Cambridge Materials: Health (2026), **1**, e4, 1–12
doi:10.1017/S2978170126000038

CAMBRIDGE UNIVERSITY PRESS

Perspective

Toward a lab-in-the-loop AI platform for RNA therapeutics

Gen Li[1] and Bowen Li[1,2,3,4]

[1]Institute of Biomedical Engineering, University of Toronto, Toronto, ON, Canada; [2]Leslie Dan Faculty of Pharmacy, University of Toronto, Toronto, ON, Canada; [3]Department of Chemistry, University of Toronto, Toronto, ON, Canada and [4]Princess Margaret Cancer Center, University Health Network, Toronto, ON, Canada

Abstract

RNA therapeutics, spanning messenger RNA vaccines, small interfering RNAs, antisense oligonucleotides, and aptamers, are advancing rapidly, yet the design space that governs whether an RNA drug succeeds remains difficult to navigate. Its principal determinants, namely sequence, chemical modification, secondary structure, delivery vehicle, and tissue tropism, are strongly coupled, and together they constitute a combinatorial space far too large for human heuristics or isolated experiments to explore effectively. Here, we envisage a platform that brings two complementary advances into concert. First, RNA and molecular foundation models can compress this design space into a learnable representation that generalizes across RNA molecules and the carriers that deliver them. Second, a lab-in-the-loop architecture narrows the gap between prediction and biological function by coupling model-generated proposals to automated wet-laboratory experiments, allowing the model to be continually recalibrated against domain-specific feedback. Calibrated uncertainty connects these components: it enables the loop to prioritize candidates whose measurement is expected to be most informative. A platform built on these principles could support joint optimization of sequence and modification, co-design of payload and delivery vehicle, off-target prediction and in silico candidate ranking, and could ultimately bring the rational design of personalized therapies within reach.

Impact Statement

RNA medicines, including the mRNA vaccines deployed during the COVID-19 pandemic, as well as gene-silencing and gene-correcting therapies, have the potential to address diseases that were until recently considered undruggable, from infections and cancers to rare genetic disorders. Yet designing an effective RNA drug remains slow and uncertain: its therapeutic sequence, chemical modifications, folded structure, and delivery vehicle interact in ways too complex for trial-and-error experimentation to explore efficiently. This perspective sets out a vision for overcoming that bottleneck by uniting two rapidly maturing technologies: artificial intelligence models that learn the underlying rules of RNA design, and self-driving laboratories that test candidate molecules automatically. Linked in a continuous "lab-in-the-loop" cycle, the AI proposes designs, the laboratory measures them, and each result sharpens the next round of predictions. We outline the data, computational architecture, and safeguards needed to make such a platform trustworthy and to scale it responsibly toward the clinic. By making RNA design faster, cheaper, and more systematic, this approach could accelerate the development of personalized therapies and broaden access to a powerful class of medicines.

Keywords: RNA therapeutics; Artificial intelligence; Lab-in-the-loop

(Received: 09 July 2026; revised: 28 August 2026; accepted: 31 August 2026)

1. A Central Question and a Dilemma

A central aim of RNA therapeutics is to translate a therapeutic objective into a molecule that is simultaneously potent, stable, safe, and deliverable to its intended target. mRNA has matured into a validated drug class, exemplified by vaccines and protein-replacement therapies [1,2]. Beyond mRNA, small interfering RNAs (siRNAs) and antisense oligonucleotides (ASOs) silence or correct target transcripts with sequence specificity [3–6], aptamers adopt defined conformations that recognize protein targets in a manner analogous to antibodies [7] (Fig. 1A–D). Across these modalities, however, the efficacy of an RNA drug seldom derives from any single property. The sequence encodes the intended function (protein expression, transcript silencing, or splicing modulation) [8], yet its function is typically shaped by several determinants, including chemical modifications that confer stability while averting innate immune activation [9,10], secondary structure that modulates translation, half-life, and silencing efficiency [11,12,13], and, for many applications, a delivery vehicle, often a lipid nanoparticle (LNP), that determines biodistribution, cellular uptake, and the dose ultimately reaching the target [14–18]. Together, these tightly coupled determinants make RNA therapeutic design a high-dimensional optimization problem.

Conventional RNA design has relied on empirical rules, such as species-specific codon-usage tables, catalogs of well-characterized motifs, and trial-and-error screening of nanoparticle formulations [12,17–20]; these heuristics often operate in isolation, capture little of the nonlinear interplay among variables, and generalize poorly

Corresponding author: Bowen Li; Email: bw.li@utoronto.ca

Cite this article: G. Li and B. Li. 'Toward a lab-in-the-loop AI platform for RNA therapeutics'. *Cambridge Materials: Health*, **1**, e4: (2026), 1–12.
https://doi.org/10.1017/S2978170126000038

© The Author(s), 2026. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (http://creativecommons.org/licenses/by/4.0), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

Downloaded from https://www.cambridge.org/core. 17 Nov 2026 at 08:57:33, subject to the Cambridge Core terms of use.

