Pharmacophore

ISSN-2229-5402

Journal home page: http://www.pharmacophorejournal.com

AUTONOMOUS LABORATORY PLANNING AGENT FOR SOLID-STATE FORM SELECTION AND POLYMORPH SCREENING

Rashid Al-Mahdi[1], Khalifa Al-Suwaidi[1]*, Mariam Al-Kuwari[2]

1. ***Department of Pharmaceutical Informatics and AI Systems, Faculty of Pharmacy, Qatar University, Doha, Qatar.***
2. ***Department of Computational Drug Engineering, Faculty of Engineering, Hamad Bin Khalifa University, Doha, Qatar.***

ARTICLE INFO

Received:
26 January 2026
Received in revised form:
10 April 2026
Accepted:
15 April 2026
Available online:
28 April 2026

Keywords: Autonomous laboratories, Solid-state form selection, Polymorph screening, Cocrystals, Salts, Bayesian optimization

ABSTRACT

Solid-form screening is a foundational step in pharmaceutical development because the physical form of an active pharmaceutical ingredient can influence stability, manufacturability, dissolution, and downstream formulation strategy. Polymorphs, cocrystals, and salts each expand the developable form landscape, but they also increase the complexity of experimental exploration. Current solid-form screening workflows often rely on expert-designed experimental grids, manual interpretation of characterization data, and sequential decision-making. These practices can be slow, material-intensive, and vulnerable to incomplete exploration of crystallization conditions. This article proposes an autonomous AI planning agent that designs solid-form screening protocols, coordinates robotic execution, analyzes diffraction and spectroscopic data in real time, and adaptively refines the next experimental campaign. The agent is conceptual and intended as a system architecture rather than a report of experimental performance. The proposed system combines a Bayesian or reasoning-based planning core, robotic crystallization and sample-handling modules, solid-state characterization tools, and AI-powered form identification. A learning loop updates the agent's internal representation of the crystallization space after each experimental batch. Such an agent would be expected to make solid-form screening more systematic, traceable, and adaptive. It could support broader exploration of polymorph, cocrystal, and salt landscapes while preserving expert oversight at critical scientific and safety decision points. An autonomous solid-state screening agent could transform solid-form selection from a predominantly empirical workflow into a data-driven, closed-loop planning process. Its value would depend on robust integration of planning, robotics, characterization, human review, and prospective validation.

To Cite This Article: Al-Mahdi R, Al-Suwaidi K, Al-Kuwari M Autonomous Laboratory Planning Agent for Solid-State Form Selection and Polymorph Screening. Pharmacophore. 2026;17(2):103-11. https://doi.org/10.51847/M9DXGmqVYy

Introduction

Selecting the optimal solid form of a drug candidate is a critical decision in pharmaceutical development because polymorphism, salt formation, and cocrystallization can affect stability, processability, dissolution, and formulation behavior. Computational descriptions of molecular crystal polymorphs have become increasingly practical for anticipating possible solid forms, but prediction alone does not replace experimental confirmation. Modern crystal structure prediction can map plausible polymorph landscapes and support form-risk assessment, yet the physical realization of forms still depends on crystallization conditions and characterization evidence. An autonomous planning agent is therefore best viewed as a bridge between computational foresight and experimental solid-form verification.

Conventional screening typically explores solvent, temperature, concentration, additives, and crystallization pathway choices through manually designed grids or high-throughput campaigns. Robotic crystallization and transmission powder X-ray diffraction screening have improved throughput, but the resulting workflows can still depend heavily on expert selection of follow-up experiments [1]. Reviews of AI-assisted pharmaceutical crystallization emphasize that crystallization remains a multidimensional decision problem in which kinetics, thermodynamics, and practical constraints are intertwined [2]. This creates a strong rationale for an adaptive agent that can learn from each batch rather than treating screening as a fixed experimental checklist.

Corresponding Author: Khalifa Al-Suwaidi; Department of Pharmaceutical Informatics and AI Systems, Faculty of Pharmacy, Qatar University, Doha, Qatar. E-mail: khalifa.suwaidi@outlook.com

Self-driving laboratories and autonomous discovery systems provide a conceptual foundation for this shift because they combine experiment selection, robotic execution, measurement, and model updating into a closed-loop workflow [3]. ChemOS demonstrated how orchestration software can coordinate instruments and decision engines for autonomous discovery, while later architectures extended this idea toward more modular laboratory control [4, 5]. Bayesian optimizers such as Phoenics show how chemical search spaces can be explored adaptively without requiring exhaustive enumeration of all possible conditions [6]. These developments suggest that solid-state form discovery could be reframed as an agentic planning problem rather than a purely manual laboratory campaign.

