OPEN ACCESS

Biomedical Engineering REVIEW

The Innovation Drug Discovery

Artificial intelligence in the rational design of lipid nanoparticles for mRNA therapeutics

Heyu Zhao,[1,5] Junchao Xu,[2,5] Xia Gao,[3] Jingcheng Zhu,[2] and Zhongshan He[4,*]

*Correspondence: udr5kr@virginia.edu

Received: January 12, 2026; Accepted: April 6, 2026; Published Online: April 8, 2026; https://doi.org/10.59717/j.xinn-drugdisc.2026.100006

GRAPHICAL ABSTRACT

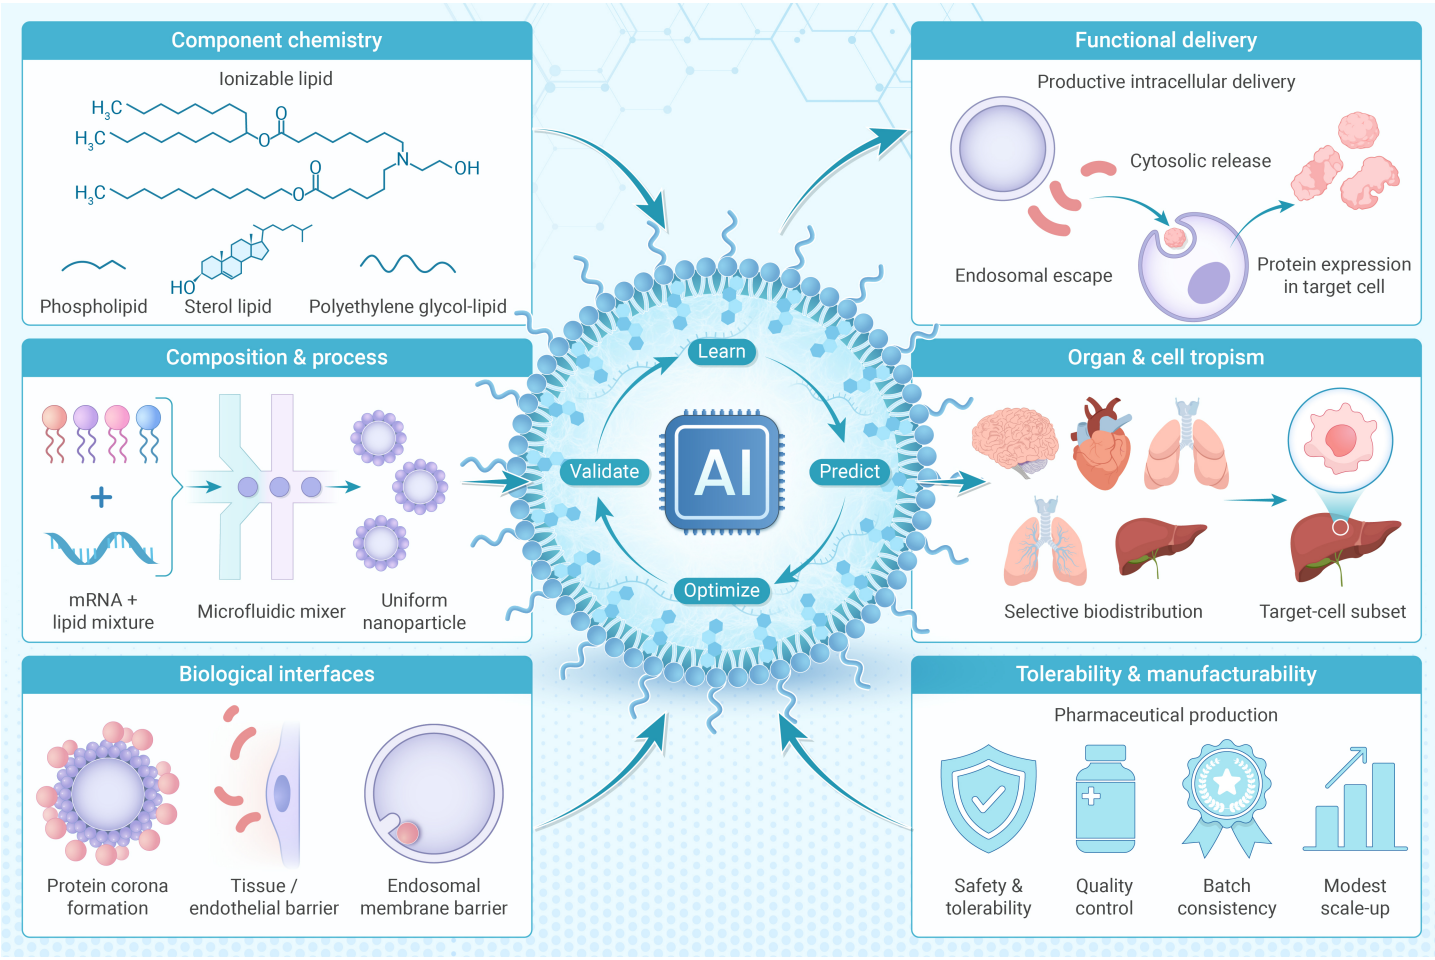

PUBLIC SUMMARY

- Lipid nanoparticles (LNPs) are not passive carriers; they shape where messenger RNA (mRNA) reaches.
- This review explains how lipid chemistry, formulation, and manufacturing jointly control delivery.
- Artificial intelligence (AI) can learn these linked factors and accelerate safer, more targeted LNP design.
- Progress now depends on shared data, better assays, and models built for real-world manufacturing.

REVIEW Biomedical Engineering

Artificial intelligence in the rational design of lipid nanoparticles for mRNA therapeutics

Heyu Zhao,[1,5] Junchao Xu,[2,5] Xia Gao,[3] Jingcheng Zhu,[2] and Zhongshan He[4,*]

[1]Institutes of Biomedical Sciences, Fudan University, Shanghai 200032, China
[2]Department of Bioengineering, University of Pennsylvania, Philadelphia, PA 19104, USA
[3]Department of Experimental Surgery, Tangdu Hospital, The Fourth Military Medical University, Shaanxi 710038, China
[4]Departments of Molecular Physiology and Biological Physics and Biomedical Engineering, University of Virginia, Charlottesville, VA 22903, USA
[5]These authors contributed equally
*Correspondence: udr5kr@virginia.edu

Received: January 12, 2026; Accepted: April 6, 2026; Published Online: April 8, 2026; https://doi.org/10.59717/j.xinn-drugdisc.2026.100006

Citation: Zhao H., Xu J., Gao X., et al. (2026). Artificial intelligence in the rational design of lipid nanoparticles for mRNA therapeutics. *The Innovation Drug Discovery* **1**:100006.

Lipid nanoparticles (LNPs) have made RNA therapeutics clinically viable, yet delivery remains the dominant constraint on efficacy and safety as applications expand beyond the liver. Performance is emergent: it reflects coupled choices in ionizable-lipid chemistry, multi-component formulation, and process history, and is further reshaped by biological interfaces including protein corona remodeling, tissue transport barriers, endocytic trafficking, and low-probability endosomal escape. This review surveys how artificial intelligence is converting this nonlinear design space from empirical iteration into data-efficient, multi-objective optimization. We highlight (i) structure-activity learning and synthesis-aware generative modeling for ionizable lipid discovery; (ii) formulation- and process-conditioned architectures that treat LNPs as composite, process-defined materials; and (iii) pooled in vivo barcoding and single-cell readouts that enable prediction and tuning of organ and cell-type tropism. We conclude by outlining three priorities for translation: standardized data and metadata, mechanistic endpoints for escape and immunogenicity, and CMC-aware optimization. Progress on these fronts will be necessary for potent, safe, and manufacturable mRNA-LNP medicines.

INTRODUCTION

Messenger RNA (mRNA) therapeutics have progressed from a delivery-limited concept to a scalable clinical modality, enabled by advances in RNA chemistry, manufacturing, and non-viral delivery systems. Nucleoside modifications that dampen innate immune sensing while preserving translational output were a pivotal breakthrough.[1–2] Decades of delivery research then converged on ionizable lipid nanoparticles (LNPs) that support systemic RNA delivery.[3–4] Patisiran established clinical feasibility for siRNA-LNP delivery,[5–6] and COVID-19 mRNA-LNP vaccines demonstrated robust efficacy in large, global phase 3 trials.[7–8] Together, these milestones established mRNA-LNP as a clinically tractable platform and accelerated investment in LNP-enabled RNA therapeutics and delivery technologies.[9–10]

Despite these milestones, delivery remains the dominant determinant of therapeutic index for many RNA modalities. Clinically validated LNPs are multi-component, process-defined assemblies: performance depends on coupled choices of ionizable lipids, phospholipids, sterols, and PEG-lipids, and on process history that shapes size distributions, internal morphology, and surface presentation.[11–15] As programs move beyond hepatic delivery into oncology, immunotherapy, protein replacement, and in vivo genome editing, requirements tighten simultaneously across potency, tissue selectivity, tolerability, and repeat-dose compatibility.[12,16–18] These pressures have motivated organ-selective LNP chemistries such as selective organ targeting (SORT).[19] Recent amidine-incorporated degradable lipids further illustrate how modular chemistry can broaden in vivo mRNA delivery performance.[20]

These requirements collide with mechanistic bottlenecks that remain only partially resolved and are difficult to measure at scale. In circulation, LNPs rapidly acquire a biomolecular corona that can redirect biodistribution and engage innate immune pathways. At the tissue level, vascular permeability and extracellular matrices impose cell-type-specific transport constraints. At the cellular level, productive cytosolic delivery is typically limited by endosomal escape and subsequent trafficking.[11–12,17,21–27] Because the delivery cascade is nonlinear and context dependent, modest changes in composition or process history can yield outsized shifts in potency, organ tropism, and tolerability.[11–15]

Traditional optimization strategies struggle in this regime. Even within a constrained synthetic scheme, ionizable lipid libraries expand combinatorially as headgroup pKa, linker chemistry, hydrophobic tail geometry, and biodegradability are varied.[4,28–30] The design space grows further when helper-lipid identity, molar ratios, buffer conditions, mixing rates, and post-assembly processing are included.[13–15] Development constraints such as manufacturability, scale-up feasibility, and stability requirements couple discovery decisions to downstream CMC realities.[5,12,17] Design-of-experiments approaches (for example, definitive screening designs) can map local structure-process-property relationships,[31] but they still scale poorly when variables interact and biological readouts are context dependent.