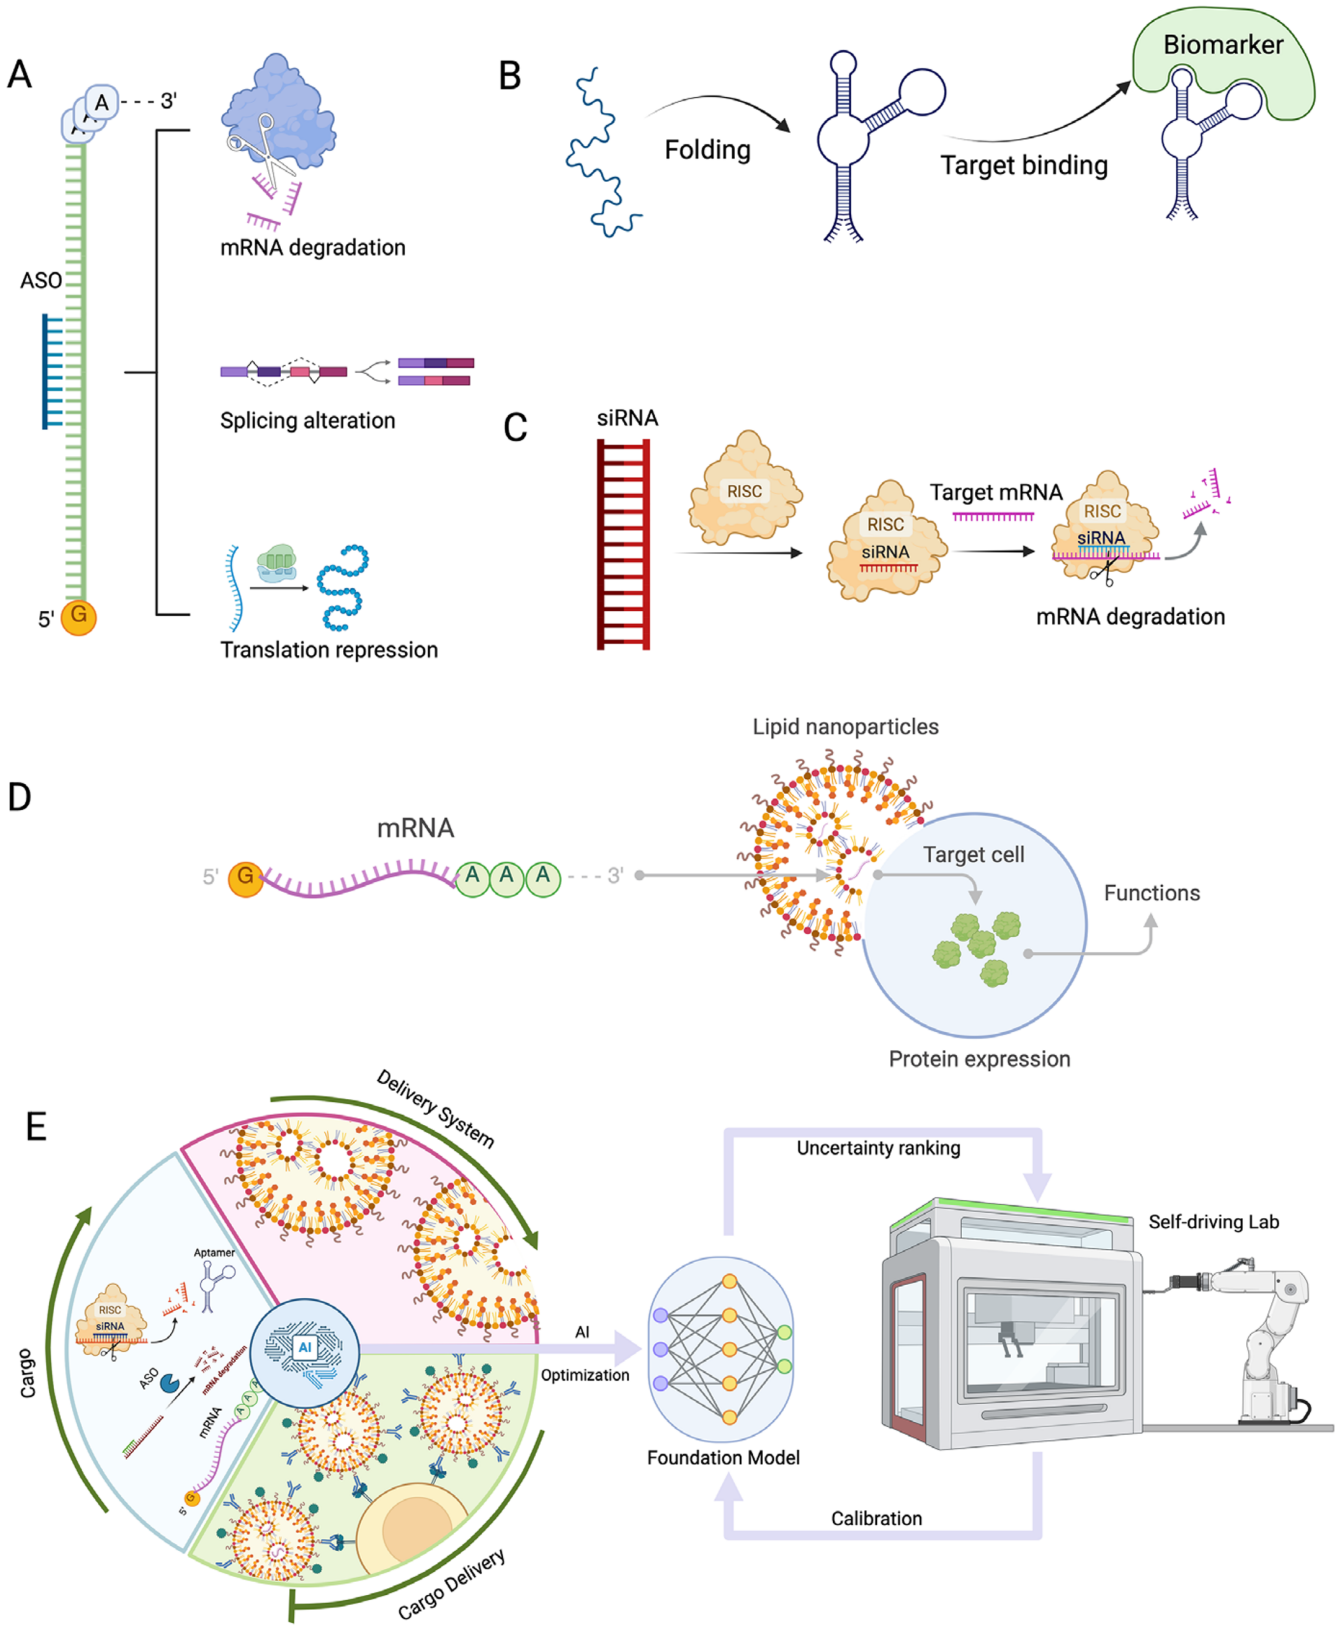


Figure 1. (A) ASOs hybridize to target mRNA through base pairing, leading to mRNA degradation, splice alteration, or translational repression. (B) Aptamers fold into a defined three-dimensional conformation that recognizes and binds a target (e.g., a biomarker) with high affinity and specificity, enabling molecular sensing or functional modulation. (C) An siRNA duplex is loaded into the RNA-induced silencing complex (RISC); the retained guide strand directs sequence-specific recognition and degradation of the target mRNA. (D) LNP-mediated delivery of mRNA to a target cell, where it is translated into a functional protein. (E) Closed-loop schematic of the lab-in-the-loop AI platform. Representations of diverse RNA cargos (e.g., siRNA, ASO, mRNA, and aptamers) and delivery vehicles (e.g., LNPs) are integrated by a foundation model for AI-guided optimization. At each round, candidates are selected by balancing exploitation against exploration and are dispatched to a self-driving laboratory for experimental validation. The resulting measurements are fed back to calibrate the model, forming an active-learning design-experiment-learn cycle.

Downloaded from https://www.cambridge.org/core. 17 Nov 2026 at 08:57:33, subject to the Cambridge Core terms of use.

across RNA classes. Two parallel advances now create an opening to move beyond them. Computationally, biological and molecular foundation models have matured across chemistry and genome biology, providing transferable priors for tasks involving nucleic acids and small molecules [21,22,23]. Experimentally, self-driving laboratories can now plan and execute experiments with minimal human intervention [24,25,26], while agentic artificial intelligence (AI) systems that organize large language models into multi-agent, generate–debate–rank workflows have begun to propose and prioritize novel research hypotheses [27]. Together, these developmentsmotivate a unified lab-in-the-loop platform for RNA therapeutics.

We therefore envisage integrating an AI model with a self-driving laboratory as a promising approach to RNA therapeutic development. The guiding strategy is to couple predictive and generative models with active learning and automated experimentation so that computation and experiment reinforce one another. Within this scheme, the AI module proposes candidate molecules and predicts their properties [28,29]; active learning then selects candidates for experimental validation by balancing exploitation, those with the highest predicted performance, against exploration, those whose evaluation is expected to reduce model uncertainty most [30,31]. The self-driving laboratory executes the chosen experiments, returning measurements as empirical feedback that recalibrate the model and correct its biases [24,25,26]. Operating recursively, the closed loop refines the model's design prior at each iteration and progressively narrows the search toward increasingly promising candidates (Fig. 1E).

The sections that follow develop this vision in turn. We first frame the foundation model and the loop as two halves of a single system and explain why neither suffices alone; we then survey the design tasks the platform could enable, from sequence and structure to delivery, before detailing the data requirements and computational components that the model and the loop each demand. We close by confronting the technical, evaluative, and ethical obstacles that remain and by sketching the collaborative ecosystem the vision will require.

2. Overview of the Platform

2.1. A Foundation Model Meets a Lab-in-the-Loop

Foundation models are deep neural networks pretrained on expansive data under self-supervised objectives, acquiring representations that transfer to a broad range of downstream tasks [32,33,34], and they have begun to reshape the modeling of nucleic acids and small molecules. Genomic language models such as DNABERT learn contextual representations of nucleotide sequence that transfer to regulatory-element and genetic-variant identification [35,36], whereas RNA-focused models such as 3UTRBERT capture the sequence determinants of mRNA subcellular localization and RNA-protein interactions within untranslated regions [37,38]; at the coding-sequence level, CodonBERT is pretrained on more than 10 million codon-tokenized mRNA coding sequences from diverse organisms and can be fine-tuned to predict properties such as protein expression, mRNA stability, and degradation, thereby supporting the evaluation and selection of coding sequences for mRNA vaccine and therapeutic design [39]; in parallel, molecular foundation models such as Uni-Mol learn three-dimensional representations of chemical structure that support property prediction [40]. Collectively, these advances support the feasibility of pretrained prior transfer within each modality. No existing model, however, jointly encodes the RNA payload together with the chemistry underlying its delivery. Bridging these hitherto separate priors is a central premise of the platform we propose.

A pretrained prior alone, however, is insufficient. The properties most critical to an RNA therapeutic, notably its in vivo efficacy and delivery performance, are difficult to learn from static datasets alone because they depend on biological context that no fixed corpus fully captures. This limitation motivates the lab-in-the-loop paradigm. A lab-in-the-loop architecture shifts the model from passive fitting to active exploration: the model proposes candidates, automated experiments evaluate them, and standardized results return to recalibrate the model [24–26,30] (Fig. 1E). The two components are complementary. Without the loop, a foundation model extrapolates unreliably to out-of-distribution designs, precisely the novel regions of greatest therapeutic interest; without a foundation model, the loop must explore afresh at every cycle, expending scarce experimental throughput on a sample-inefficient search. The foundation model thus supplies a transferable prior, and the loop supplies the continual, domain-specific supervision that keeps it calibrated.

2.2. Expected Characteristics and Architecture

Realizing this lab-in-the-loop vision places specific demands on the underlying model architecture. The platform should integrate several core capabilities within a shared representation: the modeling of sequence, the joint encoding of sequence and structure, an awareness of chemical modification, and a description of the delivery vehicle. Concretely, we envision modality-specific encoders that project their respective inputs into a shared latent space for further fusion of features. Transformer architectures offer a natural computational backbone for this task, because their attention mechanisms relate distant tokens directly and are therefore well suited to RNA, in which long-range base pairing, interactions among motifs, and the cross-modal coupling between a sequence and its delivery molecule are all consequential.

The platform should be generative as well as predictive: the models that score a candidate must also be able to propose and sample sequences, modification patterns, and delivery formulations conditioned on a target property profile. This generative capacity renders the loop proactive rather than merely evaluative, and it mirrors the trajectory of protein design, in which models progressed within a few years from scoring natural sequences to generating functional ones de novo [28,29,41]. For RNA, the generative target is broader still, as a useful proposal comprises a sequence together with its chemistry and its carrier rather than a sequence alone. Such proposals require the model to respect the hard constraints of synthesizability and manufacturability throughout exploration.

2.3. The Lab-in-the-Loop Workflow

The platform would supplement hypothesis-driven, small-scale iteration with data-driven, model-guided exploration. A single turn of the loop proceeds through five linked stages. The model first proposes candidate RNAs, selected not at random but by expected information gain or predictive uncertainty; automated synthesis then assembles the corresponding library, and high-throughput in vitro and in vivo assays measure their function. The resulting data are standardized, quality-controlled, and normalized to align with the pretraining corpus, after which the model is updated, and the cycle repeats [24–26,30,31] (Fig. 2A).

Downloaded from https://www.cambridge.org/core. 17 Nov 2026 at 08:57:33, subject to the Cambridge Core terms of use.

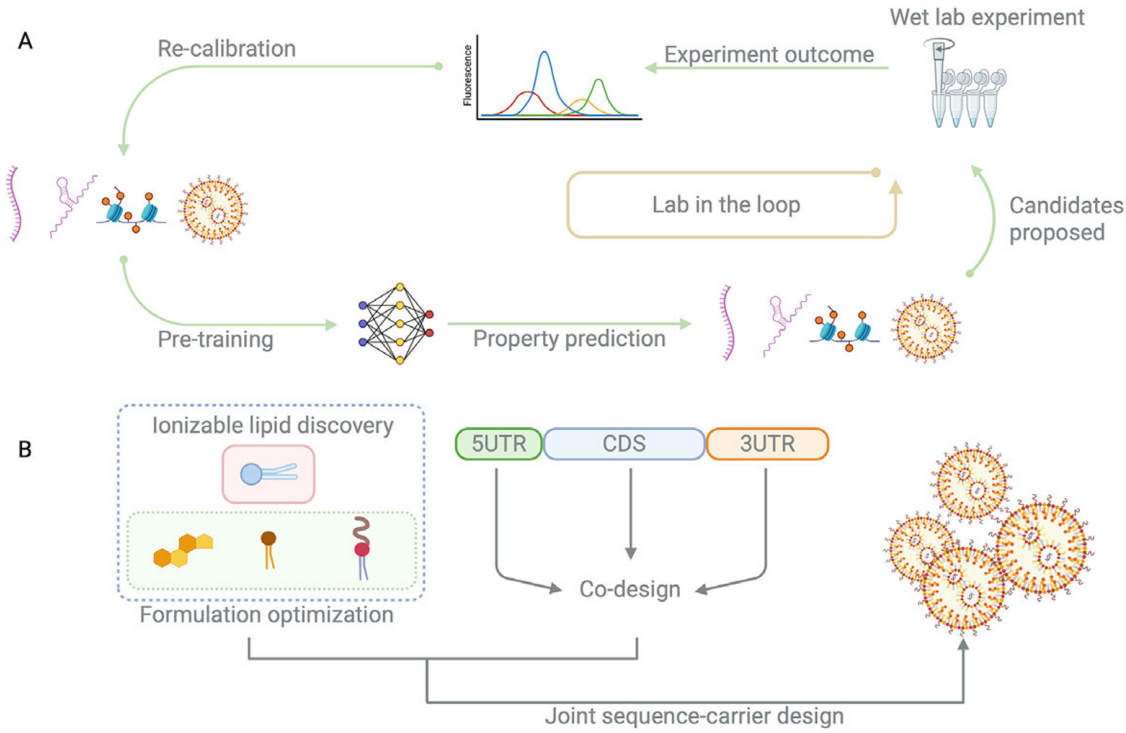


Figure 2. (A) A single turn of the lab-in-the-loop cycle. The foundation model is first pretrained on multimodal input (sequences, structures, modifications, and delivery molecules), yielding a prior that supports property prediction over proposed candidates. Candidates selected by the model using predicted performance and uncertainty are advanced to wet-lab experiments, whose functional outcomes are standardized and returned to recalibrate the model. The updated model proposes the next set of candidates, and the cycle repeats, with the foundation model serving as the persistent memory that links successive rounds. (B) Opportunities enabled by the platform. Delivery optimization comprises two distinct tasks: ionizable lipid discovery, which searches the combinatorial chemical space of the ionizable lipid, and formulation optimization across component ratios and process parameters. In parallel, the mRNA construct is optimized across its 5′ UTR, coding sequence (CDS), and 3′ UTR, which are conventionally tuned separately but are jointly optimized here. Coupling delivery optimization with sequence co-design yields joint sequence-carrier design, in which payload and delivery vehicle are optimized within a shared representation rather than sequentially.