The thesis of this article is that an autonomous laboratory planning agent could orchestrate polymorph, cocrystal, and salt screening by combining structured experimental planning, robotic execution, real-time characterization, and human-governed decision gates. Such a system would not replace expert solid-state scientists but would provide a disciplined mechanism for proposing experiments, preserving digital traceability, and flagging uncertain or potentially novel forms for review [7]. Recent examples of autonomous workflows in solid-state chemistry and AI-driven robotic crystal exploration show that this direction is technically plausible, although prospective evaluation in pharmaceutical contexts remains necessary [8, 9]. The proposed EAI agent is therefore presented as a conceptual architecture for future development rather than as a claim of achieved experimental performance.

Background

Polymorphs, Cocrystals, and Salts: The Solid-Form Landscape

The solid-form landscape of a drug substance includes polymorphs of the same molecule, multicomponent cocrystals, salts, solvates, hydrates, and related crystalline or partially ordered forms. Computational polymorph studies emphasize that multiple low-energy packing arrangements may be plausible, making form discovery a search over both molecular packing and experimental crystallization pathways. Cocrystal prediction studies further show that coformer selection can be guided by molecular descriptors, graph representations, and interaction patterns, but these tools still require experimental validation before development decisions are made [10, 11]. A planning agent must therefore treat form selection as a coupled prediction, experimentation, and confirmation process rather than a single model-output task.

Traditional and High-Throughput Approaches to Solid-Form Screening

Traditional solid-form screening uses solvent panels, slurry experiments, cooling crystallization, evaporation, antisolvent addition, and thermal or mechanical stress conditions to explore likely form outcomes. High-throughput crystallization platforms can expand the number and diversity of experiments, while automated powder diffraction analysis can accelerate triage of crystalline products [1]. Machine-learning clustering of powder diffraction and Raman outputs has been proposed as a way to organize large screening datasets and reduce the burden of manual pattern comparison [12]. However, high throughput alone does not ensure adaptive exploration unless the next experiments are informed by the accumulated results.

Autonomous Agents and Self-Driving Laboratories in Chemistry

Autonomous chemical laboratories are typically organized around a loop in which an AI system proposes experiments, laboratory hardware executes them, instruments measure outcomes, and the decision engine updates future plans [3]. Self-driving laboratory reviews in chemistry and materials science describe this loop as a general framework that can be adapted to different domains, including molecular discovery, materials optimization, and formulation-relevant workflows [13]. Universal self-driving laboratory concepts emphasize that the same architectural logic can coordinate synthesis, measurement, and learning across varied experimental targets [14]. For solid-state pharmaceutical development, the corresponding agent would need to translate abstract screening objectives into executable crystallization and characterization tasks.

Machine Learning for Polymorph Prediction and Characterization

Machine learning can support polymorph screening both before and after the experiment. Before experimentation, data-efficient models and active-learning potentials can help prioritize molecular crystal structures and reduce the cost of exploring crystal-packing landscapes. After experimentation, automated X-ray diffraction classification and adaptive diffraction analysis can assist with phase identification, quality control, and novelty detection in large datasets [15, 16]. These capabilities are complementary: prediction proposes what may exist, while characterization models help determine what has actually been produced.

Human-in-the-Loop Laboratory Automation

Human oversight remains central in autonomous laboratories because experimental recommendations, safety constraints, and scientific claims require accountable review. Human-in-the-loop perspectives emphasize that interactive AI, scientist feedback, and transparent interfaces are necessary for trustworthy laboratory autonomy [7]. In solid-form screening, this oversight is especially important when the agent flags a potentially new form, recommends additional characterization, or proposes final form selection for development. The appropriate design is therefore not full replacement of the scientist but structured collaboration between the agent and expert decision-makers.

Agent Architecture Overview

High-Level Architecture

The proposed agent consists of five interacting layers: a planning core, a robotic execution interface, a characterization data pipeline, a solid-form knowledge base, and a human-review console. The planning core draws on Bayesian optimization, rule-based constraints, and domain knowledge to propose crystallization conditions, reflecting principles already used in autonomous chemical experimentation [3, 6]. The robotic interface converts plans into executable liquid-handling, temperature-control, and sample-transfer operations, while the characterization pipeline analyzes XRPD, DSC, Raman, and related data streams. The knowledge base stores links among materials, conditions, observations, and form assignments so that each campaign becomes part of a cumulative decision record.

