



MULTIMODAL AI COPILOT FOR FORMULATION DEVELOPMENT USING PROTOCOLS, EXCIPIENT DATA, AND DISSOLUTION CURVES

Mohammed Al-Farsi^{1*}, Salim Al-Harthy¹, Nasser Al-Rawahi²

1. *Department of Pharmaceutical AI Systems, Faculty of Pharmacy, Sultan Qaboos University, Muscat, Oman.*
2. *Department of Computational Drug Sciences, Faculty of Engineering, German University of Technology in Oman, Muscat, Oman.*

ARTICLE INFO

Received:

01 November 2025

Received in revised form:

21 January 2026

Accepted:

26 January 2026

Available online:

28 February 2026

Keywords: Multimodal AI, Pharmaceutical formulation, AI copilot, Excipient databases, Dissolution curves, Retrieval-augmented generation

ABSTRACT

Formulation development remains a knowledge-intensive activity in which scientific judgment is distributed across protocols, excipient knowledge, experimental records, and dissolution interpretation. AI assistance could help organize this complexity by connecting evidence sources that are normally reviewed separately. Formulators often move manually between protocol folders, spreadsheet-based excipient records, and dissolution analysis tools. This fragmented workflow can make it difficult to identify relevant precedents, compare similar formulations, and explain why a formulation failed to meet a release target. A multimodal AI copilot could ingest internal development protocols, excipient property databases, and dissolution curves to support natural-language formulation queries. Such a system would not replace the scientist but would help retrieve evidence, suggest formulation adjustments, and generate rationale for expert review. The proposed copilot includes a document-retrieval module for protocols and development reports, an excipient-property knowledge graph, a dissolution-curve encoder, a multimodal reasoning engine, and a conversational interface. Together, these modules would allow the system to connect text, structured formulation attributes, and release-profile behavior. By providing traceable, evidence-based responses, the copilot would be expected to reduce cognitive load during formulation design and troubleshooting. It could also help preserve institutional knowledge by turning historical development experience into searchable, reusable evidence. A formulation AI copilot could support a shift from experience-based trial-and-error toward data-driven, hypothesis-guided formulation development. Its value would depend on careful validation, human oversight, and integration into regulated pharmaceutical workflows.

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: Al-Farsi M, Al-Harthy S, Al-Rawahi N. Multimodal AI Copilot for Formulation Development Using Protocols, Excipient Data, and Dissolution Curves. *Pharmacophore*. 2026;17(1):43-52. <https://doi.org/10.51847/ckHNOFgliz>

Introduction

Pharmaceutical formulation development is a complex and time-intensive process in which each design decision draws on physicochemical properties, excipient behavior, processing constraints, and product-performance targets. Machine-learning approaches have already been explored for predicting tablet breaking force, disintegration, and other formulation outcomes, showing that computational models can assist with formulation reasoning when relevant variables are represented appropriately [1]. Neural-network methods for orally disintegrating tablets further indicate that formulation composition and process descriptors can be mapped to performance expectations in a way that supports early design decisions [2]. However, these models typically operate as task-specific predictors rather than as interactive assistants that integrate the broader evidence landscape surrounding a formulation project.

Current formulation workflows are often fragmented because protocols may reside in shared document repositories, excipient properties may be tracked in spreadsheets or external references, and dissolution curves may be analyzed in separate software environments. Predictive dissolution work using spectroscopy and process data illustrates the value of connecting experimental measurements with release behavior, yet the knowledge remains difficult to reuse when the corresponding formulation rationale is locked in reports or batch records [3]. Process Analytical Technology-based dissolution modeling also demonstrates that release interpretation can benefit from integrated data streams, but such models do not automatically retrieve

Corresponding Author: Al-Farsi; Department of Pharmaceutical AI Systems, Faculty of Pharmacy, Sultan Qaboos University, Muscat, Oman. E-mail: mohammed.alfarsi@gmail.com.

the protocol context or excipient reasoning that scientists need for decision-making [4]. A copilot architecture would therefore address not only prediction but also the practical synthesis of dispersed formulation evidence.

The emergence of multimodal AI creates an opportunity to reason across text, tables, and time-series data in a single user-facing system. Multiparticulate release modeling and 3D-printability prediction show that formulation variables, process conditions, and product-performance descriptors can be encoded computationally, while broader machine-learning-directed formulation development has framed AI as a tool for navigating large design spaces [5-7]. In this context, retrieval-augmented generation would serve as a bridge between historical formulation knowledge and generative dialogue, allowing the system to answer questions with evidence rather than relying on unsupported language generation. The conceptual advance is not a new experimental result but a system design that places multimodal learning inside the daily reasoning loop of formulation scientists.

This EAI article proposes a multimodal AI copilot that integrates development protocols, excipient data, and dissolution curves to provide interactive, evidence-backed formulation guidance. The system would combine structured formulation descriptors, excipient-function knowledge, and dissolution-profile representations with document retrieval so that recommendations remain grounded in traceable sources [8]. Drug–excipient compatibility systems and real-time dissolution surrogate models suggest that individual components of such a copilot are technically plausible, although their integration into a conversational decision-support environment remains an open development challenge [9, 10]. The proposed copilot is therefore presented as an emerging architecture for human-supervised formulation intelligence rather than an autonomous formulation engine.