The pivotal requirement is that the model emit probabilistic outputs rather than point predictions. Only a calibrated estimate of uncertainty enables the loop to address its central question, which candidates should be measured next, and thereby allocate limited experimental capacity where it most efficiently reduces uncertainty. This requirement binds the model to the loop because the model's epistemic uncertainty constitutes the loop's experimental-design objective. We anticipate that this cycle of proposal, measurement, and model update could become as central to RNA therapeutics as the design-build-test-learn cycle has become for synthetic biology.

3. Opportunities

3.1. Optimization of Sequence

Among the capabilities the platform enables, sequence optimization is the most immediate, and it can be pursued across several sequence scales. The model proposes mRNA candidates by jointly optimizing the 5′ and 3′ untranslated regions and the coding sequence, either as individual components or within the full-length construct (Fig. 2B), while the self-driving laboratory synthesizes and characterizes them automatically, returning quantitative measurements that feed directly into the next design round [13,38,42].

The principal advantage of this approach lies in joint optimization. Untranslated regions, codon choice, and secondary structure are conventionally tuned by separate heuristics, although they trade off against one another in practice [11,12,43]; by learning sequence-to-function relationships directly from assay data, the model navigates these coupled axes simultaneously rather than sequentially [42,44,45]. Experimentally, massively parallel reporter assays evaluate the proposed variants in parallel and return functional measurements within days, closing the design-build-test-learn loop [46].

However, sequence optimization is incomplete if treated as a search over nucleotide identity alone. A therapeutic RNA must simultaneously fold into a structure compatible with its function, resist degradation during storage and in circulation, and avoid triggering innate immune sensing; these three properties are governed by overlapping design variables rather than by independent ones. They also compete: for instance, pairing a nucleotide protects it from hydrolysis but may impede ribosomal scanning [47]; the chemical modifications introduced to suppress innate sensing simultaneously alter base-pairing thermodynamics and translational fidelity [48], as the +1 ribosomal frameshifting recently reported for N^1-methylpseudouridine illustrates [49]. Structure, stability, and immunogenicity must therefore be optimized jointly, and each enters the loop twice: as information the model conditions on, and as a measurement the loop returns.

As model input, all three properties are supplied as conditioning information, but they differ in how reliably that information can be obtained. Structure is the case where a prediction is readily available and most easily over-trusted, and we caution against conditioning on a single predicted minimum-free-energy structure. A long mRNA does not adopt one conformation but populates an

Downloaded from https://www.cambridge.org/core. 17 Nov 2026 at 08:57:33, subject to the Cambridge Core terms of use.

ensemble, of which the minimum-free-energy structure may represent only a small fraction. More seriously, the nearest-neighbor parameters underlying standard folding algorithms were calibrated on unmodified nucleotides and do not describe the modified bases used therapeutically [50], so the predicted structure is not merely incomplete but systematically miscalibrated. We therefore favor conditioning on base-pairing probability matrices, which retain the ensemble [51], and on experimentally measured reactivity profiles from high-throughput chemical probing [52], supplied directly as per-nucleotide features rather than converted into pseudo-energies for a thermodynamic folding step, which would reintroduce the same parameter dependence. Stability and immunogenicity admit no comparable predictor, and their determinants enter as explicit features instead: for stability, the modification pattern and the ensemble-averaged fraction of unpaired nucleotides, which sets the rate of hydrolysis [47]; for immunogenicity, the identity and position of modified nucleotides together with process metadata, which the sequence alone does not determine. Across all three, the principle is the same: predicted or derived quantities should be supplied as soft constraints carrying their own uncertainty, not as ground truth, consistent with the calibration requirements set out above.

As loop measurements, each property has a plate-scale readout: accelerated-degradation and integrity assays for stability, chemical probing for structure, and interferon-stimulated-gene reporters and cytokine panels for immunogenicity [13,53]. These belong to the densely queried layer of the assay hierarchy described below, with human reactogenicity remaining an escalated endpoint that in vitro immunogenicity predicts only imperfectly. Because the three properties trade off, they should enter the acquisition function as competing objectives on the Pareto front (Fig. 3C) rather than collapsing into a single expression score.

3.2. Co-Design of the Delivery Vehicle and Sequence

3.2.1. Ionizable Lipid Discovery

The ionizable lipid is a key determinant of the performance of an LNP, and its chemical space is vast: the combinatorial assembly of amine head groups, linkers, and lipid tails yields millions of enumerable structures [14,17,18,20,26,54]. Trained on structure–activity data relating lipid chemistry to delivery efficacy, the AI module proposes candidate lipids predicted to deliver efficiently while remaining within a prespecified physicochemical window [54]. A key advantage is that combinatorial chemistry makes a substantial fraction of this space experimentally accessible in parallel: multicomponent reactions such as the Ugi four-component reaction generate one product per combination of building blocks, so 10 variants of each component yield on the order of 10^4 distinct ionizable lipids from a single reaction scheme [26,54]. Downstream evaluation is comparably scalable, as plate-reader-based transfection assays can quantify many formulations in parallel [26,54]. This efficiency, however, is confined to chemistries amenable to combinatorial assembly. For lipids that are low-yielding or require multistep synthesis, synthesis and wet-lab validation remain rate-limiting rather than negligible steps. This synthesis step is itself increasingly a target for machine learning, and folding it into the loop turns synthesizability from an afterthought into a design variable [55]. Data-driven retrosynthesis and reaction-outcome prediction can propose plausible routes to a candidate scaffold and anticipate which steps will be low-yielding or selectivity-limited [56,57], while reaction-condition recommendation and Bayesian reaction optimization tune temperature, stoichiometry, and catalyst choice with far fewer experiments than a factorial screen [58–62]. Integrated with automated experimentation, in which AI planning is coupled to robotic synthesis and online analysis [59,63], these methods let the "how to make it" be reasoned about in the same computational pass as the "what to make." Embedded in the platform, they allow the hard constraint we noted above, that a proposed lipid be not merely potent but synthesizable and manufacturable at acceptable cost, to be enforced by the model during exploration rather than discovered only after a candidate reaches the bench.

The complementary advance is on the design side, where combinatorial chemistry and machine learning together make a vast lipid space experimentally traversable. Combinatorial lipidoid and ionizable-lipid libraries, assembled from modular amine, linker, and tail building blocks, have long yielded thousands of structures from a single reaction scheme [20,54], and recent work couples this synthetic accessibility to machine-learning models that predict delivery-relevant properties directly from lipid structure, so that the fraction of the space actually synthesized and assayed is concentrated on the regions the model deems most promising [18,64]. Read together, these two lines close the design-make-test cycle on both of its axes: models that decide which ionizable lipid to pursue, and models that decide how to synthesize it, each supplying the calibrated predictions the loop needs to allocate its scarce synthesis and assay capacity. This coupling is precisely what a self-driving laboratory operationalizes, and it is where the loop's advantage over a sequential, one-lipid-at-a-time medicinal-chemistry campaign is largest.

3.2.2. Formulation Optimization

Beyond the lipid structure itself, the model can optimize the formulation of an LNP. The molar ratios of ionizable lipid, helper lipid, cholesterol, and PEG-lipid each shift particle size, encapsulation efficiency, stability, and biodistribution [14,16,65]; treating the formulation as a continuous, model-navigable space allows the platform to search component ratios and process conditions jointly, conditioned on the target tissue, rather than varying one factor at a time [15]. In addition, formulation ratios are only one determinant of the final product's critical quality attributes. Buffer identity and ionic strength during mixing [66], the manufacturing method and its process parameters [67], and the purity of the lipid inputs all contribute, and lipid quality is itself variable, as oxidative degradation of the ionizable lipid generates reactive aldehydes that form adducts with the mRNA payload and abolish translation [68,69], altering potency independently of the nominal formulation. Treating formulation as a continuous, model-navigable space is therefore useful only if that space is defined to include process and material variables alongside composition, and if lot-level characterization is recorded with each measurement. This is a direct instance of the requirement, developed below, that experimental metadata be treated as a first-class model input. It also implies a scale-dependent domain shift: a process-property mapping learned from bench-scale synthesis need not hold under manufacturing regimes [67,70].

A deeper limitation is that biodistribution cannot be predicted from formulation chemistry alone. On contact with biological fluids, an LNP acquires a protein corona that constitutes its biological identity, and it is this adsorbed layer, rather than the synthetic surface, that the body largely recognizes [71]: apolipoprotein E adsorption and LDL-receptor engagement constitute the canonical route to hepatocyte uptake [72], and deliberate manipulation of surface charge redirects tropism by reshaping corona composition [15,73]. Downstream, the fraction of an administered

Downloaded from https://www.cambridge.org/core. 17 Nov 2026 at 08:57:33, subject to the Cambridge Core terms of use.