Figure 1 presents the proposed autonomous laboratory planning architecture linking development constraints, adaptive experimental design, robotic crystallization, real-time solid-state characterization, knowledge-graph updating, and human-governed solid-form selection.

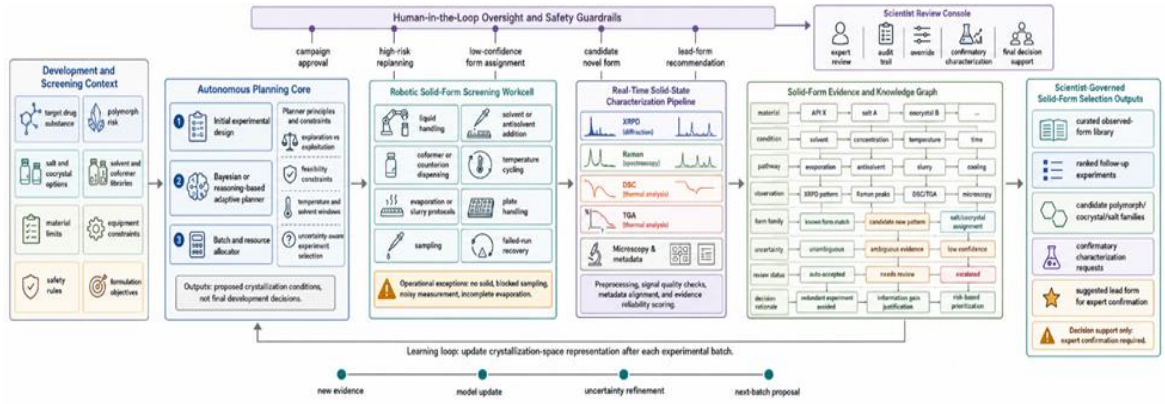

Figure 1. Autonomous laboratory planning agent for solid-state form selection and polymorph screening.

Core Inputs and Outputs

The core inputs include the target drug substance, available counterions and coformers, solvent libraries, material limits, temperature ranges, concentration windows, and constraints imposed by equipment and safety protocols. Cocrystal-prediction tools based on molecular features or graph neural networks could help rank coformers before the first robotic campaign is launched [17, 18]. The outputs would include a curated library of observed solid forms, associated characterization data, inferred relationships among forms, and a suggested lead form requiring human confirmation. Because cocrystal screening tools are increasingly available through deep-learning-assisted web systems, the agent could also integrate external prediction services into its candidate-generation workflow when validated for the intended use context [19].

Design Principles

The first design principle is adaptive breadth: the agent should explore broadly enough to avoid premature convergence while still using incoming evidence to prioritize promising regions of the experimental space. The second is traceability, because self-driving laboratory architectures require a clear record of what was proposed, executed, measured, and learned at each step [5, 13]. The third is operational safety, implemented through hard constraints on solvent compatibility, temperature ranges, material handling, and instrument access. The fourth is scientific humility: the agent may flag a pattern as potentially novel, but it should not autonomously declare a developable form without expert review and confirmatory characterization.

Table 1 defines the functional logic of the proposed agent by separating planning, robotic execution, characterization, evidence memory, and human-governed decision support into distinct architectural layers.

Table 1. Architectural logic of the autonomous solid-form screening agent

Agent layer	Primary function	Representative inputs	AI or automation logic	Solid-state decision contribution	Required human oversight
Development-context definition	Translates pharmaceutical objectives into a constrained screening problem	Drug substance, formulation goals, polymorph risk, salt/cocrystal options, material availability, safety limits	Structured objective setting; rule-based feasibility constraints; prior-knowledge encoding	Defines what the agent is allowed to explore and what counts as a scientifically useful outcome	Scientist approves campaign scope, excluded conditions, and development priorities
Candidate and condition generation	Seeds the initial solid-form search space	Solvent libraries, coformers, counterions, crystallization	Diversity-aware design; cocrystal or salt candidate ranking;	Prevents overdependence on narrow manual	Expert reviews whether proposed conditions are