These constraints explain why one-factor-at-a-time iteration and local design-of-experiments often stall in next-generation mRNA-LNP programs. The ionizable-lipid space is combinatorial, formulation composition and process history are tightly coupled, in vivo endpoints are expensive and noisy, and optimization must balance potency, tissue selectivity, tolerability, repeat-dose compatibility, stability, and manufacturability. AI-based workflows are useful in this setting because they can screen broader spaces in silico, allocate experiments with uncertainty-aware active learning, and optimize several constraints at once.

At the same time, the experimental landscape is changing in ways that make data-centric approaches feasible. High-throughput synthesis, microfluidic formulation, and pooled or barcoded in vivo screens can generate dense datasets spanning physicochemical properties, biodistribution, gene expression, and safety-relevant endpoints.[15,32–33] Yet these datasets routinely break naive models: missing annotations for process variables, batch effects across reagents and animals, label noise from biological variability, and distribution shifts when moving between species, disease states, and manufacturing scales.

AI and machine learning can learn coupled, high-dimensional relationships, propose candidates under explicit constraints, and prioritize experiments when data are expensive. In practice, value depends on more than predictive accuracy: models must capture chemically meaningful structure-property relationships, quantify uncertainty to avoid overconfident extrapolation, and incorporate mechanistic priors when direct supervision is sparse. Hybrid strategies that combine learned representations with mechanistic readouts (for example, endosomal-escape proxies), physics- or kinetics-inspired features, and decision-making frameworks such as Bayesian optimization and active learning are therefore becoming standard for allocating experimental budgets efficiently.

Across these layers, AI contributes in four practical ways: it learns nonlinear structure-process-property relationships, directs scarce experiments toward informative regions of the design space, optimizes multiple objectives under explicit synthetic and CMC constraints, and transfers information from larger chemical or assay datasets when delivery labels are sparse (Figure 1). Figure 1C also provides a short timeline of selected milestones, from barcoded in vivo screening and SORT organ targeting to recent AI-guided lipid and formulation platforms.

AI ARCHITECTURES FOR LIPID OPTIMIZATION

AI-augmented LNP optimization departs fundamentally from traditional,

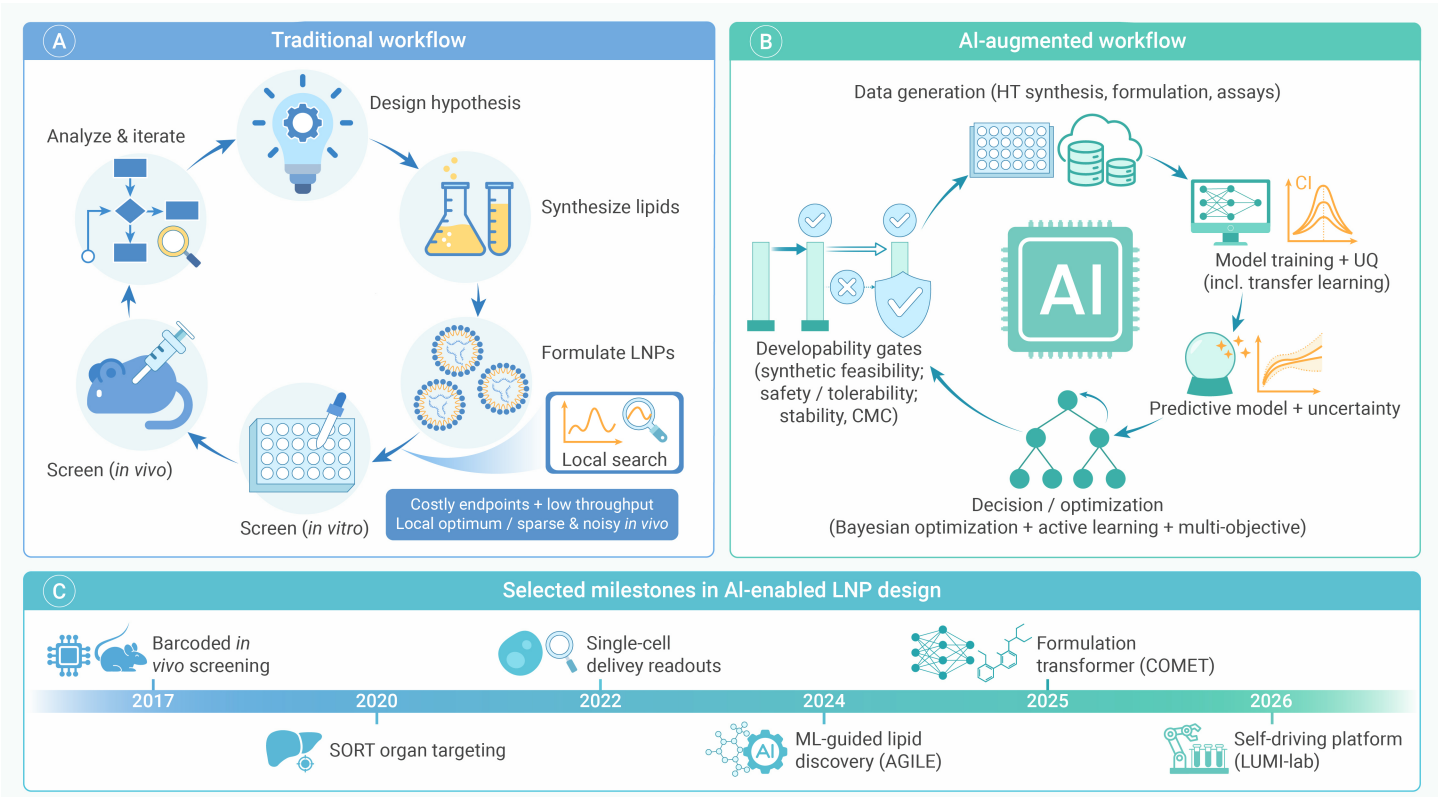

Figure 1. From empirical iteration to AI-augmented LNP engineering (A) Traditional LNP development relies on sequential cycles of design, synthesis/formulation, in vitro testing, in vivo testing, and iteration, which scale poorly in high-dimensional, multi-objective design spaces. (B) AI-augmented workflows integrate high-throughput data generation with model training, uncertainty quantification (UQ), predictive modeling, developability assessment, and Bayesian/active-learning loops to support data-efficient discovery under synthetic-feasibility, safety/tolerability, stability, and chemistry, manufacturing, and controls (CMC) constraints. (C) Selected milestones in AI-enabled LNP design, including barcoded in vivo screening, SORT organ targeting, single-cell delivery readouts, and representative AI platforms.

low-throughput design-build-test cycles by integrating high-throughput experimentation with virtual screening, uncertainty-aware active learning, and multi-objective optimization under explicit synthetic-feasibility and CMC constraints (Figure 1A–B). Selected milestones underpinning this transition are summarized in Figure 1C.

At the molecular scale, the goal of AI is not simply to rank candidates within fixed libraries, but to infer the structural determinants of key delivery bottlenecks, including RNA–lipid complex formation during formulation, productive endosomal escape, and the balance among biodegradation, tolerability, and potency across biological contexts. Major model families and their practical trade-offs are summarized in Figure 2 and Table 1.

Within this design space, most AI-enabled LNP studies rely on a relatively limited set of model classes. Conventional supervised models remain useful in chemically narrow, data-limited settings. Gaussian-process surrogates provide calibrated uncertainty for Bayesian optimization. Graph neural networks learn directly from lipid structure, transformers capture interactions among tokenized components or sequences, and constrained generative models propose candidate lipids under synthetic and safety constraints. More recently, multimodal architectures and foundation models pretrained on large molecular corpora have been adapted to sparse delivery datasets.

Model selection in LNP design is therefore governed by trade-offs among interpretability, data requirements, and robustness to scaffold or process shift. Classical models can perform well on small, well-controlled datasets and provide transparent feature attribution, but they often degrade under distribution shift. Deep architectures can generalize across broader chemical and formulation spaces, particularly when paired with self-supervised pretraining and uncertainty quantification, although their value still depends on rigorous validation and transparent reporting of experimental covariates. Transfer learning through molecular pretraining followed by delivery-specific fine-tuning is thus a practical strategy for expanding effective training data without requiring prohibitively large labeled LNP datasets.

From empirical rules to mechanistically aware representations

Foundational siRNA-era work established that delivery depends on a narrow physicochemical window.[3–4] Ionizable lipids should be largely neutral at physiological pH (pH 7.4) to reduce systemic toxicity and complement activation, yet protonate efficiently in the acidifying endosome (pH 5.5–6.5) to promote membrane destabilization through hexagonal phase transitions and cargo release. Systematic structure–activity studies demonstrated that effective pKa values of 6.2–6.8 maximize the balance between circulation stability and endosomal membrane disruption.[4] Quantitative imaging added a critical constraint: the majority of internalized cargo remains trapped in endolysosomal compartments, with productive escape occurring through transient, stochastic membrane contact events.[22] This mechanistic complexity motivates representation learning approaches that can capture non-linear structure–function relationships beyond simple pKa correlations.