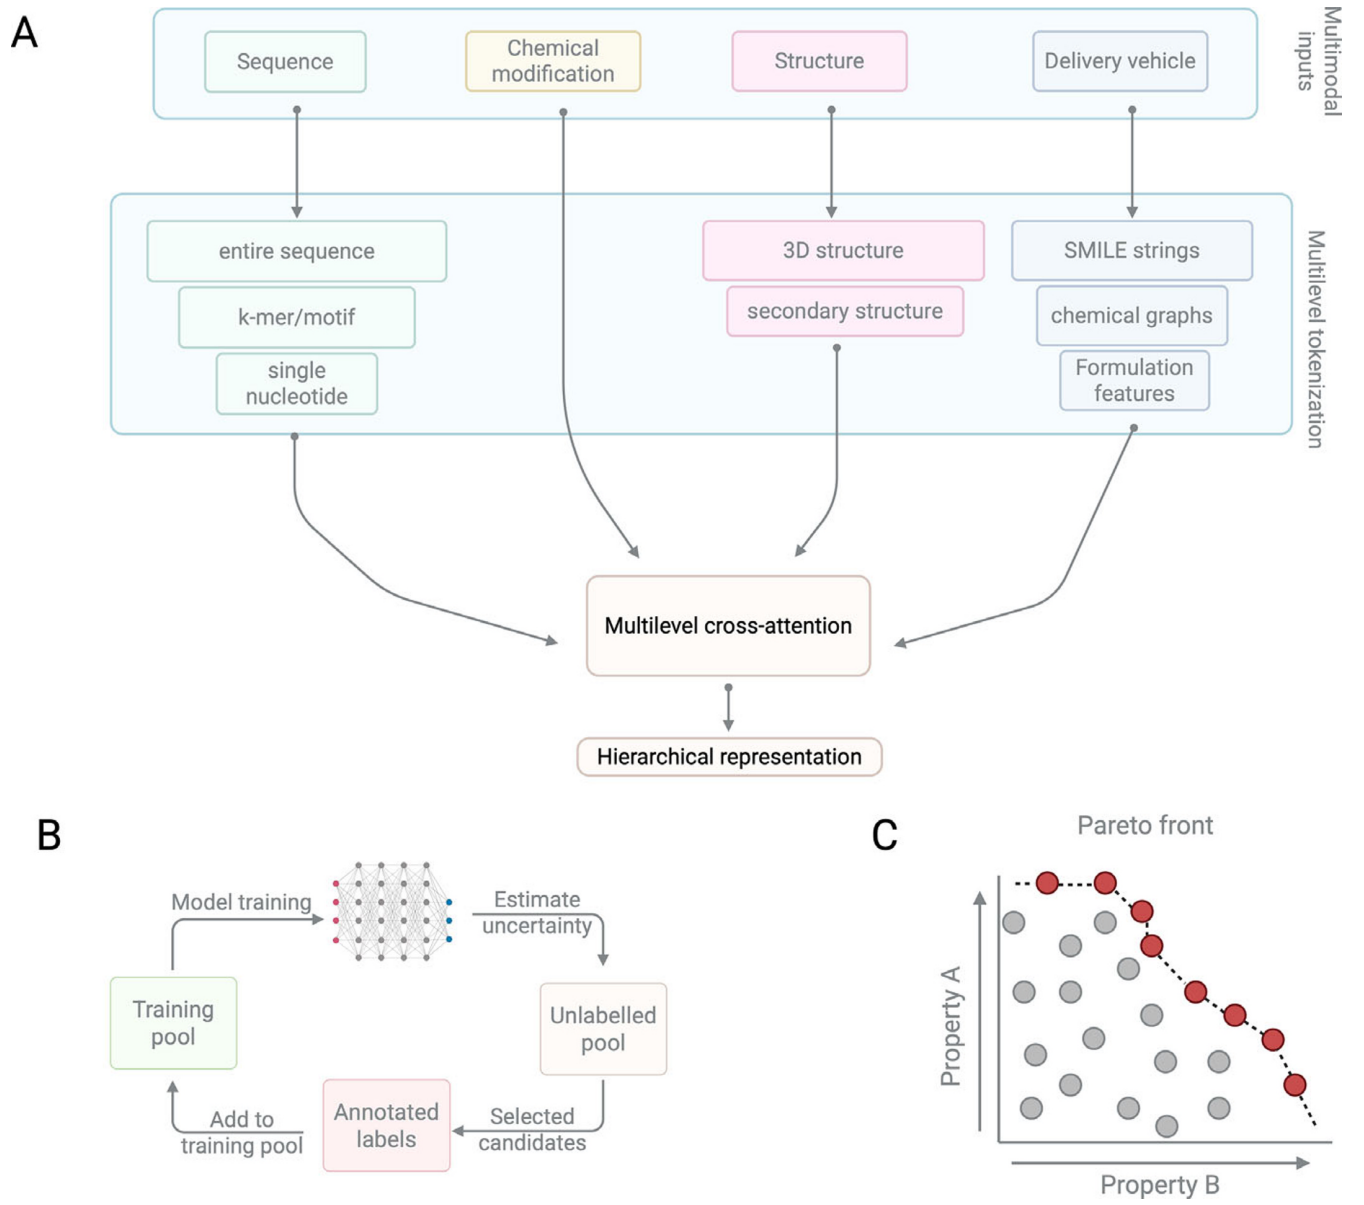


Figure 3. (A) Multimodal, multilevel representation. Four input modalities, biological sequence, chemical modification, structure and delivery vehicle, are tokenized at multiple levels of granularity: sequence is resolved from single nucleotides through k-mers and motifs to the entire sequence; structure is represented at both secondary and three-dimensional levels; and the delivery vehicle is encoded as SMILES strings and chemical graphs. In the proposed architecture, these representations are integrated early through multilevel cross-attention, yielding a shared hierarchical representation in which a sequence and its carrier are reasoned over jointly. (B) Active-learning loop. A model trained on the current training pool estimates its predictive uncertainty over an unlabeled pool of candidates; the most informative candidates are selected for experimental annotation, and the resulting labels are added back to the training pool. Iterating this cycle concentrates experimental effort where it most reduces the model's uncertainty. (C) Multi-objective selection. Since an RNA drug must satisfy multiple competing properties simultaneously, candidate selection is inherently multi-objective. Rather than maximizing a single scalar, the loop traces the Pareto front (red), the set of non-dominated candidates for which no property can be improved without degrading another, here shown for two representative properties.

dose producing a functional outcome is set above all by endosomal escape, which remains inefficient and is not predicted by the amount of material internalized [74,75]. A model mapping composition directly to organ-level expression is thus fitting a heavily confounded relationship, and the joint representation proposed above should be understood as mediated by these biological variables rather than as bypassing them.

Incorporating this layer into an automated loop is genuinely difficult, and we do not claim it can be fully automated in the near term. The relevant measurements are animal-dependent, slow, and high in inter-subject variance; their throughput is bounded by ethical and logistical constraints that no increase in computational capacity relaxes. The loop must therefore be stratified rather than uniform: in vitro potency, particle characterization and ex vivo corona proteomics can be sampled densely and serve as intermediate variables that partially explain in vivo outcome, while in vivo biodistribution and functional endpoints are reserved for an outer, human-supervised cycle.

Critically, the escalation rule described below requires knowing how well each cheap layer predicts the expensive endpoint, and for LNPs this *in vitro-in vivo* correlation is weak and context-dependent [76]. It should be estimated and reported as a measured

Downloaded from https://www.cambridge.org/core. 17 Nov 2026 at 08:57:33, subject to the Cambridge Core terms of use.

quantity within the platform. Designing acquisition strategies that remain useful under the resulting small in vivo budget is the principal open problem in extending lab-in-the-loop design from formulation chemistry to in vivo delivery.

3.2.3. Joint Sequence-Carrier Design

The deeper opportunity lies in co-designing the RNA payload together with its delivery vehicle. The physicochemical properties of the cargo govern its encapsulation and release, while the carrier, in turn, determines the effective dose and intracellular exposure kinetics [77,78]. By representing sequence and formulation within a shared feature space, the platform can model their interactions and predict the resulting biodistribution and expression, thereby capturing the payload-carrier interactions that a sequential design workflow would miss (Fig. 2B). To make co-design operational, the joint representation must be trained on paired measurements that link a molecular triplet, RNA cargo, ionizable lipid, and formulation composition, to a set of delivery-relevant readouts spanning three coupled layers: organ- and tissue-level biodistribution, cell-type-resolved uptake, and the resulting functional expression of the encoded cargo. Concretely, each training instance is a tuple of a cargo representation (a sequence- or structure-derived embedding spanning mRNA, siRNA, and ASOs), a lipid representation (molecular descriptors or a learned chemical embedding), and a formulation vector (molar ratios, N/P ratio, and helper-lipid identity), mapped to a multi-task readout vector. The representation is learned by requiring these three component encoders to predict the paired biodistribution, uptake, and expression readouts through a shared latent space, so that the latent space captures how cargo, chemistry, and formulation jointly determine delivery rather than treating them as independent factors. Because the space is combinatorially large, exhaustive measurement is infeasible. We therefore cast the problem as tensor completion over this space, regularized by the learned component representations. Unmeasured combinations are inferred not from co-occurrence alone but from generalization across chemically and structurally related components; a lipid unmeasured with a given cargo can be positioned via its behavior with related cargos. Coverage is then allocated adaptively: rather than sampling the matrix uniformly, an active-learning loop selects the combinations whose measurement most reduces predictive uncertainty across the joint space, prioritizing regions where the model is both uncertain and consequential for the co-design objective. This converts an intractable exhaustive matrix into a sparse, iteratively refined design that remains predictive over the full combinatorial space.

4. Toward Building the Platform

4.1. Data: Pretraining Data and Loop Data

The central modeling obstacle is the scarcity and dispersion of paired measurements linking RNA sequence and formulation to function. Substantial static resources exist, including encyclopedias of genomic elements and structure repositories [79,80,81], yet they seldom pair a sequence with the therapeutic readouts of interest and remain fragmented across RNA classes and species. Integration of data across RNA modalities and across organisms will therefore be essential if the model is to acquire the breadth from which transferable priors emerge [23,36].

A complementary obstacle arises on the experimental side. Data generated over successive rounds of experimentation must be standardized, quality-controlled, and normalized before they can recalibrate the model, and the associated pipelines must harmonize newly acquired measurements with the pretraining corpus through shared metadata, controls, and normalization [26]. This harmonization, however, is not a neutral preprocessing step. A central difficulty is distinguishing the correction of technical batch effects from genuine distribution shift. Successive rounds of active learning deliberately query underexplored regions of chemical or sequence space, so each new batch is expected to differ from the pretraining corpus not only technically but biologically. Batch-correction methods that assume a shared underlying distribution across batches (global scaling, quantile normalization, or purely data-driven surrogate-variable approaches) risk absorbing this genuine novelty as if it were technical variation, over-normalizing away precisely the signal the loop is designed to discover. This risk is greatest exactly for the novel, sparsely sampled points that active learning tends to select, since these are the easiest to distort during normalization. Mitigating this requires anchoring technical variation to explicit references rather than inferring it from the data alone. Bridging or spike-in controls carried across every round allow batch effects to be estimated without conflating them with biological shift, and hierarchical or mixed-effects formulations can model technical and biological variance jointly rather than subtracting the former. Framed in terms of uncertainty, the goal is to attribute technical variability to aleatoric noise while preserving out-of-distribution novelty as epistemic uncertainty, so that the acquisition strategy pursues genuine discovery rather than chasing technical noise. Where paired data are scarce, synthetic data generated by mechanistic or generative models can serve as a complementary signal, provided that their biases are characterized and that empirical measurements remain the final arbiter.

