site with a strong inspection record. Data-integrity and quality-practice studies show that compliance weaknesses can compound technical concerns by raising broader questions about the reliability of submitted evidence [13, 14]. SHAP interaction analysis would be useful conceptually because it could show whether the combination of a deficiency topic and a readiness signal contributes more strongly than either feature alone.

Counterfactual Analysis for Risk Mitigation

Counterfactual analysis would translate explanation into action by estimating how the predicted delay risk might change if a remediable factor were addressed before final agency action. For instance, the model could conceptually compare the current submission profile with a scenario in which a Warning Letter has been resolved, additional validation evidence has been submitted, or a facility-readiness signal has improved. Regulatory manufacturing literature emphasizes continuous quality improvement and the future direction of pharmaceutical quality systems, making such scenario analysis relevant to remediation

planning [3, 21]. The goal would not be to claim causality, but to help teams prioritize corrective actions that are most aligned with the model's explanation.

Explainability for Regulatory Affairs and Compliance Teams

Dashboard for Submission Risk Assessment

A submission risk dashboard would present the predicted delay risk together with the most influential deficiency topics, manufacturing readiness signals, and links to source documents. Explainable AI guidance stresses that explanations should be matched to the user's decision context, so a regulatory affairs dashboard should emphasize traceability, plain-language interpretation, and evidence review rather than technical model internals. SHAP-based explanation displays could provide ranked feature contributions while allowing users to inspect the underlying deficiency text that generated each feature. This interface would help regulatory strategists distinguish between delay risk driven by correctable CMC gaps and risk driven by broader product or review complexity.

Narrative Explanation for Senior Management

For senior management, the system should generate a concise narrative that explains the predicted delay in operational and regulatory terms. Such a narrative might state that the submission appears high risk because agency correspondence indicates unresolved aseptic validation concerns and the proposed facility has recent compliance signals requiring remediation. Prior work on interpretable machine learning emphasizes that explanations must be understandable to the intended audience, not merely mathematically faithful. In pharmaceutical manufacturing settings, this narrative framing would support governance decisions about resource allocation, escalation, and readiness review, especially as AI becomes more embedded in manufacturing and quality operations [23].

Audit Trail and Model Governance

Every prediction should be stored with the model version, input data provenance, extracted text features, structured readiness variables, explanation output, and user actions taken after review. This audit trail is important because regulatory and compliance decisions require accountability, especially when model outputs influence submission planning or remediation priorities. Explainable AI frameworks emphasize transparency and accountability as core requirements for trustworthy use in high-stakes domains, while data-integrity research in pharmaceutical manufacturing underscores the importance of reliable, attributable, and reviewable records [14]. Model governance should therefore treat each explanation as part of a controlled decision-support record rather than as an informal analytical note.

Table 4 defines the governance and decision-use requirements needed to convert explainable delay prediction into auditable regulatory affairs support rather than an opaque risk-scoring exercise.

Table 4. Explainability, Governance, and Decision-Use Framework for Regulatory Affairs Implementation

Decision-support layer	Key output	Primary users	Interpretability requirement	Governance safeguard	Example decision supported	Failure mode to monitor
Submission-level risk prediction	Predicted probability of meaningful approval delay; optional estimated additional review time	Regulatory affairs leads; submission strategy teams	Output must be framed as probabilistic decision support, not as a deterministic forecast	Require expert review before any change to submission strategy	Identify whether a submission requires intensified pre-action remediation	Overconfidence in a numerical risk score without reviewing evidence drivers
Local SHAP explanation	Ranked feature contributions for one submission	Regulatory affairs, CMC, quality, compliance teams	Each driver must be traceable to a deficiency topic, text segment, readiness signal, or product descriptor	Preserve source-document links and model version for each explanation	Determine whether delay risk is driven by aseptic validation, stability, inspection status, or data integrity	Technically correct explanations that are too vague for operational action
Global driver analysis	Cross-submission ranking of recurring delay-associated features	Portfolio governance teams; regulatory intelligence groups	Global patterns must be interpreted as associations, not universal regulatory rules	Review for agency, product-type, and time-period stratification	Identify recurring CMC or manufacturing-readiness weaknesses across a portfolio	Treating historical correlations as fixed future agency behavior
Interaction analysis	Combined effect of deficiency severity and manufacturing readiness	CMC governance; quality leadership; site-readiness teams	Interaction findings must show whether combined evidence increases risk beyond either factor alone	Require domain review before using interaction outputs for escalation	Escalate a process-validation deficiency when paired with recent unfavorable inspection history	Misreading interaction effects as causal proof