		pathways, concentration and temperature windows	knowledge-based condition selection	grids and broadens early exploration	chemically plausible and safe
Adaptive planning core	Selects the next batch of experiments after each evidence cycle	Previous outcomes, failed runs, uncertain regions, form-family assignments, unexplored conditions	Bayesian optimization, active learning, reasoning-based planning, exploration–exploitation balancing	Converts screening from a static checklist into a cumulative learning process	Scientist reviews high-risk replanning, unusual recommendations, and major changes in campaign direction
Robotic execution interface	Converts planned experiments into executable laboratory actions	Batch plans, well-plate maps, dispensing instructions, temperature protocols, sampling schedules	Instrument orchestration; resource-aware scheduling; exception handling	Makes experiments repeatable, machine-readable, and traceable across cycles	Human intervention required for hardware faults, safety exceptions, or unapproved deviations
Real-time characterization pipeline	Converts experimental products into structured evidence	XRPD patterns, Raman spectra, DSC/TGA traces, microscopy images, metadata, quality flags	Signal preprocessing; pattern matching; clustering; novelty detection; confidence scoring	Distinguishes known forms, ambiguous products, failed experiments, and candidate new forms	Expert confirms low-confidence assignments, conflicting evidence, and claims of novelty
Solid-form knowledge graph	Preserves cumulative relationships among experiments, evidence, and decisions	Material identity, conditions, pathways, observations, form assignments, uncertainty, review status	Relational evidence modeling; uncertainty tracking; redundancy detection; rationale storage	Enables transparent learning, avoids redundant experiments, and supports auditability	Scientist validates interpretation, decision rationale, and final form-family classification
Human-review console	Provides accountable oversight of autonomous recommendations	Plans, rationale, data-quality indicators, uncertainty flags, candidate outputs, audit trail	Explainable interface; decision-gate routing; override capture	Keeps autonomy bounded by scientific judgment and pharmaceutical governance	Human approval required for campaign approval, candidate novel forms, and lead-form selection
Final decision-support output	Organizes the screening outcome into development-relevant recommendations	Curated form library, ranked follow-up experiments, confirmatory tests, candidate lead form	Evidence synthesis; uncertainty-aware ranking; traceable recommendation generation	Supports form selection without allowing the agent to make autonomous regulatory or development determinations	Final form advancement remains a scientist-governed decision

Experimental Design and Planning Module

Initial Experimental Design

The initial design module would generate a diverse set of screening conditions using prior knowledge of the molecule, likely intermolecular interactions, solvent behavior, and practical crystallization routes. For cocrystal screens, machine-learning models can rank potential coformers by expected formation propensity, while complementary graph-based approaches can encode structural relationships among molecules [10, 11]. Property-driven cocrystal discovery using Gaussian processes and active learning suggests a path toward selecting candidates not only for form formation but also for downstream properties relevant to development [20]. In this conceptual agent, such models would be used to seed an initial experimental space rather than to make final form decisions.

Bayesian Optimization for Adaptive Screening

After the first batch of experiments, the planning module would shift from static design to adaptive screening. Bayesian optimization is well suited to this role because it can balance exploration of uncertain regions with exploitation of conditions that appear likely to yield informative or desirable outcomes [6]. Adaptive Bayesian approaches to crystallization can guide condition selection under different kinetic regimes, which is directly relevant to screening workflows where crystallization

behavior varies across solvents and temperatures [21]. In the proposed system, the optimizer would be evaluated conceptually by whether it proposes scientifically useful follow-up experiments, not by unsupported claims of numerical performance.

Batch Planning and Resource Allocation

A solid-state screening agent must plan in batches because robotic platforms, well plates, temperature zones, filtration steps, and characterization queues impose physical constraints. Modular multi-robot laboratory integration demonstrates how autonomous workflows can coordinate multiple instruments and operations in solid-state chemistry, making resource-aware scheduling a central architectural concern [8]. The agent would need to convert high-level experimental intent into feasible batches while respecting incompatible solvents, limited heating or cooling positions, sampling order, and instrument availability. This batch planner would also maintain contingency logic so that failed crystallizations or ambiguous measurements can be revisited in later cycles.

Robotic Integration and Closed-Loop Execution

Robotic Workflow for Crystallization and Form Screening

The robotic workflow begins when the planning core dispatches executable instructions for weighing or dispensing material, adding solvent or antisolvent, introducing coformers or salt formers, and applying temperature or evaporation protocols. AI-driven robotic crystal exploration provides a model for integrating automated crystallization with rapid form triage, although pharmaceutical deployment would require strict validation and human governance [9]. Liquid handling, plate-based crystallization, robotic transfer, and automated sampling would together allow the agent to execute repeatable campaigns with a consistent digital record. The purpose of automation in this context is not merely speed but the ability to make each experiment machine-readable for subsequent learning.

Real-Time Data Pipeline

After each experimental batch, the characterization pipeline would collect XRPD patterns, thermal traces, Raman spectra, and supporting metadata, then pass them through automated preprocessing and quality checks. Machine-learning methods for large synchrotron and laboratory X-ray diffraction datasets show how automated analysis can support phase recognition and data triage when the number of patterns becomes difficult for manual review [16, 22]. Adaptive X-ray diffraction guided by machine learning also illustrates how measurement and interpretation can be coupled to subsequent decisions rather than treated as separate stages [15]. In the proposed agent, only data that pass predefined quality criteria would be used to update the planning model without human intervention.