Discriminative learning with graph neural networks

Data-driven efforts often start with curated physicochemical descriptors (e.g., logP, polar surface area, rotatable bonds) coupled to linear or tree-based regressors. Such models can be informative within a chemical series but are sensitive to descriptor choice and can degrade under scaffold shift.[40–41] Graph neural networks (GNNs) learn representations directly from molecular graphs and, particularly when combined with self-supervised pretraining, can improve extrapolation beyond the training distribution.[35,41–42]

The AGILE (AI-Guided Ionizable Lipid Engineering) platform represents a seminal application of this paradigm.[34] AGILE employs a Graph Isomorphism Network (GIN) architecture (which updates node embeddings through neighborhood aggregation while preserving graph-level discriminative power) enhanced through MolCLR contrastive pre-training on large molecular databases (millions of structures).[42] This self-supervised pretraining learns generalizable chemical representations before task-specific fine-tuning on delivery data. The platform was trained on 1,200 synthesized ionizable lipids with in vitro transfection readouts, then applied to virtually screen 12,000

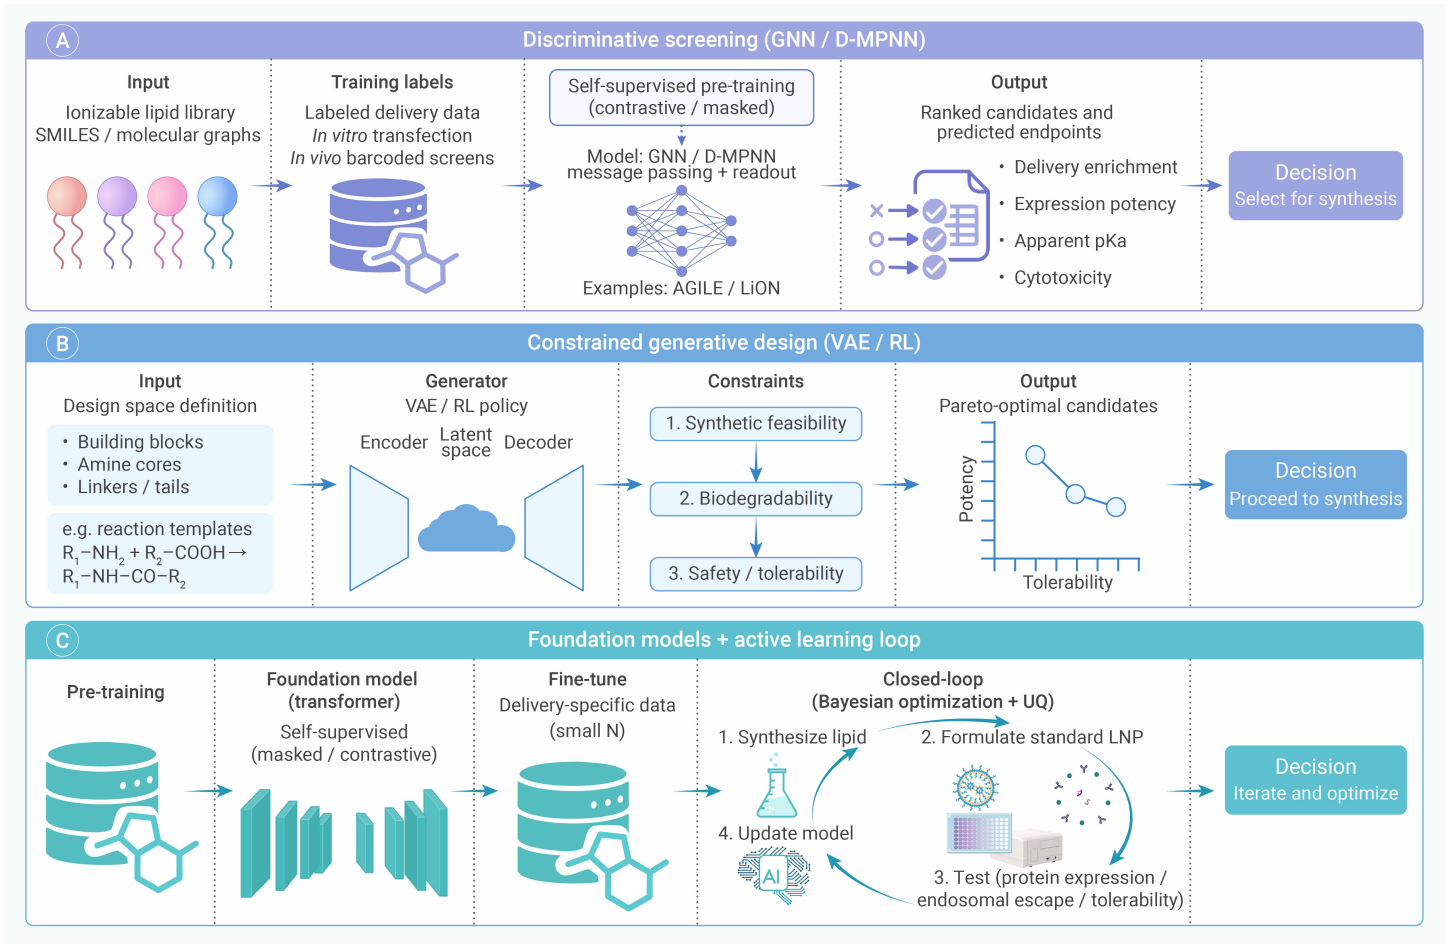

Figure 2. AI workflows for ionizable lipid discovery and optimization (A) Discriminative screening models, including graph neural networks (GNNs) and directed message passing neural networks (D-MPNNs), learn structure-activity relationships from labeled in vitro and/or in vivo datasets to rank candidates and predict delivery-related properties. Optional self-supervised pretraining can improve data efficiency and generalization. (B) Constrained generative design, including variational autoencoders (VAEs) and reinforcement learning (RL), proposes new lipids within defined chemical and reaction spaces while optimizing synthetic feasibility, biodegradability, potency, and safety/tolerability. (C) Foundation-model workflows combine self-supervised pretraining, task-specific fine-tuning, and closed-loop Bayesian optimization with UQ to guide lipid synthesis, formulation, testing, and iterative model updating.

Table 1. Representative AI platforms for LNP optimization across molecular and formulation–process scales

Platform	Institution	Scale / Architecture	Training data	Key result
AGILE[34]	University of Toronto	Molecular / GIN + MolCLR	1,200 lipids	Designed ionizable lipids achieving up to ~7.8× higher muscle expression vs MC3; revealed cell-type-specific lipid preferences.
LiON[35]	MIT	Molecular → organ / D-MPNN	>9,000 LNPs	Virtually screened ~1.6M lipids; identified FO-32 and FO-35. FO-32 matched state-of-the-art nebulized mRNA delivery to mouse lung; FO-32 and FO-35 delivered mRNA to ferret lungs; FO-35 also supported local delivery to mouse muscle and nasal mucosa.
COMET[36]	MIT/NTU	Formulation / Transformer	>3,000 LNPs	Transferred across tasks with strong agreement; outperformed baselines in predicting delivery-relevant properties.
LUMI-lab[37]	University of Toronto	Multi-scale / Foundation + SDL	Large-scale molecular pretraining	Autonomous closed loop tested >1,700 LNPs; identified brominated tails and a lead inhalable formulation with 20.3% lung epithelial gene editing in mice.
TransMA[38]	Chinese Academy of Sciences	Molecular / explainable multi-modal model	Public benchmarks	Transformer-based lipid representation improves molecular property prediction on public benchmarks; enables transferable descriptors for downstream models.
TuNa-AI[39]	Duke	Formulation / GP + automation	Iterative	Hybrid-kernel design improves nanoparticle formation success and supports tunable excipient selection for better developability.

Abbreviations: D-MPNN, directed message passing neural network; GIN, graph isomorphism network; GP, Gaussian process; MolCLR, molecular contrastive learning; SDL, self-driving laboratory.

candidates. Prospective validation reported lipid H9 achieving 7.8-fold greater muscle-specific delivery compared with DLin-MC3-DMA (MC3), and lipid R6 demonstrating 5-fold enhanced macrophage transfection; structurally distinct scaffolds that would be unlikely to emerge from intuition-driven

Table 2. Representative AI-guided LNP component designs linked to tissue/cell targeting outcomes

Target tissue/cell	Design lever	AI approach	Training/screening data	Key finding/outcome	Ref
Muscle / macrophage	Ionizable lipid structure	GIN + MolCLR (AGILE)	~1,200 lipids; in vivo and cell-type readouts	H9 increased muscle expression (~7.8× vs MC3) in mice; R6 increased macrophage transfection (~5× in RAW 264.7 cells)	34
Lung (systemic + airway)	Ionizable lipid structure	D-MPNN virtual screening (LiON)	~1.6M virtual lipids screened; >9,000 LNPs for model training	Identified FO-32 and FO-35; FO-32 matched state-of-the-art nebulized mRNA delivery to mouse lung; FO-32 and FO-35 delivered mRNA to ferret lungs; FO-35 also supported local delivery to mouse muscle and nasal mucosa.	35
Lung (inhaled)	Ionizable lipid tails	Foundation model + SDL closed loop (LUMI-lab)	Large-scale molecular pretraining; >1,700 autonomous tests	Brominated tails; lead inhalable LNP reported 20.3% lung epithelial gene editing in mice	37
Brain / neurons	Surface chemistry (ligand/PEG-lipid)	AI-assisted/validated targeting	In vivo brain expression readouts	AI-validated brain-targeted mRNA-LNPs with neuronal tropism via BBB-interacting small-molecule functionalization	48
CNS immune cells (microglia)	Formulation + immunomodulatory design	ML-assisted multi-objective optimization	Phenotype + transfection readouts	Designed immunomodulatory LNPs delivering IL10 mRNA to modulate microglial activation state in vitro; validated in human iPSC-derived microglia	49

Abbreviations: BBB, blood–brain barrier; D-MPNN, directed message passing neural network; GIN, graph isomorphism network; SDL, self-driving laboratory.

exploration.