Counterfactual scenario analysis	Estimated risk change under plausible remediation scenarios	Regulatory response teams; manufacturing leadership	Scenarios must be labeled as hypothetical and non-causal	Restrict scenarios to feasible, documented, and remediable factors	Compare the likely value of resolving a Warning Letter versus adding validation evidence	Presenting counterfactuals as guaranteed delay reduction
Regulatory affairs dashboard	Risk score, top drivers, evidence links, and action status	Regulatory affairs and cross-functional submission teams	Interface must prioritize plain-language explanation and document traceability	Log user review, comments, overrides, and action decisions	Coordinate response planning and assign issue ownership	Dashboard use without checking the underlying correspondence
Senior-management narrative	Concise explanation of risk, drivers, and recommended escalation	Executive governance; launch planning; portfolio leadership	Narrative must translate technical drivers into regulatory and operational consequences	Include uncertainty statement and human-review status	Decide whether to allocate additional CMC, quality, or regulatory resources	Oversimplifying complex regulatory uncertainty for executive reporting
Audit trail and lifecycle monitoring	Stored prediction record, inputs, features, explanations, actions, and outcomes	Model governance, quality assurance, compliance, audit teams	Every prediction must be reproducible from governed inputs and model version	Maintain controlled records and periodic performance review	Support inspection readiness, internal audit, and model recalibration	Loss of traceability between raw evidence, features, explanation, and action

Feedback Loop from Post-Approval Reviews

After final regulatory action, the actual review pathway should be fed back into the model-development process to refine feature definitions, calibration, and explanation quality. Clinical-trial termination studies show that historical outcomes can inform predictive modeling when text and structured variables are curated consistently [7, 24]. Regulatory approval prediction challenges also demonstrate that retrospective outcome data can support learning systems when the problem is carefully defined and evaluated [1]. In this framework, feedback should be governed so that updates improve future decision support without overwriting the auditability of earlier predictions and explanations.

Integration Into Submission Planning And Risk Management

Pre-Submission Risk Assessment and Resource Allocation

During submission preparation, the model could be used to identify high-risk CMC topics, weak manufacturing readiness signals, and facility-compliance concerns before the application is filed or before a major response is submitted. Prior work on pharmaceutical quality and smart manufacturing suggests that readiness assessment should be integrated into routine quality and regulatory planning rather than treated as a late-stage review activity [3, 9]. Machine-learning approaches to development and approval prediction indicate that structured historical evidence can support earlier prioritization of risk [15, 18]. In practice, the explanation would help teams allocate additional scientific, manufacturing, or regulatory resources to the issues most likely to drive delay.

Supporting Regulatory Interactions and Meeting Requests

The model's explanations could also support planning for agency interactions by identifying which deficiency categories or manufacturing readiness gaps require clarification, additional evidence, or formal discussion. Public analyses of FDA filing and refuse-to-file communications show that agency correspondence can contain actionable information about the reasons an application cannot proceed smoothly [10, 11]. Benefit-risk and regulatory-decision frameworks similarly indicate that transparent evidence organization can improve the quality of regulatory dialogue [6]. An explainable delay forecast would therefore be useful for preparing pre-submission meetings, structuring briefing materials, and prioritizing responses to the most consequential agency concerns.

Evaluation Strategy

Predictive Performance

Predictive performance should be evaluated conceptually through a temporally separated validation strategy that tests whether the model generalizes to later submissions rather than memorizing historical review patterns. Prior drug-approval prediction studies emphasize the importance of defining outcomes carefully and comparing models against simpler baselines [1, 5]. Clinical-trial termination and early-trial prediction work also shows that predictive models should be assessed against realistic decision points rather than only retrospective convenience samples [24, 25]. In this article's proposed setting, evaluation should compare text-plus-readiness models with structured-only models while avoiding unsupported claims about achieved performance.

