

AI in Lyophilised Product Development: Opportunities and Limitations

Artificial intelligence (AI) is beginning to influence many aspects of pharmaceutical development, including areas once thought too specialist or product-specific for generic digital tools. In lyophilisation (also known as freeze-drying), AI can already help teams frame protocols, suggest first-pass cycle parameters and organise early technical thinking more quickly. That makes it useful, particularly when a project is still taking shape and development teams are deciding what to test first.



Yet freeze-drying remains a process governed by formulation behaviour, critical temperature data, equipment characteristics and product quality requirements that cannot be inferred reliably from pattern recognition alone. For scientists working on biologics, vaccines, peptides and other sensitive products, the key question is not whether AI has a role, but what are its limitations. Used well, it can support the early stages of development. Used uncritically, it can create false confidence in outputs that still require experimental confirmation and expert interpretation.

Why AI is Entering Lyophilisation Process Development

Artificial intelligence has entered pharmaceutical workflows because it offers the speed research and development teams are keen to harness. AI can save valuable time by summarising information quickly, organising technical ideas into a useful framework, and generating draft protocols and project reports. As a niche part of this drug development workflow, freeze-drying is no exception. Lyophilisation projects involve recurring questions around cycle design, formulation options, analytical methods, equipment choice and scale-up strategy, so it is unsurprising that AI is now being tested as a practical support tool in this area. When asking current AI Large Language models, a scientist can receive a draft cycle, a list of analytical considerations or a summary of typical formulation factors within seconds.

However, these useful responses should not be confused with scientific understanding.

While AI can help structure analysis and identify patterns in the information available, it does not yet understand how a specific formulation behaves under freeze concentration or where a true process limit lies for lyophilisation.

“What matters in freeze-drying is not whether a lyophilisation cycle looks plausible, but whether it stands up to product-specific data, experimental testing and the realities of scalable manufacture,” confirms Dr Kevin Ward.

Where AI Can Genuinely Help

For that reason, the real value of AI is not to completely replace the role of scientists, but to act as a tool to increase efficiency in the R&D workflow. The strongest use case for AI in freeze-drying is at the front end of development, where it can help a team create a sensible starting framework before experimental work begins. That might include drafting protocols, creating Gantt charts for

project timelines, generating a list of common excipients for formulation development, or suggesting a rough first-pass cycle concept for discussion.

It can also help non-specialists understand the terminology and ask more informed questions before speaking with scientists, engineers or CDMO partners. That can be especially useful in cross-functional teams, where discussions often involve people from formulation, analytical, manufacturing and regulatory backgrounds. Even so, its role needs to stay in proportion. AI works best as a tool for early thinking, communication and preparation. It is far less reliable when its outputs are treated as evidence-based conclusions.

Why Freeze-Drying Remains Product-Specific

In early development, AI is most useful as a practical starting point: it can help teams move faster, but it cannot replace scientific



judgment. That matters in lyophilisation, where plausible outputs can be accepted too quickly if they are not rigorously tested, and where the knowledge base is niche and not readily available online. Freeze-drying remains highly product-specific, with the physicochemical behaviour of each formulation shaping what kind of cycle is safe, robust and commercially viable. Monoclonal antibodies, peptides and vaccines may all require lyophilisation for stability, but they cannot be developed using identical process. Even with a purpose-built formulation, confident cycle design still depends on understanding the thermal properties of the mix and its associated critical temperatures such as collapse temperature, eutectic melting and glass transition. Without that data, the process window is inferred rather than defined.

That is why analytical techniques such as Freeze-Drying Microscopy (FDM) – Lyostat analysis – and Modulated Differential Scanning Calorimetry (MDSC) remain central to development: they provide product-specific data rather than generalised assumptions. They help establish suitable product temperatures during the freeze-drying process; whether additional annealing steps are required for largely crystalline systems; and what margin may exist between a robust process and a risky one. Equally important is the empirical evaluation that follows lyophilisation. Cake appearance, residual moisture and reconstitution behaviour are core indicators of whether the chosen process has produced a usable and stable product. AI models can describe these considerations but cannot replace the experimental work required to define them for a novel formulation. This limitation goes to the heart of why lyophilization remains a hands-on scientific discipline rather than a text-generation exercise.