A subtle but consequential manifestation of this problem is the batch effect between newly generated data and the historical corpus. Measurements produced by different laboratories, instruments, and protocols carry systematic offsets that, if left uncorrected, may be misread by the model as biological signal. The platform should therefore treat experimental metadata, including reagent lot, assay date, operator, and instrument, as first-class inputs and incorporate spike-in controls and shared reference standards in every batch, so that the system recalibrates against a fixed anchor rather than drifting with each round of measurement.

4.2. Computational Components for the Foundation Model

The foundation model is shaped above all by how it represents scale, modality, and prior knowledge. Because RNA is organized across many scales, a tokenization that ascends from nucleotides through k-mers and motifs to structural elements, chemical modifications and ultimately the full-length construct can capture this hierarchy [23,36]. Integration across modalities is equally essential: biological sequence, chemical modification, delivery molecules (representable as chemical graphs or string encodings), and structures must occupy a shared representation [40] (Fig. 3A). The strategy of modality fusion, however, remains an open design question. Early fusion, integrating modalities at an early representation stage so that the model reasons about a sequence and its carrier in concert rather than in series, is attractive when the modalities are mechanistically coupled, as sequence and delivery vehicle plausibly are, since their interactions can be modelled directly instead of at the output. However, early fusion increases the risk that a data-rich modality dominates a data-scarce one, and it can be brittle when one modality is missing at inference. Late fusion, which combines modalities only at or near the output stage, is more robust to

Downloaded from https://www.cambridge.org/core. 17 Nov 2026 at 08:57:33, subject to the Cambridge Core terms of use.

missing modalities and to imbalanced data, while hybrid schemes that share representations only at selected layers offer intermediate trade-offs. We propose early fusion as the primary hypothesis (Fig. 3A) but benchmark it directly against late and hybrid alternatives, using predictive performance on downstream therapeutic-property tasks under matched parameter scale as the primary comparison criterion, with robustness to missing modalities and to modality-scale imbalance as secondary criteria. Finally, mechanistic prior knowledge can constrain learning: thermodynamic base-pairing rules, curated RNA-motif databases and the structure–activity relationships of delivery vehicles can be introduced as inductive biases that ground an otherwise purely data-driven model in mechanism [11,82].

4.3. Computational Components for the Loop

The loop demands its own computational machinery, which is what distinguishes the platform from a conventional foundation model. At its core lie active learning, Bayesian experimental design and uncertainty-guided selection (Fig. 3B), which determine which candidates to measure, while batch acquisition proposes a set of candidates at once, matching the throughput of parallel assays rather than evaluating a single candidate per round [30,62]. Because an RNA drug must satisfy efficacy, stability, safety, and manufacturability simultaneously, this selection is inherently multi-objective, requiring the loop to trace a Pareto front rather than to maximize a single scalar objective (Fig. 3C). These procedures rest on uncertainty quantification, which converts the model's epistemic uncertainty into a decision rule for experimentation and thereby binds the model to the loop once more. They also depend on automation infrastructure, namely, self-driving synthesis and assay platforms, that supplies the throughput without which the loop stalls [24,25,26]. A further principle is to match the granularity of acquisition to the cost of each assay: inexpensive, rapid readouts such as reporter expression are queried densely to map coarse structure, whereas costly, slow readouts such as in vivo biodistribution are reserved for the few candidates that the cheaper layers cannot resolve. Making this escalation decision well requires two things. First, the uncertainty reported at each layer must be trustworthy: when the model predicts a reporter-expression value with, say, 90% confidence, that prediction should be correct 90% of the time, which can be verified on held-out candidates and corrected if it is not. Second, a candidate should be escalated to an expensive in vivo assay not simply because its cheap readout is uncertain, but because measuring it is expected to resolve real uncertainty about the in vivo outcome. This requires knowing how well the cheap readout actually predicts the expensive one: a reporter-expression value that is confidently predicted but only loosely related to biodistribution should not trigger a costly in vivo experiment. The loop therefore escalates a candidate when the expected gain in knowledge about the high-fidelity endpoint, relative to the cost of the assay, is greatest. The loop thereby becomes a hierarchy of nested cycles rather than a single feedback line, with information propagating upward from high-throughput surrogates to low-throughput, high-fidelity endpoints, and with the foundation model serving as the shared memory that binds the layers together.

4.4. Agentic Reasoning as the Proposal Layer

The loop's proposal step need not be a purely numerical optimizer. A parallel line of work reframes candidate generation as an agentic process, in which large language models are organized into cooperating specialist agents that reason over the scientific literature, formulate hypotheses, and critique one another before committing to an experiment. Co-Scientist exemplifies this paradigm: a coalition of generation, reflection, ranking, evolution, and meta-review agents advances hypotheses through tournament-style scientific debate, with an Elo-based ranking and additional test-time computation devoted to the most promising proposals. The system has produced experimentally supported hypotheses in drug repurposing, novel-target discovery, and the mechanism of antimicrobial resistance, in one case recapitulating an unpublished experimental finding [27]. Such systems complement the foundation-model-plus-loop architecture we advocate. Their strength is breadth and interpretability, the ability to traverse a large hypothesis space and to attach an explicit, literature-grounded rationale to each proposal. Their weakness is that language-model reasoning alone is not calibrated against wet-laboratory outcomes and is prone to fluent but unfounded claims. The foundation model supplies precisely what the agents lack, a quantitative, uncertainty-aware estimate of how a candidate is likely to behave, and the self-driving laboratory supplies the empirical verdict that neither the agents nor the model can substitute.

Adapted to RNA therapeutics, an agentic layer could sit at the front of each loop iteration and take on the tasks least amenable to a numerical objective. It could translate a therapeutic goal into concrete design hypotheses, for instance proposing a chemical-modification pattern to evade innate immune sensing, a 5′ UTR motif to raise translation, or an ionizable-lipid scaffold for a target tissue, together with the mechanistic argument for each. It could reconcile these proposals against prior loop feedback and the published literature, flagging where a suggestion contradicts an earlier measurement, and it could nominate the experiments expected to be most discriminating, which the foundation model then scores and ranks by calibrated uncertainty before the self-driving laboratory executes them. Realizing this safely requires the same guardrails we develop below: every agent-generated hypothesis must remain experimentally falsifiable, be traceable to the evidence that motivated it, and be gated by calibrated uncertainty rather than by the rhetorical confidence of the language model that proposed it. Framed this way, agentic reasoning and the calibrated foundation model occupy distinct and complementary roles, one supplying interpretable hypothesis breadth and the other the quantitative ranking on which the loop's acquisition decisions ultimately depend.

5. Challenges and Limitations

5.1. Data and Automation Infrastructure

The computational loop can advance only as quickly as the automated laboratory can supply reliable data. Self-driving laboratories meet this demand but remain costly and unevenly accessible [24,25,26]; paired sequence-to-function data are still scarce; and the computational expense of pretraining large models constrains who can participate. Low-resource training and inference techniques, together with shared infrastructure, will therefore be essential to broaden access, particularly for academic laboratories operating under constrained computational budgets.

5.2. Evaluation Methodology

The field of RNA therapeutics lacks standardized benchmarks, directly limiting method comparison and model selection.

Downloaded from https://www.cambridge.org/core. 17 Nov 2026 at 08:57:33, subject to the Cambridge Core terms of use.

Feedback drawn from different layers of the loop carries different degrees of reliability, since an in vitro reporter readout, an in vivo biodistribution measurement, and a clinical endpoint are not interchangeable; the translation gaps between them oblige the loop to weight its feedback accordingly [75,83]. Benchmarks for this setting should therefore reward what the loop actually optimizes, not agreement with fixed labels. Four metrics are well suited. Prospective hit rate, the fraction of proposed candidates that succeed when tested in the next round, is the quantity the loop most directly aims to improve, and cannot be inflated by fitting existing data. Out-of-distribution generalization measures performance on chemistries or sequences absent from training, which matters because active learning deliberately explores such regions. Calibration error checks whether the model's stated confidence matches its actual accuracy, a prerequisite for the uncertainty-guided selection described above. Pareto hypervolume measures how well the proposed candidates cover the trade-offs among efficacy, stability, safety, and manufacturability, rather than collapsing them into one score. A further caution concerns evaluation metrics anchored to established design heuristics, though it requires a distinction we did not draw sharply enough. Some accumulated design knowledge rests on well-characterized biophysics: the apparent pKa window within which an LNP remains near-neutral in circulation yet protonates in the acidifying endosome [84,85], and the molecular geometry favoring non-bilayer phases that destabilize the endosomal membrane [86,87], are mechanistic constraints on ionization, packing, and self-assembly, supported by structure–activity data across many chemistries [84,88]. They belong in the platform as inductive biases and should be evaluated as critical quality attributes in their own right. Our caution is narrower and concerns the more specific rules that accompany them, such as a preferred tail length or head-group class [88], which largely describe the regions of chemical space that existing libraries happened to cover [84,88]. Rewarding agreement with these penalizes a model precisely when it proposes a design outside that region [89,90]. The two cases are empirically separable: a mechanistic constraint should continue to govern outcomes as the loop moves into unexplored space, whereas a sampling artifact should not. We therefore advocate encoding mechanistic constraints as priors, treating descriptive regularities as testable hypotheses rather than scoring rules, and grounding evaluation in prospective outcomes.