The LiON (Lipid Optimization through Neural networks) platform extends the GNN paradigm using directed message-passing neural networks (D-MPNN) that incorporate bond-level messages and directional information flow.[35] Unlike standard GNNs that pass messages only between nodes, D-MPNNs propagate information along directed edges, capturing asymmetric bond environments relevant for distinguishing constitutional isomers and stereochemical contexts that can impact biological activity. LiON was trained on >9,000 LNP measurements spanning multiple assay conditions, enabling virtual screening of 1.6 million ionizable lipid candidates. In a pulmonary-focused study, LiON identified candidates FO-32 and FO-35. In vivo, FO-32 matched state-of-the-art performance for nebulized mRNA delivery to mouse lung, and both FO-32 and FO-35 delivered mRNA to ferret lungs; FO-35 also supported local mRNA delivery to mouse muscle and nasal mucosa.

Representative examples of AI-guided ionizable lipid designs enabling tissue- and cell-type-selective delivery are summarized in Table 2. Recent studies from 2024 to 2026 further expand this molecular-design layer. Li, B. et al. combined machine learning with combinatorial chemistry to accelerate ionizable-lipid discovery for mRNA delivery.[43] Han, X. et al. developed an in situ combinatorial synthesis strategy for degradable branched lipidoids, showing how branching and degradable linkers can expand accessible chemical space while improving systemic mRNA delivery.[44] Han, X. et al. also reported a plug-and-play assembly strategy for biodegradable ionizable lipids that supports potent mRNA delivery and in vivo gene editing.[45] Wang, W. et al. reported an AI-driven virtual screening framework that links predicted apparent pKa and delivery efficiency to guide rational lipid design.[46] To address class imbalance in LNP datasets, Wu, K. et al. introduced a data-balanced transformer for accelerated ionizable-lipid/LNP screening.[47]

Generative and synthesis-aware lipid design

Generative models move beyond ranking by proposing candidate lipids directly. For translational use, however, generation must remain tied to synthetic feasibility, safety, and biodegradability rather than predicted transfection alone. Synthesis-aware approaches encode reaction templates, modular building blocks, or retrosynthetic graphs, shifting optimization from abstract structures to molecules that medicinal chemists can actually make.[50] This makes iterative design cycles experimentally actionable and supports multi-objective searches across potency, tolerability, stability, and manufacturability. A complementary *Nat. Mater.* strategy combined machine learning with combinatorial chemistry over a buildable lipid library, illustrating that useful search spaces in this field are often library-defined rather than unconstrained chemical universes.[43] VAEs, reinforcement learning, diffusion or graph-based generators, and Monte Carlo tree search can all be adapted to ionizable-lipid discovery.[51–53] Lipid-specific language models such as LipidBERT further connect proposal generation to delivery-specific scoring by pretraining on large virtual lipid corpora and transferring those representations to wet-lab screening tasks.[54] In practice, generative design is only useful when synthesis constraints, safety filters, and prospective validation are enforced from the outset.

Foundation models do not replace discriminative or generative models. Their near-term value is to provide transferable chemical priors that improve property prediction and to furnish richer latent spaces for constrained design when delivery labels are sparse.

Foundation models and self-driving laboratories

Foundation models are attractive because they can absorb general chemical regularities before seeing sparse delivery data. LUMI-lab exemplifies this shift.[37] Now published in Cell, the platform couples a transformer-based chemical foundation model with active learning and autonomous experimentation to discover ionizable lipids for mRNA delivery.[37] Lipid-focused language models such as LipidBERT suggest a related route through domain-specific pretraining on virtual lipid libraries, followed by transfer to downstream LNP property prediction.[54]

When paired with a closed-loop laboratory, LUMI-lab autonomously synthesized and screened more than 1,700 LNPs, identified brominated tails as a useful motif for pulmonary delivery, and reported 20.3% gene editing in lung epithelial cells after intratracheal administration of its lead formulation in mice.[37] The broader significance is practical: representation learning, uncertainty-aware experiment selection, and automation can be combined in one workflow to generate testable chemical hypotheses rather than only retrospective rankings.

Multimodal architectures and interpretability

To address the 'black box' nature of deep learning, newer architectures employ multi-modal fusion combining complementary molecular representations. TransMA integrates 3D geometric encodings (atomic coordinates and derived spatial features) with a sequence module to fuse conformational and SMILES information, and provides atom-level attribution to highlight determinants of activity cliffs in transfection datasets.[38]

FORMULATION-LEVEL OPTIMIZATION: LNPs AS COMPOSITE MATERIALS

At the formulation and process level, the design object is the drug product rather than the ionizable lipid alone. LNPs are soft-matter composites whose

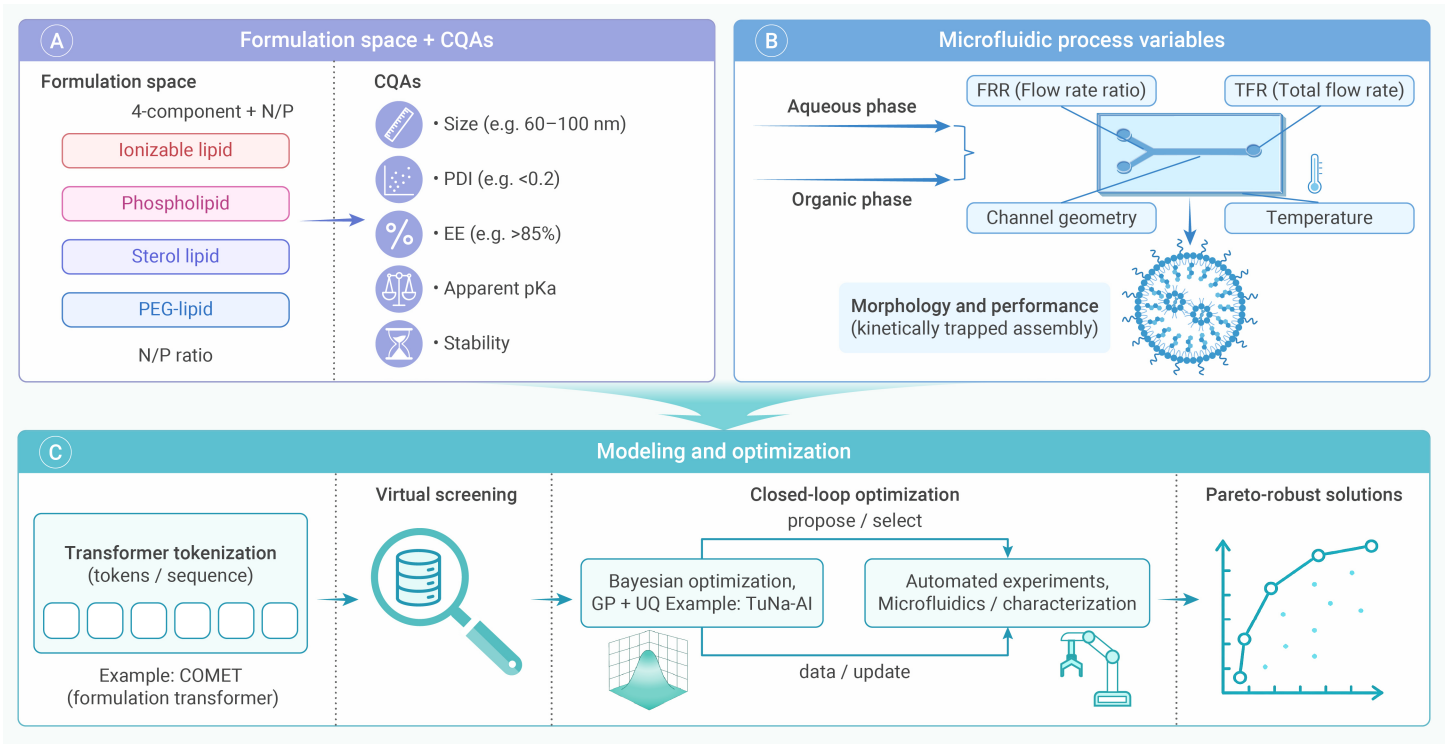

Figure 3. LNPs as composite materials: joint modeling of composition and process history (A) The formulation design space is represented by a four-component formulation plus nitrogen-to-phosphate (N/P) ratio, with component classes including ionizable lipid, phospholipid, sterol lipid (e.g., cholesterol), and PEG-lipid, and with representative critical quality attributes (CQAs) such as size, polydispersity index (PDI), encapsulation efficiency (EE), apparent pKa, and stability. Numeric targets are illustrative and depend on route, payload, and target product profile. (B) Microfluidic process variables, including phase composition, flow rate ratio (FRR), total flow rate (TFR), channel geometry, and temperature, shape kinetically trapped assembly, morphology, and performance. (C) Modeling frameworks integrate transformer-based tokenization, virtual screening, Gaussian process/Bayesian optimization with UQ, and automated microfluidic experiments to enable closed-loop formulation optimization and identification of Pareto-robust solutions.

internal morphology, surface chemistry, and batch reproducibility emerge from coupled choices: component identities (ionizable lipid, phospholipid, sterol, PEG-lipid), stoichiometries (typically 20–50:10–20:20–50:0.5–3 mol%), N/P ratio (nitrogen-to-phosphate charge ratio, typically 3–6), buffer composition and pH, and the kinetics of mixing and solvent exchange.[11,15] This is where delivery science meets device engineering, and where AI can directly reduce the gap between promising formulations and scalable, quality-controlled products. This composite-material view is schematized in Figure 3.