Background

The Formulation Development Process and Decision Points

Solid oral dosage-form development typically proceeds from pre-formulation assessment to prototype design, optimization, scale-up, and control-strategy refinement. At each stage, scientists interpret drug properties, excipient functions, process feasibility, manufacturability, and release behavior, often revisiting earlier decisions when dissolution or mechanical performance is inadequate. Machine-learning studies on 3D-printed drug delivery systems and artificial neural networks for formulation optimization show that formulation development can be represented as a structured decision process, but these approaches still require expert interpretation of when and how predictions should inform experimental planning [11, 12]. A copilot would be most useful at these decision points, where evidence from prior experiments, protocol constraints, and performance data must be considered together.

Excipient Knowledge and Databases

Excipient knowledge includes functionality, grade-specific behavior, compatibility, regulatory acceptability, manufacturability, and interactions with the active ingredient and process route. Although structured excipient information is essential for design, it is often underused because formulation scientists must manually translate database entries into practical design options. Reviews of AI technologies for solid dosage forms and models that use formulation descriptors for particle or product-performance prediction suggest that excipient attributes can be encoded as structured variables, but they also highlight the need for meaningful representation of material function rather than simple ingredient labels [13, 14]. An excipient-property knowledge graph could therefore make functionality, compatibility, and precedent use accessible to a copilot in a form that supports reasoning.

Dissolution Testing and Data Interpretation

Dissolution testing is central to solid oral dosage-form development because the shape of a release profile reflects formulation structure, excipient behavior, manufacturing effects, and potential biopharmaceutical performance. Dissolution curves can be compared through regulatory and scientific tools such as similarity metrics, but practical interpretation often requires judging whether deviations arise from disintegration, diffusion, matrix erosion, coating behavior, or process variability. Raman mapping-based dissolution prediction and active machine-learning approaches for formulation design show that dissolution behavior can be linked to material and process signals, supporting the idea that release profiles should be treated as informative model inputs rather than only final quality checks [15, 16]. A copilot could use curve-shape similarity to retrieve analogous historical cases and support scientifically grounded troubleshooting.

AI in Pharmaceutical Development

Existing AI work in pharmaceutical development has largely focused on single-task models, including printability prediction, dissolution forecasting, formulation classification, compatibility evaluation, and image-based quality-control support. Multimodal data have already been used to predict production outcomes for 3D-printed drug products, indicating that formulation systems can benefit from combining different evidence types when the relationships among materials, processes, and product attributes are complex [17]. Studies on desired release-profile design and machine-vision quality assessment further show that AI can assist both prospective formulation design and product evaluation, but these systems are not typically embedded as conversational tools for scientists [18, 19]. The gap is therefore not only algorithmic performance but the absence of an integrated, query-driven environment that helps formulators reason across the full development record.

Retrieval-Augmented Generation and Copilots in Specialized Domains

Retrieval-augmented generation and specialized copilots are conceptually attractive for pharmaceutical formulation because regulated scientific decisions require source-grounded answers, domain constraints, and transparent rationale. In formulation contexts, the same principle would mean retrieving protocol sections, previous batch outcomes, excipient compatibility evidence, and dissolution-profile analogues before generating a response. Machine-learning approaches for drug–excipient compatibility and BCS-related formulation prediction show that formulation knowledge can be organized into computationally searchable forms, which could be paired with language models for controlled question answering [20, 21]. The key design requirement is that the copilot should generate recommendations only after grounding them in retrieved evidence and presenting that evidence for human review.

Copilot System Architecture Overview

High-Level Copilot Design

At a high level, the copilot would begin with a natural-language question from the formulator, such as a request to explain a slow dissolution profile or identify alternative disintegrants for an immediate-release tablet. The system would parse the request, route it to relevant data modules, retrieve protocol and formulation-history evidence, compare dissolution behavior, and generate a response containing both narrative guidance and links to source materials. Non-linear dissolution-profile modeling and comparative data-driven release prediction provide conceptual support for treating release behavior as a searchable and predictable signal within such a workflow [22, 23]. The resulting system would function as an evidence-synthesis layer rather than as a standalone model that makes unreviewed formulation decisions.

Figure 1 presents the proposed multimodal AI copilot architecture for connecting development protocols, excipient knowledge, formulation descriptors, and dissolution-curve histories into traceable, human-reviewed formulation guidance.