5.3. Interpretability and Hallucination Risk

In a therapeutic context, the platform's interpretability and trustworthiness are essential. Any candidate the platform proposes must be experimentally falsifiable, and the model must be able to acknowledge the limits of its knowledge; in a patient-safety context, this explicit communication of uncertainty constitutes a hard constraint [33,91,92]. Large language models illustrate both the promise and the peril, as adapted models can match domain experts on certain tasks yet remain prone to fluent and confident error [93]. We accordingly require that the platform's outputs be traceable to training data and loop feedback, remain consistent with the specified context, and be accompanied by calibrated uncertainty. To be actionable rather than aspirational, these principles must be instantiated as auditable safeguards. Traceability becomes provenance reporting: each proposed candidate carries a record of the training data and prior measurements that most influenced it, so the proposal can be inspected rather than trusted. Calibrated uncertainty becomes an explicit threshold, above which a prediction is not advanced automatically but flagged for human review or additional measurement. This defines an abstention criterion: the platform declines to make a confident recommendation when its uncertainty is high or a candidate falls outside its training space, an explicit declaration of uncertainty that is safer than a confident error. Finally, escalation between assay layers is governed by predefined decision criteria, pre-registered criteria on efficacy, safety, and uncertainty that a candidate must meet before advancing to a costlier assay, so that each escalation can be audited against the criterion that authorized it.

5.4. Feedback Loop Bias

A risk that is especially acute in a closed loop is the distribution shift that the loop induces in itself. Should the model propose only candidates in regions where it is already confident, the loop will reinforce that confidence, starve the model of informative failures, and systematically overlook novel designs; this is the dynamic analogue of an evaluation metric that penalizes novelty. Mitigating this danger calls for deliberate exploration, uncertainty-seeking acquisition and the periodic inclusion of diverse or adversarial candidates, so that the loop broadens the model's competence rather than narrowing it.

5.5. Open Science and Ethical Considerations

Where privacy, security and intellectual-property constraints permit, models and curated data should be openly accessible, accompanied by clear statements of their capabilities, limitations and intended use [94]. The translation of AI-designed RNA into the clinic raises questions of intellectual property, regulatory oversight and the validation of predictions on clinical cohorts before use in patient care; the recent progress of computationally designed drugs into clinical testing renders these questions concrete rather than hypothetical [95]. Patient data privacy and representativeness across populations must likewise be safeguarded, lest the platform encode or amplify existing inequities.

5.6. A Future of Collective Innovation

The development of a lab-in-the-loop AI platform will require a collaborative ecosystem of RNA biologists, chemists, AI researchers, automation engineers and clinical scientists, each contributing data, models and validation. The foundation model distills what constitutes a good RNA design into transferable knowledge, while the lab-in-the-loop keeps that knowledge anchored to the evolving reality of the wet laboratory; together they can carry RNA therapeutics from a craft of trial and error toward a more predictable and scalable engineering discipline. The most compelling applications are already in view: personalized RNA neoantigen vaccines tailored to the individual patient [96,97,98], therapies for rare diseases underserved by conventional development models [99], and rapid-response vaccine platforms that can be reprogrammed against emerging pathogens on compressed timelines [2,100]. These prospects are nonetheless bounded. The advantage for rare diseases is clearest where a well-defined molecular target exists, but patient numbers are too small to justify conventional development economics [99]; it is more limited where the underlying mechanism is poorly characterized or where data remain too sparse for a model to generalize. Likewise, the promise of rapid-response platforms applies principally to the design phase: sequence and formulation can be reprogrammed quickly, but this compresses only one segment of the path from pathogen to approved product [101]. A realistic appraisal must therefore separate what the platform

Downloaded from https://www.cambridge.org/core. 17 Nov 2026 at 08:57:33, subject to the Cambridge Core terms of use.

accelerates from what it does not. The loop can shorten the design-test cycle and improve the quality of candidates entering development, but manufacturing, regulatory review, and clinical validation follow their own timelines and are largely independent of it [102]. Good manufacturing practice production and scale-up, the evidentiary requirements of regulatory approval, and the multi-year safety and efficacy testing of human trials remain rate-limiting regardless of how efficient candidate generation becomes. The platform's contribution is to ensure that the candidates entering these stages are better characterized and more likely to succeed, not to bypass the stages themselves. Realizing this future will require a shared scientific community in which models learn from, and experiments continually test, each successive design.

Open peer review. To view the open peer review materials for this article, please visit http://doi.org/10.1017/S2978170126000038.

Author Contributions. G.L. and B.L. conceptualized the perspective and framed the entire manuscript. G.L. and B.L. wrote the manuscript. B.L. acquired funding and supervised the project.

Financial Support. The authors acknowledge financial support from the GSK Chair Professorship, the Leslie Dan Faculty of Pharmacy Startup Fund, the Princess Margaret Cancer Centre Operating Fund, the Canada Research Chairs Program (CRC-2022-00575), the Canadian Institutes of Health Research (PJH-185722, PJT-195669), Cystic Fibrosis Canada (1188219), the New Frontiers in Research Fund (NFRFE-2023-00203), and the Natural Sciences and Engineering Research Council of Canada (RGPIN-2023-05124).

Competing Interests. The authors declare that they have no competing interests.

Data Availability Statement. This study did not involve the novel creation, generation, or analysis of primary research materials. All relevant literature is cited in the References.

References

1. U. Sahin, K. Karikó and Ö. Türeci, 'mRNA-based therapeutics — developing a new class of drugs'. *Nature Reviews. Drug Discovery*, 13(10): (2014), 759–780. https://doi.org/10.1038/nrd4278.
2. N. Chaudhary, D. Weissman and K. A. Whitehead, 'mRNA vaccines for infectious diseases: principles, delivery and clinical translation'. *Nature Reviews Drug Discovery*, 20(11): (2021), 817–838. https://doi.org/10.1038/s41573-021-00283-5.
3. R. L. Setten, J. J. Rossi and S. P. Han, 'The current state and future directions of RNAi-based therapeutics'. *Nature Reviews Drug Discovery*, 18(6): (2019), 421–446. https://doi.org/10.1038/s41573-019-0017-4.
4. S. T. Crooke, B. F. Baker, R. M. Crooke and X. H. Liang, 'Antisense technology: an overview and prospectus'. *Nature Reviews Drug Discovery*, 20(6): (2021), 427–453. https://doi.org/10.1038/s41573-021-00162-z.
5. A. Fire, S. Xu, M. K. Montgoery et al. Potent and specific genetic interference by double-stranded RNA in *Caenorhabditis elegans*. *Nature*, 3916669: (1998), 806–811. https://doi.org/10.1038/35888.
6. S. M. Elbashir, J. Harborth, W. Lendeckel et al., 'Duplexes of 21-nucleotide RNAs mediate RNA interference in cultured mammalian cells'. *Nature*, 411(6836): (2001), 494–498. https://doi.org/10.1038/35078107.
7. J. Zhou and J. Rossi, 'Aptamers as targeted therapeutics: current potential and challenges'. *Nature Reviews Drug Discovery*, 16(3): (2017), 181–202. https://doi.org/10.1038/nrd.2016.199.
8. S. T. Crooke, J. L. Witztum, C. F. Bennett and B. F. Baker, 'RNA-targeted therapeutics'. *Cell Metabolism*, 27(4): (2018), 714–739. https://doi.org/10.1016/j.cmet.2018.03.004.
9. N. Pardi, M. J. Hogan, F. W. Porter and D. Weissman, 'mRNA vaccines — a new era in vaccinology'. *Nature Reviews. Drug Discovery*, 17(4): (2018), 261–279. https://doi.org/10.1038/nrd.2017.243.
10. M. Schlee and G. Hartmann, 'Discriminating self from non-self in nucleic acid sensing'. *Nature Reviews Immunology*, 16(9): (2016), 566–580. https://doi.org/10.1038/nri.2016.78.
11. Y. Wan, K. Qu, Q. C. Zhang et al., 'Landscape and variation of RNA secondary structure across the human transcriptome'. *Nature*, 505(7485): (2014), 706–709. https://doi.org/10.1038/nature12946.
12. V. Presnyak, N. Alhusaini, Y. H. Chen et al., 'Codon optimality is a major determinant of mRNA stability'. *Cell*, 160(6): (2015), 1111–1124. https://doi.org/10.1016/j.cell.2015.02.029.
13. K. Leppek, G. W. Byeon, W. Kladwang et al., 'Combinatorial optimization of mRNA structure, stability, and translation for RNA-based therapeutics'. *Nature Communications*, 13(1): (2022), 1536. https://doi.org/10.1038/s41467-022-28776-w.
14. X. Hou, T. Zaks, R. Langer and Y. Dong, 'Lipid nanoparticles for mRNA delivery'. *Nature Reviews Materials*, 6(12): (2021), 1078–1094. https://doi.org/10.1038/s41578-021-00358-0.
15. Q. Cheng, T. Wei, L. Farbiak et al., 'Selective organ targeting (SORT) nanoparticles for tissue-specific mRNA delivery and CRISPR–Cas gene editing'. *Nature Nanotechnology*, 15(4): (2020), 313–320. https://doi.org/10.1038/s41565-020-0669-6.
16. M. J. Mitchell, M. M. Billingsley, R. M. Haley et al., 'Engineering precision nanoparticles for drug delivery'. *Nature Reviews Drug Discovery*, 20(2): (2021), 101–124. https://doi.org/10.1038/s41573-020-0090-8.
17. J. Chen, Y. Xu, M. Zhou et al., 'Combinatorial design of ionizable lipid nanoparticles for muscle-selective mRNA delivery with minimized off-target effects'. *Proceedings of the National Academy of Sciences*, 120(50): (2023), e2309472120. https://doi.org/10.1073/pnas.2309472120.
18. B. Li, I. O. Raji, A. G. R. Gordon et al., 'Accelerating ionizable lipid discovery for mRNA delivery using machine learning and combinatorial chemistry'. *Nature Materials*, 23(7): (2024), 1002–1008. https://doi.org/10.1038/s41563-024-01867-3.
19. N. T. Ingolia, S. Ghaemmaghami, J. R. S. Newman and J. S. Weissman, 'Genome-wide analysis in vivo of translation with nucleotide resolution using ribosome profiling'. *Science*, 324(5924): (2009), 218–223. https://doi.org/10.1126/science.1168978.
20. A. Akinc, A. Zumbuehl, M. Goldberg et al., 'A combinatorial library of lipid-like materials for delivery of RNAi therapeutics'. *Nature Biotechnology*, 26(5): (2008), 561–569. https://doi.org/10.1038/nbt1402.
21. J. Jumper, R. Evans, A. Pritzel et al., 'Highly accurate protein structure prediction with AlphaFold'. *Nature*, 596(7873): (2021), 583–589. https://doi.org/10.1038/s41586-021-03819-2.
22. Z. Lin, H. Akin, R. Rao et al., 'Evolutionary-scale prediction of atomic-level protein structure with a language model'. *Science*, 379(6637): (2023), 1123–1130. https://doi.org/10.1126/science.ade2574.
23. E. Nguyen, M. Poli, M. G. Durrant et al., 'Sequence modeling and design from molecular to genome scale with Evo'. *Science*, 386(6723): (2024), eado9336. https://doi.org/10.1126/science.ado9336.
24. B. Burger, P. M. Maffettone, V. V. Gusev et al., 'A mobile robotic chemist'. *Nature*, 583(7815): (2020), 237–241. https://doi.org/10.1038/s41586-020-2442-2.
25. D. A. Boiko, R. MacKnight, B. Kline and G. Gomes, 'Autonomous chemical research with large language models'. *Nature*, 624(7992): (2023), 570–578. https://doi.org/10.1038/s41586-023-06792-0.
26. Y. Xu, H. Cui, K. Pang et al., 'LUMI-lab: a foundation model-driven autonomous platform enabling discovery of new ionizable lipid designs for mRNA delivery'. *Cell*, 189(6): (2026), 1620–1635.e25. https://doi.org/10.1016/j.cell.2026.01.012.
27. J. Gottweis, W. H. Weng, A. Daryin et al., 'Accelerating scientific discovery with co-scientist'. *Nature*, 655(8122): (2026), 487–496. https://doi.org/10.1038/s41586-026-10644-y.
28. A. Madani, B. Krause, E. R. Greene et al., 'Large language models generate functional protein sequences across diverse families'. *Nature Biotechnology*, 41(8): (2023), 1099–1106. https://doi.org/10.1038/s41587-022-01618-2.
29. J. L. Watson, D. Juergens, N. R. Bennett et al., 'De novo design of protein structure and function with RF diffusion'. *Nature*, 620(7976): (2023), 1089–1100. https://doi.org/10.1038/s41586-023-06415-8.