Error Handling and Recovery

Closed-loop execution requires the agent to distinguish expected experimental outcomes from operational failures. For example, no solid formation, blocked sampling, noisy diffraction, incomplete evaporation, or instrument communication errors should trigger predefined recovery policies rather than unstructured replanning. Autonomous orchestration frameworks such as ChemOS emphasize the importance of coordinating instruments and software components so that experiments can proceed through controlled feedback loops [4]. In solid-form screening, the agent should be allowed to repeat or modify routine failed runs within approved constraints, while ambiguous safety, identity, or novelty issues should be escalated to the human scientist.

Real-Time Characterization and Form Identification

AI-Powered Pattern Matching and Novelty Detection

Each acquired XRPD pattern would be compared against the agent's internal library of known forms and against reference patterns generated from computational or experimental sources. Automated diffraction classification using deep learning can support phase recognition, while adaptive diffraction workflows show how measurement choices can be informed by machine-learning interpretation [15, 16]. For solid-form screening, the agent should treat a low-confidence match as a prompt for additional evidence rather than as proof of a new polymorph. Clustering methods that combine XRPD and Raman data provide a useful model for organizing high-throughput screening outputs into candidate form families for expert review [12].

Automated Thermophysical Analysis

Thermal analysis would complement diffraction by helping the agent distinguish polymorphs, solvates, hydrates, salts, and cocrystals that may show similar or partially overlapping XRPD features. Reviews of AI and machine learning in crystallization emphasize that characterization data must be interpreted alongside process conditions because crystallization history can influence the observed solid form [2, 23]. In the proposed architecture, DSC and TGA outputs would be preprocessed into structured events such as melting, desolvation, or recrystallization, then linked to the corresponding diffraction and spectroscopic record. The agent would be expected to flag conflicting evidence for human review rather than forcing a single automated form assignment.

Updating the Solid-Form Knowledge Graph

The agent's knowledge graph would connect starting material, solvent system, coformer or counterion, crystallization pathway, temperature history, observed form, characterization evidence, and review status. Cocrystal-screening studies using machine

learning and combined prediction workflows show that relationships among molecular descriptors, experimental outcomes, and form properties can be represented in ways that support future prioritization [19, 24]. As each experiment is completed, the graph would store not only the outcome but also the uncertainty and decision rationale associated with that outcome. This memory would allow the planner to avoid redundant experiments, revisit uncertain regions, and preserve a transparent record for later scientific interpretation.

Human-in-the-Loop Oversight and Decision Gates

Review Dashboard and Decision Points

The human-review dashboard would display the current experimental plan, the agent's rationale, recent characterization outputs, data-quality flags, and any proposed deviations from the approved campaign. Human-in-the-loop laboratory automation frameworks emphasize that autonomy should be interactive, allowing scientists to inspect recommendations and intervene when scientific judgment or safety considerations require it [7]. A minimal working example for a self-driving laboratory also highlights the need to define the smallest reliable loop connecting planning, execution, measurement, and decision-making before expanding autonomy [25]. In this agent, campaign approval, high-risk replanning, and any claim of novelty would remain explicit decision points for the scientist.

Confirmation of Novel Forms and Final Selection

The agent may recommend that a diffraction pattern, thermal signature, or Raman profile represents a potentially new form, but confirmation should require expert assessment and additional characterization. Reliable molecular crystal prediction can help determine whether a candidate form is physically plausible, while computational crystal structure landscapes can guide interpretation of ambiguous experimental patterns. For cocrystals and salts, the scientist would also evaluate whether the form is developable, stable, reproducible, and aligned with formulation needs rather than merely distinct. The final lead-form recommendation should therefore be framed as decision support, not as an autonomous regulatory or development determination.

Table 2 provides an evidence-to-decision framework showing how experimental observations, characterization signals, computational priors, uncertainty, and human review jointly support solid-form selection.