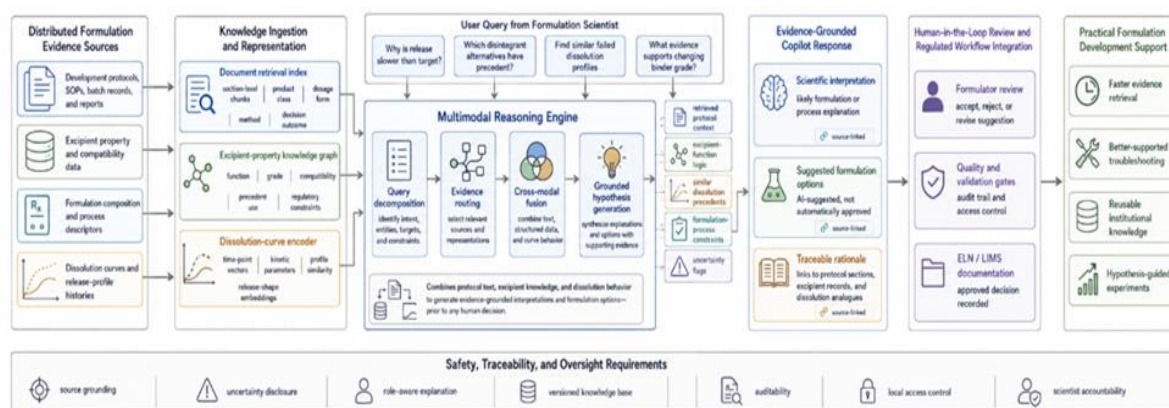


Figure 1. Multimodal AI Copilot Architecture for Evidence-Grounded Formulation Development

Core Data Modules

The proposed architecture includes three core data modules: a protocol and knowledge base, an excipient-property resource, and a dissolution-curve library connected to formulation metadata. The protocol module would support semantic search over development reports, SOPs, batch records, and lab notebooks, while the excipient module would represent functions, grades, compatibility considerations, and known use contexts. Quantitative structure–dissolution profile modeling and drug-release prediction from tablet formulation variables indicate that both molecular and formulation descriptors can be linked to release behavior, supporting the inclusion of a curve library with associated formulation attributes [24, 25]. Together, these modules would allow the copilot to answer questions by combining procedural evidence, material knowledge, and performance history.

Design Principles

The copilot should be user-centric, source-traceable, privacy-preserving, and adaptable to different product classes and organizational knowledge structures. User centrality means the formulator remains responsible for interpreting suggestions, while traceability means that each recommendation must be connected to retrieved protocol text, excipient entries, or dissolution records. Updated drug–excipient compatibility systems and integrated pre-formulation prediction frameworks illustrate the importance of domain-specific knowledge organization when AI is used to support pharmaceutical decisions [26, 27]. For regulated development environments, the architecture should also support local deployment, access controls, audit trails, and clear separation between AI-suggested options and approved scientific decisions.

Knowledge Ingestion: Protocols, Excipients Data, and Dissolution Curves

Ingesting and Indexing Development Protocols and Reports

The protocol-ingestion layer would convert SOPs, batch records, development reports, formulation rationales, and troubleshooting notes into searchable document chunks with section-level metadata. These chunks should preserve context such as product class, dosage form, manufacturing route, excipient system, dissolution method, and decision outcome, because formulation guidance depends on both what was done and why it was done. A structured oral formulation database for machine

learning illustrates the value of organizing formulation knowledge into reusable records, while a digital formulator and tableting data-factory concept points toward development environments where formulation data are continuously generated, indexed, and reused [28]. The copilot would use this indexed evidence to retrieve relevant precedent before generating any formulation suggestion.

Structuring Excipient Property Data

Excipient-property ingestion would map internal specifications, compendial descriptions, supplier information, and formulation-history annotations into a unified schema. Such a schema should represent functional categories, grade-dependent properties, compatibility risks, typical formulation roles, and links to products or batches where each excipient has been used. Earlier models for tablet performance and orally disintegrating tablet formulation prediction show that excipient-related variables can contribute to computational guidance, but a copilot requires richer semantic representation so that a binder, disintegrant, filler, or lubricant is interpreted according to its functional role rather than only as a categorical input [1, 2]. This structured representation would allow the user to ask practical questions, such as which excipient alternatives have similar function but lower compatibility concern.

Encoding Dissolution Curves

Dissolution curves may be encoded as time-point vectors, fitted kinetic parameters, similarity-metric descriptors, or learned embeddings that capture profile shape and release dynamics. These representations would support curve retrieval, clustering of similar release behavior, and conceptual forecasting of how changes in formulation composition might affect the profile. Studies using NIR-based dissolution prediction and PAT-based dissolution testing demonstrate that release profiles can be linked to measurement and formulation variables in ways that support model-assisted interpretation [3, 4]. In the copilot, a dissolution encoder would not replace formal dissolution assessment but would help identify historical analogues and formulate hypotheses about the cause of release deviations.

Table 1 defines the multimodal evidence architecture required for a formulation copilot to reason across protocols, excipient records, formulation descriptors, dissolution profiles, and expert-reviewed decision records.