Explanation Quality and Actionability

Explanation quality should be evaluated by regulatory affairs, CMC, quality, and compliance experts who review whether the model's explanations are plausible, specific, and actionable. Interpretable machine-learning literature cautions that technically valid explanations may still fail if users cannot connect them to decisions, and explainable AI reviews emphasize the importance of human-centered evaluation in high-stakes domains. Expert review should therefore examine whether highlighted deficiency topics correspond to meaningful regulatory concerns and whether suggested drivers align with known submission risks. The evaluation should also test whether explanations remain stable enough to support governance when small changes in input wording or missingness occur.

Prospective Pilot in a Regulatory Affairs Department

A prospective pilot would deploy the model as decision support within a regulatory affairs department while preserving human authority over submission strategy. Applications of machine learning in drug discovery and development show that AI systems are most useful when embedded into domain workflows and interpreted by experts rather than treated as autonomous decision-makers. Pharmaceutical manufacturing AI discussions similarly stress the need for GxP-aware implementation, governance, and validation when analytics influence quality-related decisions [23]. A pilot should therefore focus on workflow integration, explanation usefulness, escalation behavior, and whether teams can act on the model's risk narratives before formal agency action.

Limitations

Confounding and Causal Interpretation

The proposed model would identify associations between deficiency language, manufacturing readiness signals, and regulatory delay, but it should not be interpreted as estimating causal effects. A sterility-related deficiency may appear to drive predicted delay because it reflects deeper quality-system weaknesses, complex product attributes, or review-cycle timing that the model cannot fully separate. Dynamic GMP compliance models reinforce the need to view compliance status as an evolving organizational signal rather than a single causal variable [21]. For this reason, counterfactual outputs should be presented as decision-support scenarios, not as guarantees that resolving one feature would eliminate delay.

Data Availability and Agency-Specific Variability

Data availability would be a major limitation because many deficiency letters, internal response documents, facility mappings, and quality metrics are not consistently public or standardized. Comparative analyses show that regulatory decision patterns differ across agencies and time periods, meaning that a model trained in one jurisdiction may not transfer directly to another [12, 20]. Public FDA filing and correspondence studies demonstrate that some regulatory information can be analyzed, but they also imply that public disclosure is incomplete relative to the full review record [10, 11]. Agency-specific vocabulary, inspection classification systems, and product-type differences would therefore require careful adaptation before the model could be used across review environments.

Conclusion

Explainable models for regulatory approval delay could connect two forms of evidence that are often reviewed separately: the language of regulatory deficiencies and the operational signals of manufacturing readiness. By representing both sources in a single modeling framework, such systems could forecast delay risk while preserving a transparent link to the underlying submission evidence.

The principal strength of this approach is that it would make the reasons for a predicted delay visible to the teams responsible for reducing that risk. Instead of producing an opaque score, the model would identify the specific deficiency themes, readiness gaps, and compliance signals most responsible for the forecast.

Important challenges remain, especially around data access, cross-agency portability, confounding, and the need for prospective validation in real regulatory workflows. These challenges mean the model should be treated as decision support for expert teams, not as a substitute for regulatory judgment or scientific review.

Future progress would benefit from multi-company collaboration, shared regulatory-AI benchmarks, and carefully governed datasets that link submission outcomes with deficiency text and manufacturing readiness evidence. Such collaboration could strengthen predictive regulatory science while keeping transparency, auditability, and submission quality at the center of model development.