High-throughput formulation and screening platforms increasingly provide the data layer that makes these models useful. By generating larger, more diverse datasets with standardized process metadata, automation and parallelization can connect microfluidic library generation, robotic formulation, and rapid analytics to machine-learning-guided LNP optimization.[55–57] This strategy is already being used to improve mRNA vaccine formulations.[57]

Tokenizing composite formulations: Transformer architectures

A central conceptual shift is treating LNPs as composite materials rather than ingredient lists. The COMET (COmposite MatErial Transformer) platform exemplifies this direction by jointly encoding molecular structures, molar percentages, and process parameters in a unified transformer architecture.[36] Each formulation is represented as a sequence of component tokens (molecular graph embeddings concatenated with continuous stoichiometry values) processed through self-attention layers that learn pairwise and higher-order interactions. Because the model represents formulations as structured token sequences, it captures non-linear component interactions (synergies, antagonisms) that would be difficult to discover by factorial design or intuition.

COMET was trained on the LANCE (Lipid nanoparticle formulation ANalysis for mRNA dElivery) dataset comprising >3,000 unique formulations, including non-canonical systems with dual ionizable lipids, polymer–lipid hybrids, and unconventional helper lipid combinations. The model achieved Spearman ρ = 0.873 between predicted and experimentally observed transfection efficiencies on held-out test sets.[36] The multi-component architecture enabled virtual screening of nearly 50 million LNP combinations by systematically varying component identities and ratios. In prospective validation (subcutaneous administration), COMET-optimized LNPs achieved >40-fold and >5-fold higher bioluminescence intensity than the clinical benchmark formulations DLin-MC3-DMA and SM-102, respectively.[36] More broadly, transformer attention mechanisms and related attribution tools provide a route to interrogate which component identities/ratios drive predicted performance, helping prioritize candidate dual-ionizable lipid pairings and other compositional interactions for experimental testing.

Process history as a first-class variable

Microfluidic mixing work established that mixing timescales and solvent exchange kinetics control particle size, encapsulation efficiency, and reproducibility even at fixed composition.15 In rapid mixing (millisecond timescales), supersaturation drives kinetically trapped assembly, with particle characteristics depending on flow rate ratio (FRR) between the aqueous and organic phases, total flow rate (TFR), channel geometry (staggered herringbone, toroidal, simple T-junction), and temperature. Development programs therefore treat process parameters as design variables determining critical quality attributes (CQAs), with common product-dependent targets including size distribution (often in the ~60–100 nm range with low PDI), high encapsulation efficiency, residual-solvent control within applicable regulatory limits, and storage stability. AI models ingesting process metadata can map manufacturing conditions into the same optimization space as composition, enabling search for robust formulations insensitive to manufacturing variability.

Automated formulation optimization

TuNa-AI illustrates a complementary strategy: combining automated experimentation with machine learning to optimize both excipient selection and compositional ratios in a unified loop.[39] The platform employs Gaussian process regression with custom kernels capturing both discrete (component identity) and continuous (molar ratio) variables, coupled with robotic liquid handling for high-throughput formulation preparation. Although developed as a general nanoparticle framework, the approach translates directly to LNP optimization. TuNa-AI achieved a 42.9% increase in successful particle

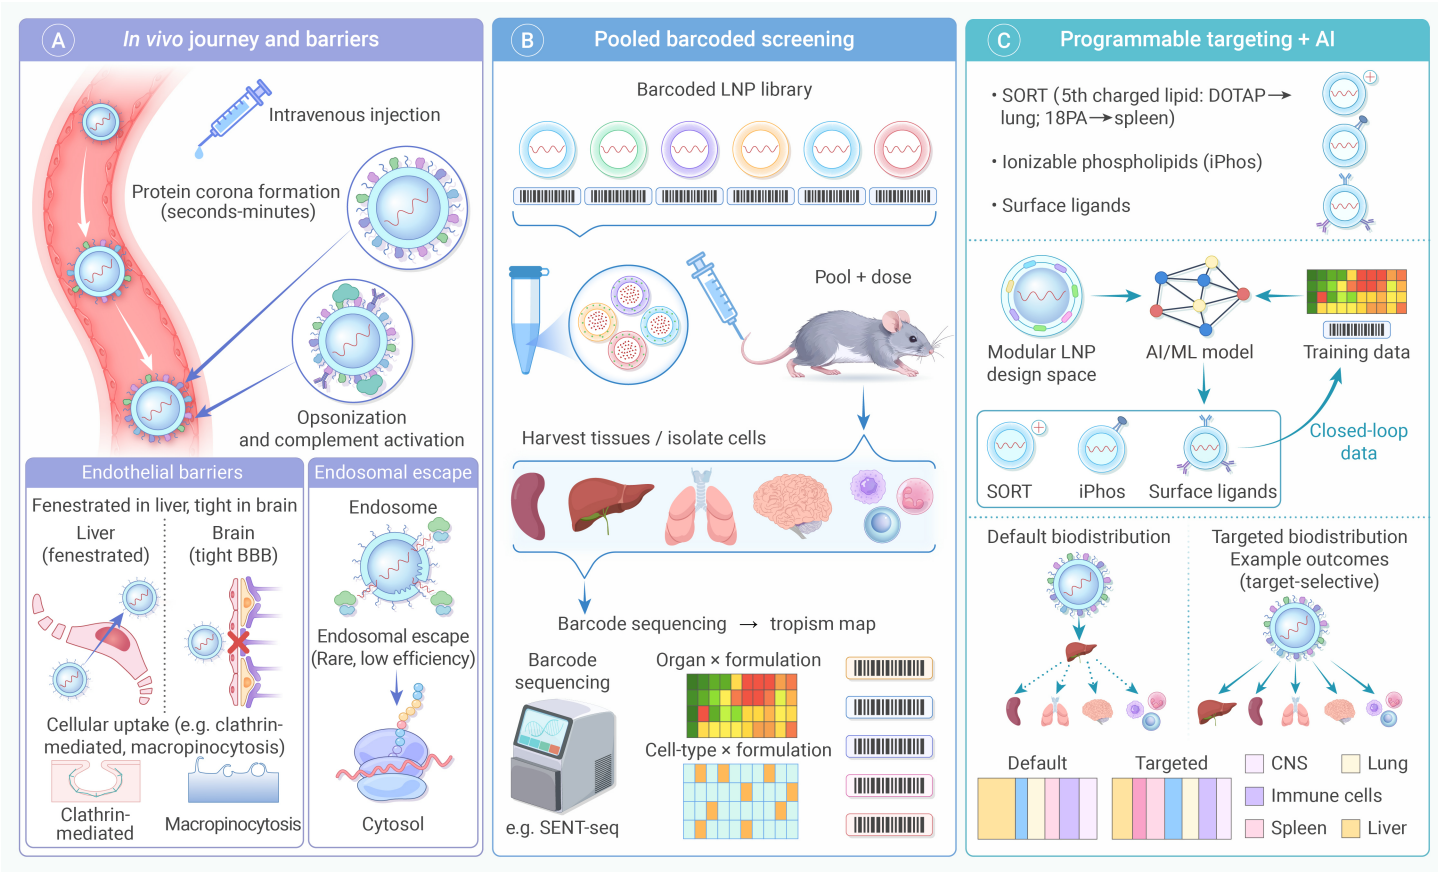

Figure 4. Learning organism-level targeting from pooled in vivo experiments (A) After intravenous administration, LNPs rapidly acquire a protein corona and encounter opsonization, endothelial barriers, cellular uptake pathways, and inefficient endosomal escape before productive cytosolic delivery. (B) In pooled barcoded screening, each formulation carries a unique nucleic-acid barcode; pooled dosing followed by tissue harvest and/or cell isolation plus sequencing generates organ-by-formulation and cell-type-by-formulation delivery maps. (C) Programmable targeting strategies, including selective organ targeting (SORT), ionizable phospholipids (iPhos), and surface ligands, expand the modular LNP design space beyond default biodistribution patterns and can be coupled with AI/ML models and closed-loop data generation to redirect delivery toward selected tissues and cell populations.

formation rate while simultaneously optimizing target properties, demonstrating that formulation is a constrained search problem where feasibility and robustness are as important as potency. For mRNA-LNPs, this aligns with Quality-by-Design (QbD) principles: AI can prioritize design space regions simultaneously satisfying CQAs and biological activity.

ORGANISM-LEVEL TARGETING AND BIODISTRIBUTION

At the organism scale, delivery performance emerges from coupled interactions between LNP composition, process history, and host biology. AI connects pooled in vivo libraries and single-cell readouts to organ and cell-type tropism, transforming biodistribution from an empirical outcome into a learnable design target (Figure 4).