Table 1. Multimodal Evidence Architecture for Formulation Copilot Reasoning

| Evidence modality | Primary formulation content | Computational representation | Formulation reasoning function | Typical user question supported | Traceability requirement |
|---|--|---|--|---|---|
| Development protocols and SOPs | Experimental procedures, dissolution methods, formulation constraints, acceptance criteria, decision rationales | Section-level semantic chunks with metadata for product class, dosage form, process route, dissolution method, and decision outcome | Retrieves procedural context and prevents recommendations that conflict with prior protocol constraints | “Which previous protocol used a similar dissolution method for this dosage form?” | Response must link to the specific protocol section, method description, or development report passage used to support the answer |
| Batch records and development reports | Actual formulation compositions, process parameters, deviations, failed trials, troubleshooting notes, and final decisions | Structured formulation records linked to narrative report sections and outcome labels | Identifies comparable historical cases and distinguishes intended formulation design from actual manufacturing execution | “Have we seen this failure pattern after changing compression force or granulation conditions?” | Response must separate documented batch evidence from inferred formulation interpretation |
| Excipient property data | Function, grade, particle size, viscosity, compatibility, regulatory acceptability, supplier attributes, and prior use | Excipient-property knowledge graph connecting material attributes, functional roles, compatibility risks, and product precedents | Supports material substitution reasoning and explains why an excipient may affect dissolution, manufacturability, or stability | “Which alternative disintegrants have similar function but fewer compatibility concerns?” | Response must cite the excipient record, grade-specific attribute, or compatibility note used in the suggestion |
| Formulation composition and process descriptors | API load, excipient ratios, binder level, lubricant level, manufacturing route, compression parameters, coating conditions | Structured tabular features and formulation-process vectors linked to batch and performance outcomes | Enables comparison of formulation variants and helps identify design variables associated with release behavior | “Which composition differences separate passing and failing prototypes?” | Response must identify whether the comparison is based on composition, process, or both |
| Dissolution curves | Time-point release data, curve shape, incomplete release, lag phase, burst release, similarity to target profile | Time-series vectors, fitted kinetic parameters, similarity metrics, clusters, or learned release-profile embeddings | Retrieves analogous release profiles and supports hypotheses about disintegration, diffusion, matrix effects, coating behavior, or process variability | “Which historical batches had the most similar slow-release pattern?” | Response must show which profiles were considered analogous and whether similarity is empirical, model-derived, or qualitative |

| | | | | | |
|-----------------------------------|--|---|---|--|---|
| User query and dialogue history | Natural-language formulation questions, follow-up constraints, role-specific information needs, and decision context | Intent classification, query decomposition, sub-task routing, and conversation-state representation | Determines which evidence modules should be consulted before a response is generated | “Explain this result for a junior scientist, but show the evidence sources.” | Response must preserve the evidence base while adapting only the level of explanation |
| Expert review and decision record | Human assessment, accepted or rejected AI suggestions, rationale for final action, and experimental follow-up | Review status, audit trail, versioned decision log, and ELN/LIMS-linked record | Converts AI output into accountable formulation knowledge without allowing autonomous decision-making | “Which AI suggestions were accepted in this project and why?” | Response must distinguish AI-suggested options from scientist-approved decisions |

Multimodal Fusion and Reasoning Engine

Query Decomposition and Routing

The reasoning engine would first decompose a user query into sub-tasks that determine whether protocol retrieval, excipient lookup, dissolution-curve comparison, or compatibility assessment is needed. For example, a question about delayed release after changing a binder grade would require retrieval of similar protocols, structured comparison of binder properties, and curve-shape analysis across relevant historical batches. Machine-learning-directed formulation development and AI-based solid dosage-form reviews both emphasize that formulation decisions involve linked variables rather than isolated features, supporting a routing strategy that coordinates several evidence modules [7, 13]. The routing layer would therefore act as a planner that decides which sources are needed before any answer is generated.

Cross-Modal Reasoning

Cross-modal reasoning would combine retrieved text snippets, structured excipient records, formulation composition, process descriptors, and dissolution embeddings into a coherent interpretation of the user’s formulation question. A multimodal model or ensemble could identify whether a dissolution issue is more consistent with excipient incompatibility, insufficient disintegration, matrix effects, process variability, or mismatch between method and formulation design. Evidence from multiparticulate release modeling, multimodal 3D-printing prediction, and release-profile design studies suggests that combining formulation variables with performance data can support more useful predictions than treating each data type separately [5, 17, 18]. In a copilot, the goal would be conceptual synthesis and hypothesis generation rather than autonomous selection of a final formulation.

Response Generation with Citations

The final response-generation layer would translate the fused evidence into a concise conversational answer that explains the likely issue, proposes options for expert consideration, and links each claim to its source. For example, a recommendation to evaluate an alternative disintegrant would cite the protocol section showing previous use, the excipient-property record describing functionality, and the dissolution curves of similar formulations. Drug–excipient compatibility prediction and formulation-database approaches show why source-grounded evidence is essential when AI suggestions influence material selection or formulation strategy [20, 28]. The language model should therefore operate under strict grounding rules, generating only claims supported by retrieved evidence and clearly marking all recommendations as AI-suggested until reviewed by a formulator.

User Interaction and Conversational Formulation Assistance

Conversational Interface Design for Formulators

The conversational interface would allow formulators to interact with the system through practical scientific questions rather than predefined database queries. A user could ask why a formulation showed incomplete release, request similar historical dissolution profiles, or ask which excipient substitutions have precedent in comparable products. Reviews of artificial neural networks for formulation optimization and machine-learning applications in solid oral dosage-form development indicate that AI outputs must be translated into interpretable scientific guidance if they are to support real formulation work [8, 12]. The interface should therefore support multi-turn dialogue in which the user can challenge a suggestion, request the underlying evidence, or narrow the answer to a specific dosage form, process route, or excipient class.

Table 2 shows key functionalities and capabilities that a conversational interface should provide to support practical formulation queries.