Acknowledgments: None

Conflict of interest: None

Financial support: None

Ethics statement: None

References

1. Siah KW, Kelley NW, Ballerstedt S, Holzhauer B, Lyu T, Mettler D, et al. Predicting drug approvals: The Novartis data science and artificial intelligence challenge. *Patterns*. 2021;2(8):100312.
2. Watanabe T, Nemoto S, Kato H, Isomura T. Prediction of drug approvals in oncology using explainable artificial intelligence. *Sci Rep*. 2025;15(1):12345.
3. Lawrence XY, Kopcha M. The future of pharmaceutical quality and the path to get there. *Int J Pharm*. 2017;528(1-2):354-9.
4. Fisher AC, Liu W, Schick A, Ramanadham M, Chatterjee S, Brykman R, et al. An audit of pharmaceutical continuous manufacturing regulatory submissions and outcomes in the US. *Int J Pharm*. 2022;622:121778.
5. Lo AW, Siah KW, Wong CH. Machine learning with statistical imputation for predicting drug approvals. SSRN; 2019.
6. Sun S, Heske S, Mercadel M, Wimmer J. Predicting regulatory product approvals using a proposed quantitative version of fda's benefit-risk framework to calculate net-benefit score and benefit-risk ratio. *Ther Innov Regul Sci*. 2021;55(1):129-37.
7. Follett L, Geletta S, Laugerman M. Quantifying risk associated with clinical trial termination: A text mining approach. *Inf Process Manag*. 2019;56(3):516-25.
8. Geletta S, Follett L, Laugerman M. Latent Dirichlet Allocation in predicting clinical trial terminations. *BMC Med Inform Decis Mak*. 2019;19(1):242.
9. 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.
10. Chahal HS, Szeto D, Chaudhry AH, Sigelman DW, Kim S, Lurie PG. Public disclosure of the filing of new drug and therapeutic biologics applications with the US Food and Drug Administration. *JAMA Intern Med*. 2019;179(8):1144-6.
11. Chahal HS, Mukherjee S, Sigelman DW, Temple R. Contents of US Food and Drug Administration refuse-to-file letters for new drug applications and efficacy supplements and their public disclosure by applicants. *JAMA Intern Med*. 2021;181(4):522-9.
12. Rohr UP, Iovino M, Rudofsky L, Li Q, Juritz S, Girceys A, et al. A decade comparison of regulatory decision patterns for oncology products to all other non-oncology products among Swissmedic, European Medicines Agency, and US Food and Drug Administration. *Clin Transl Sci*. 2023;16(9):1569-81.
13. Fellows M, Friedli T, Li Y, Maguire J, Rakala N, Ritz M, et al. Benchmarking the quality practices of global pharmaceutical manufacturing to advance supply chain resilience. *AAPS J*. 2022;24(6):111.
14. Charoo NA, Khan MA, Rahman Z. Data integrity issues in pharmaceutical industry: Common observations, challenges and mitigations strategies. *Int J Pharm*. 2023;631:122503.
15. Aliper A, Kudrin R, Polykovskiy D, Kamya P, Tutubalina E, Chen S, et al. Prediction of clinical trials outcomes based on target choice and clinical trial design with multi-modal artificial intelligence. *Clin Pharmacol Ther*. 2023;114(5):972-80.
16. Ciray F, Doğan T. Machine learning-based prediction of drug approvals using molecular, physicochemical, clinical trial, and patent-related features. *Expert Opin Drug Discov*. 2022;17(12):1425-41.
17. Beinse G, Tellier V, Charvet V, Deutsch E, Borget I, Massard C, et al. Prediction of drug approval after phase I clinical trials in oncology: RESOLVED2. *JCO Clin Cancer Inform*. 2019:1-10.
18. Feijoo F, Palopoli M, Bernstein J, Siddiqui S, Albright TE. Key indicators of phase transition for clinical trials through machine learning. *Drug Discov Today*. 2020;25(2):414-21.
19. Cho C, Lee S, Bang D, Piao Y, Kim S. ChemAP: predicting drug approval with chemical structures before clinical trial phase by leveraging multi-modal embedding space and knowledge distillation. *Sci Rep*. 2024;14(1):23010.
20. Lau F, Seifert R. Comparison of drug approvals of the FDA and EMA between 2013 and 2023. *Naunyn Schmiedeberg's Arch Pharmacol*. 2026;399(1):279-99.
21. Bao Y, Buhay N, Zheng Q. A dynamic model for gmp compliance and regulatory science. *J Pharm Innov*. 2024;19(3):25.
22. Park M, Kim D, Kim I, Im SH, Kim S. Drug approval prediction based on the discrepancy in gene perturbation effects between cells and humans. *EBioMedicine*. 2023;94:104725.
23. Nakka G, Gupta S, Reddy KK. Artificial intelligence in pharmaceutical manufacturing: Applications, case studies, and GxP implementation considerations. *J Pharm Sci*. 2026;104240.
24. Elkin ME, Zhu X. Predictive modeling of clinical trial terminations using feature engineering and embedding learning. *Sci Rep*. 2021;11(1):3446.
25. Kavalci E, Hartshorn A. Improving clinical trial design using interpretable machine learning based prediction of early trial termination. *Sci Rep*. 2023;13(1):121.