Formulation, Scale-Up and Regulatory Limits

The same principle applies to formulation development. Though AI can be a useful co-pilot in the earliest stages of cycle thinking, formulation work is where that role narrows considerably. A wide range of excipients are used to modify certain aspects of a liquid formulation to make it 'lyo-ready', with a goal of improving functionality, stability, mechanical strength and solubility in the dry state. Formulation design requires an understanding of molecular behaviour, excipient interactions, physical stability, chemical stability and how the product must perform after drying and reconstitution. How these excipients interact with the liquid



formulation is extremely product specific, and current AI systems lack the intricate knowledge to design these formulations from scratch. A proposed formulation may sound plausible in theory but prove unsuitable once tested because it protects one quality attribute while compromising another; supports acceptable cake structure while introducing problems in reconstitution; or appears stable under only specific conditions.

These problems cannot be resolved reliably through generic pattern matching but depend on data, interpretation and trade-offs. Formulation development therefore remains rooted in laboratory work and expert review, with AI-generated formulation suggestions treated with caution. Unlike a broad process outline, a formulation recommendation can imply a level of precision and suitability that simply has not been earned.

Another area where AI reaches its limits is equipment and scale-up. Freeze-drying happens on real systems with their own behaviour, constraints and performance characteristics. Not all freeze-dryers operate in the same way, and those differences become increasingly important when a process moves from laboratory feasibility to pilot, clinical or GMP manufacture. Shelf heat transfer, condenser capacity and chamber pressure all influence how a cycle performs in practice. A generic cycle may look technically possible but still fail to deliver the desired balance of product quality, reproducibility and efficiency when transferred to a specific machine or different laboratory environment.

This is one of the clearest examples of the gap between plausibility and robustness. An AI-generated cycle may include temperatures

and pressures that appear sensible at first glance and even avoid obvious technical errors. Yet if it cannot account for the behaviour of a given system, it cannot reliably predict how long drying will take, how evenly the load will behave, or whether product temperature will track safely through the critical parts of the cycle. All of these variables will change at different scales, with larger batch sizes having different thermal properties (greater thermal lag and longer drying times), meaning a cycle developed for 1000 vials may not be suitable for a batch of 10000 vials. Predictions can be made by machine learning models to determine additional drying times/temperature changes, but these conditions will always need to be validated in real-world environments.

Differences in equipment type and batch size affect throughput, cost, transferability and confidence in the process. A development team needs a cycle to perform consistently and do so within an acceptable commercial framework. That level of judgment still depends on experience, data and equipment-specific understanding.

Regulatory and quality considerations reinforce the same conclusion. One of AI's most convincing traits is that it can produce language that sounds polished, structured and professional. In technical and regulated environments, that can be helpful for drafting and organisation but can also be misleading. AI models do not always distinguish well between robust scientific information and weaker source material, especially as peer-reviewed literature in the field of lyophilisation is not as abundant as other sections of the pharmaceutical industry. The result can be inconsistency, overconfidence

or recommendations that require substantial correction. In pharmaceutical development, regulators and quality teams are interested in whether the underlying process understanding is real, traceable and supported by data. Documentation is only as strong as the development logic behind it.

This is particularly relevant in areas such as CMC (Chemistry, Manufacturing, and Controls) strategy, analytical validation planning and technology transfer. AI may assist with drafting internal text or helping teams structure a submission narrative, but it does not remove the need for scientific ownership. Development decisions must still be justified through evidence and process choices understood in context. Transfer activities must still reflect what the product and equipment have actually shown in practice. In that sense, AI can speed up the packaging of information, but it cannot replace the responsibility carried by scientists, engineers and quality professionals. The more regulated the environment, the more important that distinction becomes.