Table 2. Features of a Conversational Interface for Formulation Support

| Functionality | Description / Use Case | Scientific Rationale / Supporting Evidence |
|--|--|---|
| Query on incomplete release | Users can ask why a formulation did not achieve expected dissolution | Allows targeted investigation of formulation performance; aligns with ANNs and ML approaches for formulation optimization [8, 12] |
| Retrieval of historical dissolution profiles | Users can request similar past formulations and their release profiles | Supports data-driven decision making by referencing precedent outcomes |

| | | |
|---------------------------------|--|---|
| Excipient substitution guidance | Users can ask which excipients have precedent in similar products | Enables formulation adjustments with historical evidence for safety and efficacy |
| Multi-turn dialogue | Users can challenge AI suggestions, request underlying evidence, or narrow answers to specific dosage forms, processes, or excipient classes | Ensures AI outputs are interpretable, actionable, and relevant to real-world formulation work [8, 12] |

Proactive Suggestions and What-If Experimentation

A formulation copilot could also provide proactive suggestions when the evidence base indicates possible incompatibility, release-risk patterns, or inconsistency between protocol design and target product profile. In a what-if mode, the user could ask how changing a disintegrant grade, binder level, or matrix former might be expected to influence dissolution behavior, with the system responding conceptually from retrieved precedents and model-based analogues. Compatibility-prediction systems and release-profile modeling studies support this kind of hypothesis generation, provided that the output is framed as an option for expert evaluation rather than as a validated experimental outcome [9, 22, 25]. This conversational sandbox would help scientists explore alternatives before laboratory work while preserving the need for experimental confirmation.

Personalization to User Role and Experience

The copilot should adapt its explanations to the user's role, experience level, and immediate decision context. A junior formulator may need a fuller explanation of why a superdisintegrant or lubricant affects release, whereas a senior scientist may prefer a compact summary of comparable cases and caveats. AI systems developed for formulation classification, product-performance prediction, and quality-control support show that technical outputs can be useful only when presented in a form that matches the decision being made [19, 21, 24]. Role-aware personalization would therefore adjust the level of detail without altering the evidence base, ensuring that the same source-grounded recommendation remains available for scientific and regulatory scrutiny.

Explainability, Traceability, and Human-In-The-Loop Oversight

Source-Grounded Explainability

Source-grounded explainability is essential because formulation recommendations can affect material selection, process design, and the interpretation of product-performance failures. Each answer should state which evidence supports it, such as a protocol precedent, an excipient-property entry, a compatibility warning, or a set of analogous dissolution profiles. Dissolution surrogate modeling and Raman-based dissolution prediction show that AI can connect formulation or analytical signals to release behavior, but those connections must be communicated in a way that scientists can inspect and challenge [10, 15]. The copilot should therefore explain not only what it suggests but why the retrieved evidence makes that suggestion plausible.

Quality Gates and Expert Review

Human-in-the-loop oversight would require every recommendation to remain advisory until reviewed and accepted by an accountable formulation scientist. The system should distinguish low-risk exploratory guidance from decisions that affect pivotal batches, regulatory commitments, or approved material specifications. Integrated pre-formulation prediction frameworks and digital formulator concepts suggest that AI may become embedded in development workflows, but regulated pharmaceutical use requires review gates, auditability, and clear ownership of final decisions [27]. The copilot should therefore preserve expert authority by recording the evidence reviewed, the AI-suggested options, and the human decision that followed.

Integration Into Pharmaceutical Development Workflow

Embedding in ELN and LIMS

The copilot would be most useful if embedded directly into electronic lab notebooks, laboratory information management systems, and formulation data environments rather than operating as a separate destination. During experiment planning, a sidebar or chat panel could retrieve relevant protocol precedents, excipient constraints, and dissolution analogues without forcing the user to leave the workflow. Studies on data-driven release prediction and structured oral formulation databases indicate that formulation knowledge becomes more reusable when experimental records are represented consistently and connected to performance outcomes [23, 28]. Integration into existing systems would also support access control, traceability, and alignment with established documentation practices.

Table 3 shows the potential benefits of embedding a copilot directly into laboratory data environments.

Table 3. Benefits of Integrating a Copilot into Laboratory Data Systems

| Integration Point | Functional Advantage | Example Use Case |
|--|---|--|
| Electronic Lab Notebooks (ELNs) | Immediate retrieval of relevant protocols and prior experiment data | Sidebar or chat panel suggesting excipient constraints or dissolution analogues during experiment planning |
| Laboratory Information Management Systems (LIMS) | Streamlined workflow and traceability | Automatic linking of experiment metadata with performance outcomes |
| Formulation Data Environments | Enhanced knowledge reuse and predictive modeling | Consistent representation of experimental records for structured oral formulation databases |

| | | |
|-------------------------|---|---|
| System-wide Integration | Alignment with documentation standards and access control | Ensures secure, compliant data sharing across teams |
|-------------------------|---|---|

Knowledge Retention and Continuous Learning

As new experiments are documented, the copilot’s knowledge base could continuously index protocols, formulation compositions, analytical results, and dissolution curves for future retrieval. This would help preserve institutional knowledge that is otherwise dispersed across individual scientists, project folders, and historical reports. Machine-learning-directed formulation development and active-learning approaches for precision formulation both support the broader principle that new experimental evidence should be fed back into the formulation knowledge cycle [7, 16]. Continuous learning in this setting should be carefully governed, with new records curated and versioned before they are used to support future recommendations.