Downloaded from https://www.cambridge.org/core. 17 Nov 2026 at 08:57:33, subject to the Cambridge Core terms of use.

30. K. K. Yang, Z. Wu and F. H. Arnold, 'Machine-learning-guided directed evolution for protein engineering'. *Nature Methods*, 16(8): (2019), 687–694. https://doi.org/10.1038/s41592-019-0496-6.
31. S. Biswas, G. Khimulya, E. C. Alley, K. M. Esvelt and G. M. Church, 'Low-N protein engineering with data-efficient deep learning'. *Nature Methods*, 18 (4): (2021), 389–396. https://doi.org/10.1038/s41592-021-01100-y.
32. Y. LeCun, Y. Bengio and G. Hinton, 'Deep learning'. *Nature*, 521(7553): (2015), 436–444. https://doi.org/10.1038/nature14539.
33. M. Moor, O. Banerjee, Z. S. H. Abad et al., 'Foundation models for generalist medical artificial intelligence'. *Nature*, 616(7956): (2023), 259–265. https://doi.org/10.1038/s41586-023-05881-4.
34. R. Bommasani, D. A. Hudson, E. Adeli, E. and others. On the opportunities and risks of foundation models. arXiv https://doi.org/10.48550/arXiv.2108.07258 (2021).
35. Y. Ji, Z. Zhou, H. Liu and R. V. Davuluri, 'DNABERT: pre-trained bidirectional encoder representations from transformers model for DNA-language in genome'. *Bioinformatics*, 37(15): (2021), 2112–2120. https://doi.org/10.1093/bioinformatics/btab083.
36. H. Dalla-Torre, L. Gonzalez, J. Mendoza-Revilla et al., 'Nucleotide transformer: building and evaluating robust foundation models for human genomics'. *Nature Methods*, 22(2): (2025), 287–297. https://doi.org/10.1038/s41592-024-02523-z.
37. Y. Yang, G. Li, K. Pang et al., 'Deciphering 3'UTR mediated gene regulation using interpretable deep representation learning'. *Advanced Science*, 11(39): (2024), e2407013. https://doi.org/10.1002/advs.202407013.
38. Y. Chu, D. Yu, Y. Li et al., 'A 5′ UTR language model for decoding untranslated regions of mRNA and function predictions'. *Nature Machine Intelligence*, 6(4): (2024), 449–460. https://doi.org/10.1038/s42256-024-00823-9.
39. S. Li, S. Moayedpour, R. Li et al., 'CodonBERT large language model for mRNA vaccines'. *Genome Research*, 34(7): (2024), 1027–1035. https://doi.org/10.1101/gr.278870.123.
40. Zhou, G. Uni-mol: a universal 3D molecular representation learning framework. Preprint at ChemRxiv (2023).
41. J. B. Ingraham, M. Baranov, Z. Costello et al., 'Illuminating protein space with a programmable generative model'. *Nature*, 623(7989): (2023), 1070–1078. https://doi.org/10.1038/s41586-023-06728-8.
42. P. J. Sample, B. Wang, D. W. Reid et al., 'Human 5′ UTR design and variant effect prediction from a massively parallel translation assay'. *Nature Biotechnology*, 37(7): (2019), 803–809. https://doi.org/10.1038/s41587-019-0164-5.
43. H. Zhang, L. Zhang, A. Lin et al., 'Algorithm for optimized mRNA design improves stability and immunogenicity'. *Nature*, 621(7978): (2023), 396–403. https://doi.org/10.1038/s41586-023-06127-z.
44. N. M. Angenent-Mari, A. S. Garruss, L. R. Soenksen, G. Church and J. J. Collins, 'A deep learning approach to programmable RNA switches'. *Nature Communications*, 11(1): (2020), 5057. https://doi.org/10.1038/s41467-020-18677-1.
45. J. A. Valeri, K. M. Collins, P. Ramesh et al., 'Sequence-to-function deep learning frameworks for engineered riboregulators'. *Nature Communications*, 11(1): (2020), 5058. https://doi.org/10.1038/s41467-020-18676-2.
46. R. P. Patwardhan, J. B. Hiatt, D. M. Witten et al., 'Massively parallel functional dissection of mammalian enhancers in vivo'. *Nature Biotechnology*, 30(3): (2012), 265–270. https://doi.org/10.1038/nbt.2136.
47. H. K. Wayment-Steele, D. S. Kim, C. A. Choe et al., 'Theoretical basis for stabilizing messenger RNA through secondary structure design'. *Nucleic Acids Research*, 49(18): (2021), 10604–10617. https://doi.org/10.1093/nar/gkab764.
48. K. Karikó, M. Buckstein, H. Ni and D. Weissman, 'Suppression of RNA recognition by toll-like receptors: the impact of nucleoside modification and the evolutionary origin of RNA'. *Immunity*, 23(2): (2005), 165–175. https://doi.org/10.1016/j.immuni.2005.06.008.
49. T. E. Mulroney, T. Pöyry, J. C. Yam-Puc et al., 'N1-methylpseudouridylation of mRNA causes +1 ribosomal frameshifting'. *Nature*, 625(7993): (2024), 189–194. https://doi.org/10.1038/s41586-023-06800-3.
50. D. H. Turner and D. H. Mathews, 'NNDB: the nearest neighbor parameter database for predicting stability of nucleic acid secondary structure'. *Nucleic Acids Research*, 38(suppl_1): (2010), D280–D282. https://doi.org/10.1093/nar/gkp892.
51. R. Lorenz, S. H. Bernhart, C. Höner zu Siederdissen et al., 'ViennaRNA package 2.0'. *Algorithms for Molecular Biology*, 6(1): (2011), 26. https://doi.org/10.1186/1748-7188-6-26.
52. N. A. Siegfried, S. Busan, G. M. Rice, J. A. E. Nelson and K. M. Weeks, 'RNA motif discovery by SHAPE and mutational profiling (SHAPE-MaP)'. *Nature Methods*, 11(9): (2014), 959–965. https://doi.org/10.1038/nmeth.3029.
53. J. Nelson, E. W. Sorensen, S. Mintri et al., 'Impact of mRNA chemistry and manufacturing process on innate immune activation'. *Science Advances*, 6 (26): (2020), eaaz6893. https://doi.org/10.1126/sciadv.aaz6893.
54. Y. Xu, S. Ma, H. Cui et al., 'AGILE platform: a deep learning powered approach to accelerate LNP development for mRNA delivery'. *Nature Communications*, 15(1): (2024), 6305. https://doi.org/10.1038/s41467-024-50619-z.
55. J. C. Fromer and C. W. Coley, 'An algorithmic framework for synthetic cost-aware decision making in molecular design'. *Nature Computational Science*, 4(6): (2024), 440–450. https://doi.org/10.1038/s43588-024-00639-y.
56. Z. Tu, S. J. Choure, M. H. Fong et al., 'ASKCOS: open-source, data-driven synthesis planning'. *Accounts of Chemical Research*, 58(11): (2025), 1764–1775. https://doi.org/10.1021/acs.accounts.5c00155.
57. J. F. Joung, M. H. Fong, N. Casetti et al., 'Electron flow matching for generative reaction mechanism prediction'. *Nature*, 645(8079): (2025), 115–123. https://doi.org/10.1038/s41586-025-09426-9.
58. M. H. S. Segler, M. Preuss and M. P. Waller, 'Planning chemical syntheses with deep neural networks and symbolic AI'. *Nature*, 555(7698): (2018), 604–610. https://doi.org/10.1038/nature25978.
59. C. W. Coley, III D. A. Thomas, J. A. M. Lummiss et al., A robotic platform for flow synthesis of organic compounds informed by AI planning. *Science*, 3656453: (2019), eaax1566. https://doi.org/10.1126/science.aax1566.
60. C. W. Coley, R. Barzilay, T. S. Jaakkola, W. H. Green and K. F. Jensen, 'Prediction of organic reaction outcomes using machine learning'. *ACS Central Science*, 3(5): (2017), 434–443. https://doi.org/10.1021/acscentsci.7b00064.
61. H. Gao, T. J. Struble, C. W. Coley et al., 'Using machine learning to predict suitable conditions for organic reactions'. *ACS Central Science*, 4(11): (2018), 1465–1476. https://doi.org/10.1021/acscentsci.8b00357.
62. B. J. Shields, J. Stevens, J. Li et al., 'Bayesian reaction optimization as a tool for chemical synthesis'. *Nature*, 590(7844): (2021), 89–96. https://doi.org/10.1038/s41586-021-03213-y.
63. G. Wu, T. Jin, A. Alexander-Katz and C. W. Coley, 'Autonomous discovery of functional random heteropolymer blends through evolutionary formulation optimization'. *Matter*, 8(12): (2025), 102336. https://doi.org/10.1016/j.matt.2025.102336.
64. J. Witten, I. Raji, R. S. Manan et al., 'Artificial intelligence-guided design of lipid nanoparticles for pulmonary gene therapy'. *Nature Biotechnology*, 43 (11): (2025), 1790–1799. https://doi.org/10.1038/s41587-024-02490-y.
65. P. R. Cullis and P. L. Felgner, 'The 60-year evolution of lipid nanoparticles for nucleic acid delivery'. *Nature Reviews Drug Discovery*, 23(9): (2024), 709–722. https://doi.org/10.1038/s41573-024-00977-6.
66. M. H. Y. Cheng, J. Leung, Y. Zhang et al., 'Induction of bleb structures in lipid nanoparticle formulations of mRNA leads to improved transfection potency'. *Advanced Materials*, 35(31): (2023), e2303370. https://doi.org/10.1002/adma.202303370.
67. D. M. Strelkova Petersen, N. Chaudhary, M. L. Arral, R. M. Weiss and K. A. Whitehead, 'The mixing method used to formulate lipid nanoparticles affects mRNA delivery efficacy and organ tropism'. *European Journal of Pharmaceutics and Biopharmaceutics*, 192: (2023), 126–135. https://doi.org/10.1016/j.ejpb.2023.10.006.
68. M. Packer, D. Gyawali, R. Yerabolu, J. Schariter and P. White, 'A novel mechanism for the loss of mRNA activity in lipid nanoparticle delivery systems'. *Nature Communications*, 12(1): (2021), 6777. https://doi.org/10.1038/s41467-021-26926-0.
69. R. E. Birdsall, D. Han, K. DeLaney et al., 'Monitoring stability indicating impurities and aldehyde content in lipid nanoparticle raw material and