A Practical Comparison: AI Machine Learning Model Versus Expert Design

A useful way to assess these boundaries is through direct comparison. In a case study ChatGPT 4 was prompted as follows: "Simulate a freeze-drying cycle trace for a 10% Dextran solution in a DIN10R Vial with a 2ml fill and present as a graph". The output was then directly compared to a recipe designed by experienced lyophilisation scientists for the same concentration, fill volume and container size.¹ The comparison is valuable because it moves the discussion away from abstract debate and into a real development scenario. Instead of asking whether AI can provide recommendations about the freeze-drying process, it asks whether AI can fully simulate a cycle that is genuinely efficient and robust in practice.




Figure 1: Graph showing cycle designed at Biopharma Group

Based on that comparison, the answer is nuanced. The AI-generated output was not worthless: on the contrary, it provided a reasonable starting point and included viable shelf temperatures and chamber pressures. This is significant, as it shows that AI can already contribute something useful at the beginning of a project. It can produce a shape for discussion, and in some cases that may save time. Where the limitations became more obvious was in the detail that separates a draft from a development-ready process. In the case study, the simulated recipe was less reliable in its representation of product temperature change over time, the length of the secondary drying step and the overall cycle duration. Those issues impact directly on cycle efficiency, process confidence and how practical the recipe would be in a real setting. This can also lead to product failure by providing process parameters that are not just inefficient, but unsuitable for the specific product being freeze-dried.

The reported outcome was that the expert-designed cycle produced visually excellent cakes using a process approximately half the length of the simulated AI cycle. That is the most telling part of the exercise. It suggests that AI can help establish a starting framework, but not yet the kind of optimised, experimentally informed process that experienced scientists can create when working from data and practical understanding. Put another way, the AI cycle was plausible, but the expert cycle was purposeful.

"Efficiency matters in lyophilisation," emphasises Dr. Ward. "A cycle that is unnecessarily long can affect capacity, cost and commercial viability. For a development team, the difference between a workable concept and an efficient, scalable process is often where the real value lies."

The case study also highlights a broader lesson. AI is most convincing when it produces outputs that are close enough to good practice to look dependable. That is precisely why disciplined use matters. When a result is broadly reasonable, the temptation is to move forward without sufficient challenge. In highly specialised development work, this is where risk begins. The real danger is that AI can generate answers that appear sufficiently credible to delay the moment when deeper scientific scrutiny is applied.

Where AI Fits Today

A more productive way forward is therefore neither uncritical enthusiasm nor blanket rejection. AI should be used in lyophilisation in proportion to what it can genuinely do




Figure 1: Graph showing simulate freeze-drying cycle conditions using ChatGPT 4

well. At present, that means helping teams begin faster, structure ideas more clearly and prepare for experimentation more efficiently. It does not mean bypassing formulation studies, critical temperature


determination, empirical cycle development or equipment-specific review. A sensible workflow is AI-assisted rather than AI-directed. The tool can support preparation and communication, but scientific decisions

should remain firmly anchored in data-led human judgment.

Viewed in that light, current machine learning models already have a meaningful place in freeze-drying cycle development. It can reduce friction at the start of a project and make some technical tasks more accessible. At the same time, lyophilisation remains a product-specific, evidence-driven process in which performance, quality and transferability cannot be secured through generic outputs alone. For all the progress AI has made, there is still no substitute for the combination of analytical testing, practical development knowledge and critical scientific interpretation that robust freeze-drying requires.


REFERENCES

1. 'Leveraging AI as a tool for freeze-drying process development' at <https://biopharmagroupcdmo.com/articles-resources/leveraging-ai-as-a-tool-for-freeze-drying-process-development/>



Freeze Drying Cycle Journey Pore Formulation Comparison

Ice crystal size drives pore structure, which determines drying performance



**Sam
Woodyard**

Sam Woodyard is a Research Scientist at Biopharma Group, specialising in freeze-drying formulation, cycle development, project management and pre- and post-lyophilisation analysis.



**Erin
Kuhn**

Erin Kuhn is Head of Pharma R&D Sales at Biopharma Group, supporting pharmaceutical companies with the handling, processability and scalability of liquid and freeze-dried formulations.



**Dr. Kevin
Ward**

Dr. Kevin Ward is Director of R&D at Biopharma Group, with extensive experience in freeze-drying process development, formulation support and analytical characterisation across pharmaceutical applications.