Table 2. Evidence-to-decision framework for polymorph, cocrystal, and salt screening

Evidence domain	What the agent observes	Interpretation challenge	Agent-level reasoning task	Decision gate	Development relevance
Crystallization outcome	Solid formation, no solid, oiling out, amorphous material, mixed solids, failed sampling	Experimental failure may look like negative chemical evidence if not properly classified	Separate operational failures from meaningful form-screening outcomes	Routine rerun may be automated; repeated or unexplained failure escalates to scientist	Prevents false conclusions about inaccessible forms or unsuitable conditions
XRPD evidence	Peak position, peak intensity, crystallinity, pattern similarity, mixed-phase indicators	Similar forms may have overlapping or low-quality diffraction patterns	Match against known forms, cluster related patterns, flag low-confidence or potentially novel patterns	Low-confidence, mixed, or novel patterns require expert review and confirmatory testing	Core evidence for identifying polymorphs and differentiating crystalline forms
Spectroscopic evidence	Raman or related spectral fingerprints linked to molecular environment	Spectral changes may reflect salt/cocrystal formation, solvation, impurity, or physical mixture	Compare spectral signatures with diffraction and condition metadata	Conflicting spectroscopy and XRPD evidence triggers review	Supports distinction among polymorphs, cocrystals, salts, solvates, and physical mixtures
Thermal evidence	Melting, desolvation, recrystallization, dehydration, glass transition, decomposition events	Thermal events may overlap or reflect sample history rather than a distinct form	Convert thermal traces into structured events and link them to candidate form families	Unexpected or contradictory thermal events require human interpretation	Helps assess stability, solvation, hydrate status, and developability
Process-condition evidence	Solvent, coformer, counterion, concentration, temperature, evaporation, slurry, cooling, aging	A form assignment is incomplete without knowing how it was generated	Link each observation to the exact pathway that produced it	Scientist reviews whether the pathway is reproducible and development-relevant	Supports reproducibility, scale-up reasoning, and quality-by-design documentation
Computational prior evidence	Predicted polymorph landscape, coformer	Computational plausibility does	Use predictions to prioritize	Computationally suggested forms	Bridges crystal-structure

	ranking, salt/cocrystal feasibility, plausible packing motifs	not equal experimental confirmation	experiments and interpret ambiguous observations	require experimental evidence before advancement	prediction and empirical screening
Uncertainty and novelty evidence	Low match confidence, rare pattern clusters, conflicting modalities, sparse condition coverage	Novelty may reflect noise, impurity, mixture, or genuine new form	Assign uncertainty, recommend confirmatory characterization, and avoid premature claims	Any candidate new form requires expert confirmation	Protects scientific validity and avoids overclaiming autonomous discovery
Human decision evidence	Scientist approval, override, comment, rejection, confirmation, requested follow-up	Human decisions are often lost outside the experimental data stream	Store review status and rationale as part of the knowledge graph	Required for campaign changes, novelty claims, and lead-form recommendation	Creates an auditable record suitable for pharmaceutical development governance
Final selection evidence	Integrated form identity, stability, reproducibility, manufacturability, dissolution relevance, formulation fit	The most novel form may not be the most developable form	Synthesize evidence into ranked decision-support outputs	Final selection remains a human-governed development decision	Aligns autonomous screening with practical form-selection strategy

Integration into Pharmaceutical Development

Early-Stage Solid-Form Selection

During early development, the agent could support rapid exploration of the solid-form landscape before formulation decisions become difficult to reverse. Self-driving laboratory architectures are particularly relevant here because they can coordinate experiment proposal, robotic execution, measurement, and learning within a single digital workflow [3, 13]. Cocrystal and salt screening modules could expand the developable options when the neutral crystalline form has limitations, using prediction tools to prioritize candidates for experimental confirmation [17, 20]. The expected benefit is a more systematic understanding of form risk and opportunity, not an unsupported promise of a specific discovery rate.

Supporting Quality-by-Design and Regulatory Submissions

The agent's digital record could support quality-by-design thinking by documenting the search space, constraints, experimental sequence, characterization evidence, and human decisions that shaped final form selection. Orchestration systems such as ChemOS and ChemOS 2.0 illustrate how autonomous laboratories can preserve links between plans, instruments, measurements, and model updates [4, 5]. In a pharmaceutical setting, this traceability would help explain why certain solid forms were advanced, rejected, or reserved for further study. The agent should therefore be designed as an auditable decision-support system whose outputs can be reviewed, challenged, and reproduced within validated laboratory procedures.

Evaluation Strategy

Screening Efficiency and Novelty Yield

Evaluation should compare the agent's screening process with established manual or high-throughput workflows using retrospective compounds, simulated campaigns, or carefully governed prospective studies. Robotic crystallization and high-throughput PXRD workflows provide relevant baselines because they already generate structured data streams suitable for comparison [1]. The assessment should focus on whether the agent could recover known forms, prioritize informative conditions, and flag candidate novel forms for review without relying on unsupported performance claims. Any evaluation of novelty yield should distinguish between automated pattern novelty, expert-confirmed new forms, and developable forms.