Evaluation Strategy

Answer Quality and Accuracy

Evaluation should begin with expert review of the copilot’s answers to realistic formulation queries covering excipient selection, dissolution troubleshooting, protocol interpretation, and formulation comparison. Experts would judge whether the answer is relevant, scientifically plausible, complete enough for decision support, and correctly grounded in the cited source material. Prior work on drug–excipient compatibility prediction and expert-system design shows that formulation AI should be assessed not merely by output generation but by the usefulness and reliability of its domain-specific reasoning [9, 26]. The evaluation should also test whether the copilot refuses or qualifies answers when the retrieved evidence is weak, contradictory, or outside its validated knowledge base.

Table 4 provides a validation and governance framework for determining whether the formulation copilot is reliable, traceable, workflow-ready, and appropriately constrained by human oversight.

Table 4. Validation, Governance, and Implementation Readiness Framework for a Formulation AI Copilot

| Evaluation or governance domain | Core validation question | Recommended assessment approach | Evidence of readiness | Failure mode addressed | Human oversight requirement |
|-----------------------------------|---|---|--|---|--|
| Retrieval fidelity | Does the copilot retrieve the most relevant protocol, batch, excipient, and dissolution evidence for the user’s question? | Expert-curated benchmark queries; top-k retrieval review; comparison against manual search by formulation scientists | High relevance of retrieved sources across protocols, excipient records, and dissolution histories | Irrelevant or incomplete evidence leading to weak formulation guidance | Scientist must be able to inspect retrieved sources before acting on the answer |
| Source-grounded response accuracy | Are generated explanations and suggestions fully supported by retrieved evidence? | Expert review of answer-source alignment; unsupported-claim detection; hallucination audit | Recommendations consistently linked to specific protocol sections, excipient records, or dissolution analogues | Plausible but unsupported AI-generated formulation claims | Unsupported recommendations must be blocked, qualified, or marked as speculative |
| Cross-modal reasoning quality | Does the system correctly combine text, structured excipient data, formulation composition, and dissolution-curve behavior? | Case-based evaluation using known formulation problems; review of whether the correct evidence modules were routed and integrated | Accurate synthesis of procedural, material, process, and release-profile evidence | Overreliance on one modality while ignoring conflicting evidence from another | Formulator must confirm whether the synthesized explanation is scientifically plausible |
| Dissolution analogue reliability | Are similar dissolution profiles retrieved and interpreted appropriately? | Curve-similarity benchmarking; expert assessment of analogue relevance; sensitivity analysis across similarity metrics | Retrieved analogues reflect meaningful release-pattern similarity and relevant formulation context | Misleading comparison between curves with similar shape but different mechanism or method | Similarity results must be presented with method, formulation context, and uncertainty |
| Excipient recommendation safety | Are suggested substitutions or adjustments compatible with function, grade, process, and regulatory constraints? | Rule-based checks against excipient knowledge graph; compatibility review; formulation scientist assessment | Suggestions are constrained by excipient role, compatibility history, and prior formulation use | Unsafe or inappropriate material substitution based only on superficial similarity | AI suggestions must remain advisory until reviewed by an accountable scientist |
| User-centered decision support | Does the copilot reduce manual search burden while improving confidence in evidence review? | Controlled workflow study comparing formulation tasks with and without the copilot; user confidence and | Faster retrieval of relevant evidence and better-documented rationale without replacing expert judgment | Tool adoption failure due to poor workflow fit or unclear usefulness | Users must retain control over follow-up questions, evidence inspection, and final decisions |

| evidence-completeness scoring | | | | | |
|---|---|---|---|--|--|
| Auditability and regulated workflow fit | Can the system document what evidence was retrieved, what answer was generated, and what human decision followed? | Audit-trail review; ELN/LIMS integration testing; version-control inspection | Complete record of query, retrieved evidence, generated response, reviewer action, and final decision | Untraceable AI influence on regulated formulation decisions | Final decisions must be recorded as human-approved, rejected, or revised |
| Knowledge-base maintenance | Are new records curated, versioned, and validated before influencing future recommendations? | Data-curation workflow review; versioned knowledge-base testing; change-impact assessment | New protocols, excipient updates, and dissolution data are indexed only after quality checks | Continuous learning from incomplete, erroneous, or unreviewed records | Knowledge updates require governance approval or defined curation responsibility |
| Uncertainty and boundary handling | Does the copilot recognize sparse, contradictory, novel, or out-of-scope evidence? | Stress testing with incomplete records, novel excipient systems, conflicting reports, and unsupported questions | System qualifies answers, asks for missing context, or refuses unsupported recommendations | Overconfident guidance in poorly represented formulation spaces | Scientist must receive explicit uncertainty statements and evidence limitations |
| Role-aware explanation | Does the system adapt explanation depth without altering the underlying evidence? | Review of outputs for junior formulators, senior scientists, QA reviewers, and project leads | Same evidence base presented at different levels of detail and regulatory formality | Oversimplification for junior users or excessive detail for expert users | Users must be able to request more evidence, less detail, or full source trace |

Decision-Making Efficiency

A second evaluation dimension would examine whether the copilot helps formulators reach better-supported decisions with less manual searching across documents, spreadsheets, and dissolution systems. Such studies should compare the experience of answering formulation questions with and without the copilot, focusing on confidence, evidence retrieval, and perceived usefulness rather than unsupported claims of productivity gains. AI applications in solid oral dosage-form development and formulation optimization reviews emphasize that practical adoption depends on how well models fit into expert workflows and improve interpretability of complex formulation information [8, 12]. Any efficiency evaluation should therefore be framed as a user-centered assessment rather than as proof that the copilot can replace formulation judgment.