Downloaded from https://www.cambridge.org/core. 17 Nov 2026 at 08:57:33, subject to the Cambridge Core terms of use.

formulated drugs'. *Journal of Chromatography B*, 1234: (2024), 124005. https://doi.org/10.1016/j.jchromb.2024.124005.
70. T. Bethiana, A. Aljabbari, Y. Li et al., 'Effects of different mixing techniques on mRNA lipid nanoparticle physicochemistry and biological performance'. *Nature Communications*, 17(1): (2026), 5890. https://doi.org/10.1038/s41467-026-72499-1.
71. M. P. Monopoli, C. Åberg, A. Salvati and K. A. Dawson, 'Biomolecular coronas provide the biological identity of nanosized materials'. *Nature Nanotech*, 7(12): (2012), 779–786. https://doi.org/10.1038/nnano.2012.207.
72. A. Akinc, W. Querbes, S. de et al., 'Targeted delivery of RNAi therapeutics with endogenous and exogenous ligand-based mechanisms'. *Molecular Therapy*, 18(7): (2010), 1357–1364. https://doi.org/10.1038/mt.2010.85.
73. S. A. Dilliard, Q. Cheng and D. J. Siegwart, 'On the mechanism of tissue-specific mRNA delivery by selective organ targeting nanoparticles'. *Proceedings of the National Academy of Sciences*, 118(52): (2021), e2109256118. https://doi.org/10.1073/pnas.2109256118.
74. J. Gilleron, W. Querbes, A. Zeigerer et al., 'Image-based analysis of lipid nanoparticle–mediated siRNA delivery, intracellular trafficking and endosomal escape'. *Nature Biotechnology*, 31(7): (2013), 638–646. https://doi.org/10.1038/nbt.2612.
75. G. Sahay, W. Querbes, C. Alabi et al., 'Efficiency of siRNA delivery by lipid nanoparticles is limited by endocytic recycling'. *Nature Biotechnology*, 31 (7): (2013), 653–658. https://doi.org/10.1038/nbt.2614.
76. K. Paunovska, C. D. Sago, C. M. Monaco et al., 'A direct comparison of in vitro and in vivo nucleic acid delivery mediated by hundreds of nanoparticles reveals a weak correlation'. *Nano Letters*, 18(3): (2018), 2148–2157. https://doi.org/10.1021/acs.nanolett.8b00432.
77. K. Paunovska, D. Loughrey and J. E. Dahlman, 'Drug delivery systems for RNA therapeutics'. *Nature Reviews Genetics*, 23(5): (2022), 265–280. https://doi.org/10.1038/s41576-021-00439-4.
78. A. Raguram, S. Banskota and D. R. Liu, 'Therapeutic in vivo delivery of gene editing agents'. *Cell*, 185(15): (2022), 2806–2827. https://doi.org/10.1016/j.cell.2022.03.045.
79. ENCODE Project Consortium, 'An integrated encyclopedia of DNA elements in the human genome'. *Nature*, 489(7414): (2012), 57–74. https://doi.org/10.1038/nature11247.
80. The RNAcentral Consortium, 'RNAcentral 2021: secondary structure integration, improved sequence search and new member databases'. *Nucleic Acids Research*, 49(D1): (2021), D212–D220. https://doi.org/10.1093/nar/gkaa921.
81. I. Kalvari, E. P. Nawrocki, N. Ontiveros-Palacios and others, 'Rfam 14: expanded coverage of metagenomic, viral and microRNA families'. *Nucleic Acids Research*, 49(D1): (2021), D192–D200. https://doi.org/10.1093/nar/gkaa1047.
82. K. Sato, M. Akiyama and Y. Sakakibara, 'RNA secondary structure prediction using deep learning with thermodynamic integration'. *Nature Communications*, 12(1): (2021), 941. https://doi.org/10.1038/s41467-021-21194-4.
83. S. N. Bhatia, X. Chen, M. A. Dobrovolskaia and T. Lammers, 'Cancer nanomedicine'. *Nature Reviews Cancer*, 22(10): (2022), 550–556. https://doi.org/10.1038/s41568-022-00496-9.
84. M. Jayaraman, S. M. Ansell, B. L. Mui et al., 'Maximizing the potency of siRNA lipid nanoparticles for hepatic gene silencing in vivo'. *Angewandte Chemie International Edition*, 51(34): (2012), 8529–8533. https://doi.org/10.1002/anie.201203263.
85. K. J. Hassett, K. E. Benenato, E. Jacquinet et al., 'Optimization of lipid nanoparticles for intramuscular administration of mRNA vaccines'. *Molecular Therapy Nucleic Acids*, 15: (2019), 1–11. https://doi.org/10.1016/j.omtn.2019.01.013.
86. I. M. Hafez, N. Maurer and P. R. Cullis, 'On the mechanism whereby cationic lipids promote intracellular delivery of polynucleic acids'. *Gene Therapy*, 8(15): (2001), 1188–1196. https://doi.org/10.1038/sj.gt.3301506.
87. P. R. Cullis and M. J. Hope, 'Lipid nanoparticle systems for enabling gene therapies'. *Molecular Therapy*, 25(7): (2017), 1467–1475. https://doi.org/10.1016/j.ymthe.2017.03.013.
88. K. A. Whitehead, J. R. Dorkin, A. J. Vegas et al., 'Degradable lipid nanoparticles with predictable in vivo siRNA delivery activity'. *Nature Communications*, 5 (1): (2014), 4277. https://doi.org/10.1038/ncomms5277.
89. K. A. Hajj, R. L. Ball, S. B. Deluty et al., 'Branched-tail lipid nanoparticles potently deliver mRNA in vivo due to enhanced ionization at endosomal pH'. *Small*, 15(6): (2019), e1805097. https://doi.org/10.1002/smll.201805097.
90. D. H. Brookes, H. Park and J. Listgarten. Conditioning by adaptive sampling for robust design. Preprint at https://doi.org/10.48550/arXiv.1901.10060 (2021).
91. D. Truhn, J. S. Reis-Filho and J. N. Kather, 'Large language models should be used as scientific reasoning engines, not knowledge databases'. *Nature Medicine*, 29(12): (2023), 2983–2984. https://doi.org/10.1038/s41591-023-02594-z.
92. A. J. Thirunavukarasu, D. S. J. Ting, K. Elangovan et al., 'Large language models in medicine'. *Nature Medicine*, 29(8): (2023), 1930–1940. https://doi.org/10.1038/s41591-023-02448-8.
93. D. Van Veen, 'Adapted large language models can outperform medical experts in clinical text summarization'. *Nature Medicine*, 30(4): (2024), 1134–1142. https://doi.org/10.1038/s41591-024-02855-5.
94. R. Bommasani, S. Kapoor, K. Klyman et al., 'Considerations for governing open foundation models'. *Science*, 386(6718): (2024), 151–153. https://doi.org/10.1126/science.adp1848.
95. F. Ren, A. Aliper, J. Chen et al., 'A small-molecule TNIK inhibitor targets fibrosis in preclinical and clinical models'. *Nature Biotechnology*, 43(1): (2025), 63–75. https://doi.org/10.1038/s41587-024-02143-0.
96. U. Sahin, E. Derhovanessian, M. Miller et al., 'Personalized RNA mutanome vaccines mobilize poly-specific therapeutic immunity against cancer'. *Nature*, 547(7662): (2017), 222–226. https://doi.org/10.1038/nature23003.
97. L. A. Rojas, Z. Sethna, K. C. Soares et al., 'Personalized RNA neoantigen vaccines stimulate T cells in pancreatic cancer'. *Nature*, 618(7963): (2023), 144–150. https://doi.org/10.1038/s41586-023-06063-y.
98. M. T. Villanueva, 'Double boost for mRNA cancer vaccines'. *Nature Reviews Drug Discovery*, 18(12): (2019), 902–902. https://doi.org/10.1038/d41573-019-00192-8.
99. J. Kim, C. Hu, C. Moufawad el Achkar et al., 'Patient-customized oligonucleotide therapy for a rare genetic disease'. *New England Journal of Medicine*, 381(17): (2019), 1644–1652. https://doi.org/10.1056/NEJMoa1813279.
100. First self-amplifying mRNA vaccine approved'. *Nature Biotechnology*, 42 (1): (2024), 4–4. https://doi.org/10.1038/s41587-023-02101-2.
101. M. Saville, J. P. Cramer, M. Downham et al., 'Delivering pandemic vaccines in 100 days — what will it take?'. *New England Journal of Medicine*, 387(2): (2022), e3. https://doi.org/10.1056/NEJMp2202669.
102. M. Youssef, C. Hitti, J. Puppin Chaves Fulber and A. A. Kamen, 'Enabling mRNA therapeutics: current landscape and challenges in manufacturing'. *Biomolecules*, 13(10): (2023), 1497. https://doi.org/10.3390/biom13101497.

Downloaded from https://www.cambridge.org/core. 17 Nov 2026 at 08:57:33, subject to the Cambridge Core terms of use.