From hepatic default to programmable organ tropism

Canonical clinical LNPs exhibit strong hepatic bias after intravenous dosing.[24–25] A biomolecular corona forms within seconds to minutes of blood contact,[26] and apolipoproteins within that corona, especially ApoE, can promote hepatocyte uptake through LDL receptor (LDLR)-linked pathways in the context of fenestrated hepatic sinusoidal endothelium (~100 nm fenestrae).[24] This default hepatic tropism has been advantageous for liver diseases (transthyretin amyloidosis, hypercholesterolemia) but acts as a confounder for AI models trained on historical data, which risk learning rules explaining liver delivery without generalizing to extrahepatic targets.

A central mechanistic insight is that tropism is an emergent property of the assembled composite and its route-dependent biological interfaces, not a fixed property of an ionizable lipid. The Selective Organ Targeting (SORT) approach demonstrated that adding a fifth charged lipid at defined molar fractions can redirect expression toward lung (cationic DOTAP) or spleen (anionic 18PA) without redesigning the mRNA or ionizable lipid.[19,58–59] SORT nanoparticles achieved lung-selective delivery through altered protein corona composition (reduced ApoE, increased vitronectin and fibrinogen) and uptake by pulmonary endothelium. Related strategies employing ionizable phospholipids, biodegradable linker chemistry, cholesterol analogs, and more radical reformulation of canonical four-component recipes further illustrate that whole-formulation context, not any single parameter like pKa or zeta potential, determines in vivo fate.[13,59–60] More recently, peptide–ionizable lipid nanoparticles achieved predictable tissue-selective delivery across lung, liver, spleen, thymus, and bone and enabled liver/lung prime editing, reinforcing that extrahepatic tropism can be engineered through encoded lipid chemistry and structure–selectivity rules rather than empirical component swapping alone.[61]

High-throughput in vivo mapping with barcoded libraries

The field's most consequential methodological shift has been pooled in vivo screening with nucleic-acid barcodes.[32,62–63] Each LNP formulation encapsulates a unique DNA or RNA barcode; after systemic administration, organ or cell-type-specific delivery is quantified by next-generation sequencing of harvested tissues. This transforms organ targeting from low-throughput single-formulation-per-animal studies into learnable mappings with thousands of formulation–tissue observations from minimal animal numbers.

Barcoding can be coupled to higher-resolution readouts (e.g., single-cell transcriptomics or other multi-omic measurements) to resolve delivery across cell types within a tissue and to train models that predict organ- and cell-type delivery from formulation descriptors.[32,62–63]

Immune cell targeting

Beyond hepatocyte delivery, targeting immune cells represents a critical

frontier for cancer immunotherapy, autoimmune disease modulation, and vaccine development. Combinatorial LNP libraries have been specifically optimized for leukocyte targeting.[64–65] Beyond bulk tropism, immune-cell selectivity can be achieved by tuning lipid head group chemistry and tail architecture, and by adding targeting moieties (e.g., antibodies or peptide ligands) to engage defined immune subsets.[66–67] For example, systemic RNA-lipoplexes can preferentially deliver mRNA to dendritic cells,[68] and mannose receptor-directed designs further enhance dendritic cell uptake.[69] Meanwhile, antibody-enabled targeting and engineered LNP formulations have enabled preferential delivery to T-cell subsets, supporting in vivo CAR-T engineering in preclinical models and early clinical feasibility studies.[66–67,70] Notably, screening in *Nat. Biomed. Eng.* optimized helper-lipid identity and component ratios to tune dendritic-cell maturation, antigen presentation, and the downstream balance of type-1/type-2 helper T-cell responses, showing that immune programme—not only cell entry—can be a direct optimization target.[71] Together, these results indicate that immune-cell delivery depends not only on organ tropism, but also on immune microenvironment navigation, activation-state-dependent uptake, and tissue-resident heterogeneity.

Machine learning has also been applied to optimize immunomodulatory LNP designs for defined therapeutic immune outcomes (beyond transfection or organ tropism), for example repolarizing hyperactivated microglia after mRNA delivery.[49]

Central nervous system targeting

Overcoming the blood–brain barrier (BBB) remains a grand challenge for nucleic acid delivery, with tight junctions, efflux pumps, and limited transcytosis restricting nanoparticle access to the central Nervous System (CNS). Recent work demonstrates AI-designed ionizable lipids with improved CNS accumulation and therapeutically relevant CNS expression.[72] In parallel, receptor-mediated transcytosis strategies (e.g., transferrin receptor/nicotinic acetylcholine receptor targeting) are being explored for systemic brain delivery.[73] Brain-targeting ligands such as rabies virus glycoprotein 29 (RVG29) have shown analogous promise in other nanoparticle systems, underscoring the broader utility of receptor- and peptide-mediated BBB engagement.[74] Focused ultrasound with microbubble-mediated BBB disruption can transiently enhance brain uptake of mRNA–LNPs.[75] Intranasal administration has been investigated as a non-invasive nose-to-brain route that may bypass some BBB constraints.[76]

In addition, Sela et al. reported AI-validated, brain-targeted mRNA-LNPs with neuronal tropism, highlighting progress from bulk biodistribution toward functional cell-type specificity in the CNS.[48]

Endosomal escape: The rate-limiting bottleneck

Subcellular trafficking constraints fundamentally limit AI model development. Mechanistic studies have characterized endosomal escape as a rare, stochastic event, with typically <10% of internalized LNP-derived cargo reaching the cytosol.[22–23,27] Most cargo remains trapped in endolysosomal compartments, undergoing acidification, lipid degradation, and eventual lysosomal destruction. The physicochemical determinants of productive escape involve membrane destabilization through lipid mixing, hexagonal phase formation, and transient pore creation at contact sites between LNPs and endosomal membranes, yet these processes remain mechanistically undercharacterized despite their central importance.

Investigations of cholesterol analogues (e.g., β-sitosterol, oxidized cholesterol derivatives) and helper lipid variants (DOPE vs. DSPC) indicate that modest changes in helper-lipid chemistry can reshape LNP morphology and intracellular trafficking, with downstream impacts on mRNA delivery efficiency.[13,77] Recent transition temperature-guided design studies further support the view that helper-lipid phase behavior is a key determinant of delivery performance.[78] For example, β-sitosterol-containing LNPs can adopt polymorphic shapes and enhance intracellular mRNA delivery, while oxidized cholesterol-containing formulations have been reported to modulate hepatic delivery profiles toward the liver microenvironment at clinically relevant doses. These findings underscore a fundamental challenge: the biological processes most critical for therapeutic efficacy occur at subcellular scales where high-throughput experimental readouts remain limited, creating a significant knowledge gap that constrains model training and validation.

Accordingly, the central unmet need is no longer recognition that endosomal escape is limiting, but scalable mechanistic measurements that can be linked to predictive models across formulations. We return to this challenge in Critical Gaps.

END-TO-END CO-DESIGN OF PAYLOAD, CARRIER, AND PROCESS

A central opportunity is moving beyond factorized design, where the carrier is improved while RNA is treated as interchangeable, or sequences are optimized assuming delivery context is fixed. In practice, payload, carrier, and process jointly determine particle state, innate immune sensing, and in vivo phenotype. The optimization target is co-optimizing payload–carrier–process triplets under efficacy, safety, and manufacturability constraints.

An important next step is a unified representation space in which RNA sequence, nucleotide modification state, ionizable-lipid structure, excipient composition, and process parameters are embedded jointly and optimized against multi-endpoint objectives. Such a framework would move beyond treating payload and carrier as separable modules, enabling optimization of payload-carrier-process triplets conditioned on route, tissue, and target product profile. In parallel, transfer-learning strategies may improve the translatability of immunogenicity models by pretraining on large murine delivery and cytokine datasets and then fine-tuning on smaller human immune-cell datasets or ex vivo readouts, with domain adaptation and uncertainty estimation used to flag mouse-specific features that are unlikely to transfer. This combination would better align abundant preclinical data with the limited but clinically decisive human immune measurements that determine tolerability and repeat-dose feasibility. Recent work on comprehensive co-engineering of ionizable LNPs and mRNA elements for next-generation vaccines underscores the practical value of this coupled optimization perspective.[79]

Generative RNA design

Generative RNA models have progressed from codon optimization to de novo sequence design that jointly optimizes coding sequences and regulatory elements (e.g., 5′/3′ UTRs) for expression, stability, and immunogenicity. Zhang et al. developed GEMORNA, a generative-explainable framework that optimizes mRNA sequences, achieving up to 41-fold higher luciferase expression in vivo versus standard codon-optimized baselines.[80] The approach was also extended to circular RNA (circRNA) design, where optimized constructs supported prolonged protein expression detectable out to 168 hours in vivo and improved antitumor cytotoxicity in CAR-T cells.[80] Together, these results illustrate that sequence-level optimization can benefit from context-aware generative models that account for local sequence–structure relationships beyond codon frequency heuristics.