Prospective Impact on Formulation Cycle Time

Prospective evaluation could examine how the copilot influences formulation planning during a pilot development project, particularly when teams are attempting to achieve a target dissolution profile or resolve recurring release failures. The appropriate question is whether the system helps scientists generate better hypotheses, retrieve more relevant precedents, and design more focused experiments, not whether it independently optimizes a formulation. Studies on dissolution-profile prediction, release design, and machine-learning-guided drug delivery systems provide conceptual support for evaluating AI as part of the design-make-test cycle [11, 18, 25]. Such evaluation should remain cautious, using qualitative and process-oriented evidence until the copilot is validated in specific organizational and product contexts.

Limitations

Data Quality and Completeness

The copilot’s usefulness would be bounded by the quality, completeness, and representativeness of the ingested formulation records. Historical protocols may omit failed experiments, excipient substitutions, informal rationales, or processing deviations that are crucial for interpreting dissolution behavior. Models using multimodal data and formulation databases show that structured information can support prediction and reasoning, but they also imply that incomplete or inconsistent records limit what AI systems can infer [17, 28]. The copilot should therefore expose uncertainty when evidence is sparse and avoid presenting incomplete historical knowledge as comprehensive guidance.

Handling of Novel Formulation Approaches

The system may struggle with novel excipient combinations, manufacturing technologies, or drug-product designs that are not well represented in its knowledge base. In such cases, similarity search and precedent retrieval may identify weak analogues but cannot establish that a new approach will perform as intended. Work on machine-learning prediction for emerging dosage forms and formulation technologies suggests that AI can assist exploration, yet unfamiliar formulation spaces still require mechanistic understanding, experimental creativity, and expert challenge [6, 14, 24]. The copilot should therefore support innovation by organizing evidence and generating hypotheses, while leaving final scientific interpretation to the formulation team.

Conclusion

A multimodal AI copilot for formulation development would unify development protocols, excipient data, and dissolution curves within a single evidence-oriented interface. By allowing scientists to ask natural-language questions across these sources, the system could assist formulation design, troubleshooting, and knowledge reuse. Its purpose would be to support expert reasoning rather than to automate formulation decisions. This makes the copilot best understood as a human-supervised layer for scientific evidence synthesis.

The central strength of the proposed copilot is integrative reasoning across modalities that are usually handled separately. Protocol text, excipient attributes, formulation composition, and dissolution behavior each provide partial insight, but their combined interpretation is often what guides useful formulation action. A conversational interface would make this combined evidence easier to access during planning and review. Traceable outputs would also help align AI assistance with the expectations of regulated development environments.

Important challenges remain before such a system could be deployed responsibly. Historical data quality, inconsistent documentation, incomplete excipient metadata, and fragmented dissolution records could all limit the reliability of generated recommendations. Integration into existing laboratory systems would require attention to access control, auditability, validation, and change management. Building user trust would depend on rigorous evaluation and clear communication of uncertainty.

Pilot deployment in pharmaceutical development laboratories would be a practical next step. Such pilots should examine how the copilot affects formulation reasoning, evidence retrieval, hypothesis generation, and project communication. The most meaningful outcome would not be autonomous formulation design but better-supported human decision-making. With careful governance, a multimodal formulation copilot could help shift development toward more transparent, data-driven, and hypothesis-guided experimentation.