Predictive Accuracy of the Planner

The planner should be evaluated by comparing its predicted experimental outcomes with observed crystallization behavior, characterization assignments, and subsequent human decisions. Bayesian optimization methods and adaptive crystallization studies provide a conceptual basis for assessing whether proposed conditions become more informative as evidence accumulates [6, 21]. For computational guidance, machine-learned crystal-structure-prediction potentials and data-efficient crystal prediction methods could be evaluated by how well they prioritize plausible forms for experimental pursuit. The goal is not to claim universal predictive accuracy but to determine whether the planner improves the scientific usefulness of each screening round.

Usability and Trust

Usability evaluation should examine whether scientists can understand the agent's plans, override recommendations, trace evidence, and confidently interpret flagged forms. Human-centered self-driving laboratory perspectives emphasize trust, interaction quality, and scientist control as central requirements for laboratory autonomy [7]. Reviews of autonomous

laboratories also indicate that reliable operation depends on transparent integration among software, instruments, and decision engines, not only on model quality [3, 13]. The agent should therefore be assessed as a socio-technical system in which trust depends on explainability, reproducibility, safety behavior, and the clarity of human decision gates.

Limitations

Hardware Dependence and Vendor Lock-In

The agent's practical behavior would depend strongly on the robotic platform, liquid handler, crystallization hardware, sample-transfer mechanism, and characterization instruments available in a given laboratory. Modular multi-robot integration shows how powerful autonomous workflows can become, but it also implies that portability requires careful abstraction of hardware commands, error states, and scheduling constraints [8]. A workflow designed for one laboratory may not transfer directly to another without revalidation of sample handling, measurement quality, and software-instrument communication. Standardized interfaces and shared reference implementations would therefore be important for avoiding vendor lock-in and improving reproducibility.

Limited to Observable Solid Forms

The proposed agent can only learn from forms that are accessible under the conditions, timescales, and characterization modalities included in its workflow. Crystal structure prediction and reliable computational ranking can help reveal plausible but unobserved polymorphs, while frontiers in molecular crystal prediction show that computation remains an essential complement to experimental screening. Forms requiring unusual pressures, long aging times, trace impurities, specialized surfaces, or manufacturing-scale conditions may remain outside the agent's practical search horizon. The agent should therefore report the boundaries of its exploration rather than implying that its campaign exhausts the entire solid-form landscape.

Conclusion

An autonomous AI planning agent for solid-state form selection would combine experimental design, robotic crystallization, real-time characterization, adaptive learning, and human review into a single coordinated workflow. Its purpose would be to make polymorph, cocrystal, and salt screening more systematic, traceable, and responsive to incoming evidence.

The major strength of this approach is that it could transform screening from a sequence of disconnected experiments into a cumulative learning process. By preserving a complete digital audit trail, the agent would help scientists understand not only which forms were observed but also why particular experiments were selected and how decisions evolved.

Important challenges remain before such an agent could become routine in pharmaceutical development. Hardware integration, safety guardrails, uncertainty communication, data standardization, and prospective validation across diverse drug candidates would all need careful attention.

Open reference implementations and collaborative benchmark studies would help establish shared expectations for autonomous solid-form screening. With appropriate validation and expert oversight, this class of agent could become a practical foundation for data-driven solid-state pharmaceutical development.

Acknowledgments: None

Conflict of interest: None

Financial support: None

Ethics statement: None

References

1. Rosso VW, Yin Z, Abourahma H, Furman A, Sharif S, Werneth A, et al. High-throughput crystallization screening technique with transmission PXRD analysis. Org Process Res Dev. 2023;27(8):1437-44.
2. Zhu Z, Zhang Y, Wang Z, Tang W, Wang J, Gong J. Artificial intelligence assisted pharmaceutical crystallization. Cryst Growth Des. 2024;24(10):4245-70.
3. Seifrid M, Pollice R, Aguilar-Granda A, Morgan Chan Z, Hotta K, Ser CT, et al. Autonomous chemical experiments: challenges and perspectives on establishing a self-driving lab. Acc Chem Res. 2022;55(17):2454-66.
4. Roch LM, Häse F, Kreisbeck C, Tamayo-Mendoza T, Yunker LPE, Hein JE, et al. ChemOS: An orchestration software to democratize autonomous discovery. PLoS One. 2020;15(4):e0229862.
5. Sim M, Vakili MG, Strieth-Kalthoff F, Hao H, Hickman RJ, Miret S, et al. ChemOS 2.0: an orchestration architecture for chemical self-driving laboratories. Matter. 2024;7(9):2959-77.
6. Hase F, Roch LM, Kreisbeck C, Aspuru-Guzik A. Phoenics: a Bayesian optimizer for chemistry. ACS Cent Sci. 2018;4(9):1134-45.