Complementing GEMORNA, Tang, X. et al. developed a deep generative framework to design 5′ UTRs specifically for N1-methyl-pseudouridine-modified mRNAs, underscoring that payload co-design may need to be conditioned on nucleotide chemistry rather than treating untranslated regions as universally transferable.[81] Task-specific deep learning has likewise optimized 5′ UTRs for mRNA-delivered gene editing, indicating that sequence-design objectives may need to be conditioned on use case rather than generic expression alone.[82] Foundation-model approaches are further expanding this layer of payload engineering: UTR-LM learned translation-linked features from endogenous 5′ UTRs and prospectively generated designs with improved protein output, while mRNABERT extended pretraining to full-length mRNAs and improved performance across 5′ UTR, CDS, and whole-transcript benchmarks.[83–84]

Immunogenicity and safety as system-level objectives

Safety and immunogenicity emerge from interactions between RNA motifs (double-stranded regions, unmodified nucleosides, specific sequence patterns) and lipid composition (ionizable lipid structure, PEG-lipid content, formulation purity). Innate activation can be beneficial for vaccination (adjuvanting adaptive responses) but detrimental for repeat-dosed protein replacement or genome editing applications requiring immune evasion. This principle is especially evident in vaccine-focused studies, where adjuvanticity can be tuned by ionizable lipidoid substitution, PEG-lipid ratio, and phospholipid composition.[85–86]

Studies have shown that empty LNPs (no mRNA cargo) can elicit innate

immune activation, including TLR4-dependent NF-κB/IRF signaling, indicating that lipid composition itself can drive inflammatory sensing independent of RNA.[21,87–88] COVID-19 vaccination cohorts showed 13-fold increases in anti-PEG IgG antibodies, potentially compromising efficacy and safety of subsequent PEGylated therapeutics including LNP-based drugs.[89] Multi-task learning over cytokine panels, single-cell immune transcriptomes, and longitudinal antibody responses can reveal indication-specific Pareto frontiers balancing immunostimulation against tolerability, enabling rational navigation of the efficacy–safety trade-off space.

CRITICAL GAPS IN AI-ENABLED LNP DESIGN

Despite remarkable advances, fundamental gaps limit translational impact. These extend beyond algorithmic limitations to encompass data infrastructure, biological complexity, and ecosystems required to translate predictions into clinical reality. These gaps can be organized into three interdependent layers.

Data ecosystem gaps

The poor correlation between in vitro transfection and in vivo delivery represents a persistent bottleneck. Systematic evaluation of 281 LNP formulations demonstrated that cellular assay performance did not predict in vivo efficacy,[25,90] attributable to factors absent in simplified systems: dynamic protein corona formation and remodeling during circulation, hemodynamic shear forces affecting particle integrity, complement activation and opsonization, tissue-specific endothelial barriers, and cell-type-dependent uptake mechanisms. AI platforms trained predominantly on in vitro data inherit these limitations, achieving high predictive accuracy on test sets that may not translate to therapeutic relevance.

As comprehensively reviewed,[17] the field lacks standardized endpoints capturing therapeutic relevance: most studies optimize transfection efficiency while ignoring immunogenicity profiles, expression durability, biodistribution kinetics, and off-target effects. Furthermore, the absence of shared, curated datasets with comprehensive metadata (formulation composition, process parameters, assay conditions) precludes independent validation and cross-laboratory benchmarking.[91] Unlike computer vision or NLP, where benchmark datasets catalyzed rapid progress, LNP research still operates with fragmented and often proprietary data. Minimal metadata standards are therefore needed to support meaningful cross-study comparison and transfer learning.

A practical obstacle is that delivery-labeled datasets are rarely centralized. Large public chemical corpora (e.g., PubChem/ChEMBL/ZINC) and lipid-structure resources (e.g., LIPID MAPS) can support representation learning and pretraining, but delivery-relevant labels (formulation composition, process history, in vitro potency, in vivo biodistribution, and safety) are typically dispersed across individual publications, supplementary tables, and proprietary screening campaigns. As a result, even when data are available, they often require manual curation to standardize identifiers, units, and experimental covariates, which strengthens the case for community data commons and machine-readable reporting standards. Encouragingly, dedicated resources are beginning to emerge: LNP Atlas compiles standardized composition, process, physicochemical, and bioactivity metadata from published formulations, and LNPDB extends this logic toward a larger structure–function database intended for modeling and simulation.[92–93]

Beyond missing metadata, LNP datasets are frequently imbalanced (for example, sparse high-performing 'hits' amid many low-performers, and over-representation of specific routes, tissues, or lipid scaffolds), which can bias models toward dominant regimes and degrade out-of-distribution generalization. Data-balancing strategies within transformer architectures have been proposed to improve robustness under such imbalance.[47]

Mechanistic understanding gaps

The main problem is not recognizing endosomal escape as a bottleneck. It is the lack of scalable labels that let models distinguish formulations that truly release cargo from those that enter cells but stall in endolysosomal trafficking.

Current assays offer either limited throughput or limited mechanistic resolution, so escape remains poorly represented in training data. Molecular dynamics simulations could help connect lipid structure to membrane destabilization, lipid mixing, and water penetration at endosomal pH.[22–23] In practice, however, the simulation timescales required for meaningful membrane events and the difficulty of validating which observables transfer across formulations still limit their use as routine labels for model development.

Immunogenicity presents parallel challenges. Work characterizing how lipid chemistry influences innate immune engagement has identified structural features associated with TLR activation, complement binding, and cytokine induction.[21,89] Yet no current model can prospectively predict inflammatory outcomes or anti-PEG antibody responses for novel formulations. Species differences further complicate translation: murine immune signatures (cytokine profiles, innate receptor expression, complement activity) do not reliably predict human responses, creating fundamental uncertainty about model transferability across preclinical species.

Translation infrastructure gaps

The path from AI-generated predictions to regulatory-accepted evidence remains undefined. Recent perspectives on trustworthy AI for medicines provide an initial framework,[94] but LNP-specific applications remain to be elaborated. Key unresolved questions include: Under what circumstances can computational predictions substitute for experimental studies? What validation standards must AI-designed formulations meet beyond conventionally developed products? How should model updates and retraining be managed within GMP quality systems? What documentation demonstrates that AI-recommended formulations satisfy CMC requirements?

To make this gap more actionable, Table 3 outlines a minimal best-practice checklist for AI-in-LNP studies. At minimum, studies should report machine-readable metadata for cargo, lipid chemistry, excipients, and process parameters; predefine endpoints and evaluation protocols for the intended route, tissue, species, and target product profile; use validation schemes that test scaffold, process, or screen transfer rather than random splits within a narrow series; report uncertainty or out-of-distribution safeguards; and prospectively validate top candidates under representative CMC conditions.

Manufacturability constraints are frequently overlooked in academic optimization. LNP drug products must satisfy product-dependent CQAs, commonly including size distribution (often in the ~60–100 nm range with low PDI), high encapsulation efficiency, residual-solvent levels within applicable regulatory limits,[95] storage stability under intended storage conditions, and batch-to-batch consistency across scale-up from mL to hundreds of liters. The experience scaling mRNA vaccine production during COVID-19 demonstrated both feasibility and challenges of large-scale LNP manufacturing.[96] Digital twins for unit operations (physics-based models of microfluidic mixers, tangential flow filtration systems, and fill-finish processes) are emerging to accelerate process development, but integration with composition-focused AI remains limited. Ensuring that AI-discovered formulations are not only potent but also producible under real-world manufacturing conditions requires explicitly incorporating manufacturability constraints into optimization objectives.

CLINICAL TRANSLATION: ADVANTAGES, LIMITATIONS, AND FUTURE PERSPECTIVES

AI-guided LNP design offers three practical advantages for translation. It can shorten the design-build-test cycle by prioritizing lipid scaffolds and formulation or process windows under joint potency, safety, stability, and manufacturability constraints. It can reuse information across structure, process, and in vivo readouts to improve data efficiency, especially when paired with uncertainty-aware active learning. It can also optimize composition and process together so that CMC constraints enter early rather than being retrofitted at late-stage development.

The main challenges are dataset shift between discovery assays, scale-up manufacturing, and human biology; weak mechanistic labels that encourage optimization on surrogate endpoints; the degree of interpretability and auditability needed for regulatory review; limited reproducibility when models rely on proprietary datasets; and the absence of shared validation protocols. For that reason, uncertainty quantification, out-of-distribution detection, and prospective validation under representative CMC conditions are not

Table 3. Minimal metadata for reproducible, cross-study benchmarking of AI-designed mRNA–LNP systems

Layer	Minimal metadata	Rationale for AI and translation
Cargo	Sequence, cap analog, modified nucleosides (type and fraction), length, purity (% intact), and dose	Payload sequence and chemistry shape translation, innate sensing, dose normalization, and cross-study comparability.
Ionizable lipid	SMILES, measured pKa, biodegradability motifs, synthesis purity	Structure, pKa, purity, and biodegradability influence encapsulation, endosomal escape, potency, and tolerability.
Excipients	Phospholipid, sterol, PEG-lipid identities; exact mol%; N/P ratio	Helper lipids and molar ratios modulate morphology, protein corona, stability, and tissue tropism.
Process	Mixing device/geometry, solvent/aqueous composition, FRR, TFR, temperature, buffer pH	Mixing and solvent-exchange conditions determine CQAs, batch reproducibility, and scale-up transferability.
In vivo	Species/strain/sex/age, route, dose, dosing schedule, timepoints, quantification method	Route, species, dose, and timepoint define external validity and are essential for benchmarking translational relevance.

Abbreviations: CMC, Chemistry, Manufacturing, and Controls; CQAs, critical quality attributes.

optional.[94]

Several developments would make AI-designed LNPs more clinically useful: benchmark datasets with machine-readable metadata and negative results (Table 3), synthesis-aware generative models constrained by plausible routes and degradability, integration of barcoded in vivo tropism maps with mechanistic biomarkers such as endosomal escape or innate sensing, and hybrid mechanistic-ML frameworks for process scale-up. In this setting, AI is best viewed as a documented decision-support tool within a GxP-aligned development program, not as a substitute for experimental evidence.