Acknowledgments: None

Conflict of interest: None

Financial support: None

Ethics statement: None

References

1. Akseli I, Xie J, Schultz L, Ladyzhynsky N, Bramante T, He X, et al. A practical framework toward prediction of breaking force and disintegration of tablet formulations using machine learning tools. *J Pharm Sci.* 2017;106(1):234-47.
2. Han R, Yang Y, Li X, Ouyang D. Predicting oral disintegrating tablet formulations by neural network techniques. *Asian J Pharm Sci.* 2018;13(4):336-42.
3. Baranwal Y, Román-Ospino AD, Keyvan G, Ha JM, Hong EP, Muzzio FJ, et al. Prediction of dissolution profiles by non-destructive NIR spectroscopy in bilayer tablets. *Int J Pharm.* 2019;565:419-36.
4. Nagy B, Petra D, Galata DL, Démuth B, Borbás E, Marosi G, et al. Application of artificial neural networks for Process Analytical Technology-based dissolution testing. *Int J Pharm.* 2019;567:118464.
5. Manda A, Walker RB, Khamanga SM. An artificial neural network approach to predict the effects of formulation and process variables on prednisone release from a multipartite system. *Pharmaceutics.* 2019;11(3):109.
6. Elbadawi M, Castro BM, Gavins FK, Ong JJ, Gaisford S, Pérez G, et al. M3DISEEN: A novel machine learning approach for predicting the 3D printability of medicines. *Int J Pharm.* 2020;590:119837.
7. Bannigan P, Aldeghi M, Bao Z, Häse F, Aspuru-Guzik A, Allen C. Machine learning directed drug formulation development. *Adv Drug Deliv Rev.* 2021;175:113806.
8. Lou H, Lian B, Hageman MJ. Applications of machine learning in solid oral dosage form development. *J Pharm Sci.* 2021;110(9):3150-65.
9. Wang N, Sun H, Dong J, Ouyang D. PharmDE: A new expert system for drug-excipient compatibility evaluation. *Int J Pharm.* 2021;607:120962.
10. Galata DL, Könyves Z, Nagy B, Novák M, Mészáros LA, Szabó E, et al. Real-time release testing of dissolution based on surrogate models developed by machine learning algorithms using NIR spectra, compression force and particle size distribution as input data. *Int J Pharm.* 2021;597:120338.
11. Castro BM, Elbadawi M, Ong JJ, Pollard T, Song Z, Gaisford S, et al. Machine learning predicts 3D printing performance of over 900 drug delivery systems. *J Control Release.* 2021;337:530-45.
12. Wang S, Di J, Wang D, Dai X, Hua Y, Gao X, et al. State-of-the-art review of artificial neural networks to predict, characterize and optimize pharmaceutical formulation. *Pharmaceutics.* 2022;14(1):183.
13. Jiang J, Ma X, Ouyang D, Williams III RO. Emerging artificial intelligence (AI) technologies used in the development of solid dosage forms. *Pharmaceutics.* 2022;14(11):2257.
14. Wang F, Elbadawi M, Tsilova SL, Gaisford S, Basit AW, Parhizkar M. Machine learning predicts electrospray particle size. *Mater Des.* 2022;219:110735.

15. Galata DL, Zsiros B, Mészáros LA, Nagy B, Szabó E, Farkas A, et al. Raman mapping-based non-destructive dissolution prediction of sustained-release tablets. *J Pharm Biomed Anal.* 2022;212:114661.
16. McCoubrey LE, Seegobin N, Elbadawi M, Hu Y, Orlu M, Gaisford S, et al. Active machine learning for formulation of precision probiotics. *Int J Pharm.* 2022;616:121568.
17. Abdalla Y, Elbadawi M, Ji M, Alkahtani M, Awad A, Orlu M, et al. Machine learning using multi-modal data predicts the production of selective laser sintered 3D printed drug products. *Int J Pharm.* 2023;633:122628.
18. Mazur H, Erbrich L, Quodbach J. Investigations into the use of machine learning to predict drug dosage form design to obtain desired release profiles for 3D printed oral medicines. *Pharm Dev Technol.* 2023;28(2):219-31.
19. Sun S, Alkahtani ME, Gaisford S, Basit AW, Elbadawi M, Orlu M. Virtually possible: enhancing quality control of 3D-printed medicines with machine vision trained on photorealistic images. *Pharmaceutics.* 2023;15(11):2630.
20. Hang NT, Long NT, Duy ND, Chien NN, Van Phuong N. Towards safer and efficient formulations: Machine learning approaches to predict drug-excipient compatibility. *Int J Pharm.* 2024;653:123884.
21. Wu Z, Wang N, Ye Z, Xu H, Chan G, Ouyang D. FormulationBCS: a machine learning platform based on diverse molecular representations for Biopharmaceutical Classification System (BCS) class prediction. *Mol Pharm.* 2024;22(1):330-42.
22. Lourenco AS, Schuster T, Lopes JA, Kirsch A. A non-linear modelling approach to predict the dissolution profile of extended-release tablets. *Eur J Pharm Sci.* 2025;204:106976.
23. Sousa AS, Serra J, Estevens C, Costa R, Ribeiro AJ. A comparative study of two data-driven modeling approaches to predict drug release from ER matrix tablets. *Int J Pharm.* 2025;671:125230.
24. Au-Yeung L, Tseng CY, Tam YK, Tsai PA. Machine learning based quantitative structure–dissolution profile relationship. *J Chem Inf Model.* 2025;65(12):6273-86.
25. Protopapa C, Siamidi A, Eneli AA, Elbadawi M, Vlachou M. Machine learning predicts drug release profiles and kinetic parameters based on tablets' formulations. *AAPS J.* 2025;27(5):124.
26. Jia S, Wang N, Yang R, Zhao X, Ding Y, Ouyang D. FormulationDE: an updated artificial intelligence system for drug-excipient compatibility prediction. *AAPS Open.* 2026;12(1):20.
27. Hamaguchi M, Adachi T, Arai N. Application of AI in tablet development: An integrated machine learning framework for pre-formulation property prediction. *Pharmaceutics.* 2026;18(4):452.
28. Zhou J, Li C, Zan P, Wang Z, Wang B, Zhao Y, et al. A structured oral formulation database for machine learning: Uncovering data-informed design strategies to facilitate effective formulation development. *Int J Pharm.* 2026:126957.