7. Hysmith H, Foadian E, Padhy SP, Kalinin SV, Moore RG, Ovchinnikova OS, et al. The future of self-driving laboratories: from human in the loop interactive AI to gamification. Digital Discovery. 2024;3(4):621-36.
8. Lunt AM, Fakhruldeen H, Pizzuto G, Longley L, White A, Rankin N, et al. Modular, multi-robot integration of laboratories: an autonomous workflow for solid-state chemistry. Chem Sci. 2024;15(7):2456-63.
9. Lee EC, Salley D, Sharma A, Cronin L. AI-driven robotic crystal explorer for rapid polymorph identification. Digital Discovery. 2026.
10. Jiang Y, Yang Z, Guo J, Li H, Liu Y, Guo Y, et al. Coupling complementary strategy to flexible graph neural network for quick discovery of coformer in diverse co-crystal materials. Nat Commun. 2021;12(1):5950.
11. Wang D, Yang Z, Zhu B, Mei X, Luo X. Machine-learning-guided cocrystal prediction based on large data base. Cryst Growth Des. 2020;20(10):6610-21.
12. Yin Z, Rosso VW, Roberts F, Furman A, Stevens JM, Watts T, et al. EZClust: a robust machine learning-based powder X-ray diffraction and Raman cluster analysis model for efficient high-throughput crystallization polymorph screening. Org Process Res Dev. 2026.
13. Tom G, Schmid SP, Baird SG, Cao Y, Darvish K, Hao H, et al. Self-driving laboratories for chemistry and materials science. Chem Rev. 2024;124(16):9633-732.
14. Epps RW, Volk AA, Ibrahim MY, Abolhasani M. Universal self-driving laboratory for accelerated discovery of materials and molecules. Chem. 2021;7(10):2541-5.
15. Szymanski NJ, Bartel CJ, Zeng Y, Diallo M, Kim H, Ceder G. Adaptively driven X-ray diffraction guided by machine learning for autonomous phase identification. npj Comput Mater. 2023;9(1):31.
16. Salgado JE, Lerman S, Du Z, Xu C, Abdolrahim N. Automated classification of big X-ray diffraction data using deep learning models. npj Comput Mater. 2023;9(1):214.
17. Mswahili ME, Lee MJ, Martin GL, Kim J, Kim P, Choi GJ, et al. Cocrystal prediction using machine learning models and descriptors. Appl Sci. 2021;11(3):1323.
18. Guo J, Sun M, Zhao X, Shi C, Su H, Guo Y, et al. General graph neural network-based model to accurately predict cocrystal density and insight from data quality and feature representation. J Chem Inf Model. 2023;63(4):1143-56.
19. Guo J, Yang S, Wang C, Liu J, Guo Y, Yang Z, et al. Cocrystal Prediction Tool (CCPT): a web server for deep learning-assisted cocrystal screening and density evaluation. Cryst Growth Des. 2024;24(20):8407-14.
20. Appiah SA, McDonald MA. Property-driven cocrystal discovery via Gaussian processes and active learning. Cryst Growth Des. 2026;26(7):2688-702.
21. Pickles T, Mustoe C, Brown CJ, Florence AJ. Comparative study on adaptive Bayesian optimization for batch cooling crystallization for slow and fast kinetic regimes. Cryst Growth Des. 2024;24(3):1245-53.
22. Zhao X, Luo Y, Liu J, Liu W, Rosso KM, Guo X, et al. Machine learning automated analysis of enormous synchrotron X-ray diffraction datasets. J Phys Chem C. 2023;127(30):14830-8.
23. Xiouras C, Cameli F, Quillo GL, Kavousanakis ME, Vlachos DG, Stefanidis GD. Applications of artificial intelligence and machine learning algorithms to crystallization. Chem Rev. 2022;122(15):13006-42.
24. Birolo R, Bravetti F, Alladio E, Priola E, Bianchini G, Novelli R, et al. Speeding up the cocrystallization process: machine learning-combined methods for the prediction of multicomponent systems. Cryst Growth Des. 2023;23(11):7898-911.
25. Baird SG, Sparks TD. What is a minimal working example for a self-driving laboratory? Matter. 2022;5(12):4170-8.