CONCLUSION

mRNA-lipid nanoparticle therapeutics have moved from proof-of-concept delivery experiments to a clinically validated platform, but performance is still set by material and process variables within the LNP envelope. The next phase will depend on treating delivery as a cross-scale design problem in which lipid chemistry, whole-formulation composition, manufacturing history, and payload properties are optimized together. AI is useful here because it can learn from multi-endpoint datasets and search spaces that are too large for intuition-driven iteration.

Its strongest role is likely to be in workflows that combine pretraining, mechanistically informative assays, uncertainty-aware generalization, and closed-loop experimentation under realistic manufacturing constraints. Progress will therefore hinge on better data ecosystems, better measurements of subcellular bottlenecks, and clearer translation infrastructure. If those pieces mature, the field can move from optimizing isolated lipids or sequences to optimizing payload-carrier-process triplets that are effective, durable, and manufacturable at scale.

REFERENCES

1. Karikó K. et al. (2005). Suppression of RNA recognition by Toll-like receptors: the impact of nucleoside modification and the evolutionary origin of RNA. *Immunity* **23**:165–175. DOI:10.1016/j.immuni.2005.06.008
2. Karikó K. et al. (2008). Incorporation of pseudouridine into mRNA yields superior nonimmunogenic vector with increased translational capacity and biological stability. *Mol. Ther.* **16**:1833–1840. DOI:10.1038/mt.2008.200
3. Cullis P.R. and Hope M.J. (2017). Lipid nanoparticle systems for enabling gene therapies. *Mol. Ther.* **25**:1467–1475. DOI:10.1016/j.ymthe.2017.03.013
4. Semple S.C. et al. (2010). Rational design of cationic lipids for siRNA delivery. *Nat. Biotechnol.* **28**:172–176. DOI:10.1038/nbt.1602
5. Akinc A. et al. (2019). The Onpattro story and the clinical translation of nanomedicines. *Nat. Nanotechnol.* **14**:1084–1087. DOI:10.1038/s41565-019-0591-y
6. Adams D. et al. (2018). Patisiran, an RNAi Therapeutic, for Hereditary Transthyretin Amyloidosis. *N. Engl. J. Med.* **379**:11–21. DOI:10.1056/NEJMoa1716153
7. Polack F.P. et al. (2020). Safety and Efficacy of the BNT162b2 mRNA Covid-19 Vaccine. *N. Engl. J. Med.* **383**:2603–2615. DOI:10.1056/NEJMoa2034577
8. Baden L.R. et al. (2021). Efficacy and Safety of the mRNA-1273 SARS-CoV-2 Vaccine. *N. Engl. J. Med.* **384**:403–416. DOI:10.1056/NEJMoa2035389
9. Pardi N. et al. (2018). mRNA vaccines—a new era in vaccinology. *Nat. Rev. Drug Discov.* **17**:261–279. DOI:10.1038/nrd.2017.243
10. Alameh M.-G. et al. (2021). Lipid nanoparticles enhance the efficacy of mRNA and protein subunit vaccines by inducing robust T follicular helper cell and humoral responses. *Immunity* **54**:2877–2892.e7. DOI:10.1016/j.immuni.2021.11.001
11. Hou X. et al. (2021). Lipid nanoparticles for mRNA delivery. *Nat. Rev. Mater.* **6**:1078–1094. DOI:10.1038/s41578-021-00358-0
12. Chaudhary N., Weissman D. and Whitehead K.A. (2021). mRNA vaccines for infectious diseases: principles, delivery and clinical translation. *Nat. Rev. Drug Discov.* **20**:817–838. DOI:10.1038/s41573-021-00283-5
13. Patel S. et al. (2020). Naturally-occurring cholesterol analogues in lipid nanoparticles induce polymorphic shape and enhance intracellular delivery of mRNA. *Nat. Commun.* **11**:983. DOI:10.1038/s41467-020-14527-2
14. Ball R.L. et al. (2018). Lipid nanoparticle formulations for enhanced co-delivery of siRNA and mRNA. *Nano Lett.* **18**:3814–3822. DOI:10.1021/acs.nanolett.8b01101
15. Maeki M. et al. (2018). Advances in microfluidics for lipid nanoparticles and extracellular vesicles and applications in drug delivery systems. *Adv. Drug Deliv. Rev.* **128**:84–100. DOI:10.1016/j.addr.2018.03.008
16. Sahin U., Karikó K. and Türeci Ö. (2014). mRNA-based therapeutics—developing a new class of drugs. *Nat. Rev. Drug Discov.* **13**:759–780. DOI:10.1038/nrd4278
17. Hajj K.A. and Whitehead K.A. (2017). Tools for translation: non-viral materials for therapeutic mRNA delivery. *Nat. Rev. Mater.* **2**:17056. DOI:10.1038/natrevmats.2017.56
18. Gillmore J.D. et al. (2021). CRISPR-Cas9 In Vivo Gene Editing for Transthyretin Amyloidosis. *N. Engl. J. Med.* **385**:493–502. DOI:10.1056/NEJMoa2107454
19. Cheng Q. et al. (2020). Selective organ targeting (SORT) nanoparticles for tissue-specific mRNA delivery and CRISPR–Cas gene editing. *Nat. Nanotechnol.* **15**:313–320. DOI:10.1038/s41565-020-0669-6
20. Han X. et al. (2024). Fast and facile synthesis of amidine-incorporated degradable lipids for versatile mRNA delivery in vivo. *Nat. Chem.* **16**:1687–1697. DOI:10.1038/s41557-024-01557-2
21. Ndeupen S. et al. (2021). The mRNA-LNP platform's lipid nanoparticle component used in preclinical vaccine studies is highly inflammatory. *iScience* **24**:103479. DOI:10.1016/j.isci.2021.103479
22. Gilleron J. et al. (2013). Image-based analysis of lipid nanoparticle-mediated siRNA delivery, intracellular trafficking and endosomal escape. *Nat. Biotechnol.* **31**:638–646. DOI:10.1038/nbt.2612
23. Sahay G. et al. (2013). Efficiency of siRNA delivery by lipid nanoparticles is limited by endocytic recycling. *Nat. Biotechnol.* **31**:653–658. DOI:10.1038/nbt.2614
24. Akinc A. et al. (2010). Targeted delivery of RNAi therapeutics with endogenous and exogenous ligand-based mechanisms. *Mol. Ther.* **18**:1357–1364. DOI:10.1038/mt.2010.85
25. Mitchell M.J. et al. (2021). Engineering precision nanoparticles for drug delivery. *Nat. Rev. Drug Discov.* **20**:101–124. DOI:10.1038/s41573-020-0090-8
26. Monopoli M.P. et al. (2012). Biomolecular coronas provide the biological identity of nanosized materials. *Nat. Nanotechnol.* **7**:779–786. DOI:10.1038/nnano.2012.207
27. Wittrup A. et al. (2015). Visualizing lipid-formulated siRNA release from endosomes and target gene knockdown. *Nat. Biotechnol.* **33**:870–876. DOI:10.1038/nbt.3298
28. Jayaraman M. et al. (2012). Maximizing the potency of siRNA lipid nanoparticles for hepatic gene silencing in vivo. *Angew. Chem. Int. Ed.* **51**:8529–8533. DOI:10.1002/anie.201203263
29. Whitehead K.A. et al. (2014). Degradable lipid nanoparticles with predictable in vivo siRNA delivery activity. *Nat. Commun.* **5**:4277. DOI:10.1038/ncomms5277
30. Dong Y. et al. (2014). Lipopeptide nanoparticles for potent and selective siRNA delivery in rodents and nonhuman primates. *Proc. Natl. Acad. Sci. U.S.A.* **111**:3955–3960. DOI:10.1073/pnas.1322937111
31. Kauffman K.J. et al. (2015). Optimization of lipid nanoparticle formulations for mRNA delivery in vivo with definitive screening designs. *Nano Lett.* **15**:7300–7306. DOI:10.1021/acs.nanolett.5b02497
32. Dahlman J.E. et al. (2017). Barcoded nanoparticles for high throughput in vivo discovery of targeted therapeutics. *Proc. Natl. Acad. Sci. U.S.A.* **114**:2060–2065. DOI:10.1073/pnas.1620874114
33. Hamilton A.G. et al. (2026). High-throughput in vivo screening using barcoded mRNA identifies lipid nanoparticles with extrahepatic tropism for in situ immunoengineering. *Adv. Mater.* **38**:e14370. DOI:10.1002/adma.202514370
34. Xu Y. et al. (2024). AGILE platform: a deep learning powered approach to accelerate

LNP development for mRNA delivery. *Nat. Commun.* **15**:6305. DOI:10.1038/s41467-024-50619-z
35. Witten J. et al. (2025). Artificial intelligence-guided design of lipid nanoparticles for pulmonary gene therapy. *Nat. Biotechnol.* **43**:1790–1799. DOI:10.1038/s41587-024-02490-y
36. Chan A. et al. (2025). Designing lipid nanoparticles using a transformer-based neural network. *Nat. Nanotechnol.* **20**:1491–1501. DOI:10.1038/s41565-025-01975-4
37. Xu Y. et al. (2026). LUMI-lab: A foundation model-driven autonomous platform enabling discovery of ionizable lipid designs for mRNA delivery. *Cell* **189**:1620–1635.e25. DOI:10.1016/j.cell.2026.01.012
38. Wu K. et al. (2025). TransMA: an explainable multi-modal deep learning model for predicting properties of ionizable lipid nanoparticles in mRNA delivery. *Brief. Bioinform.* **26**:bbaf307. DOI:10.1093/bib/bbaf307
39. Zhang Z. et al. (2025). TuNa-AI: a hybrid kernel machine to design tunable nanoparticles for drug delivery. *ACS Nano* **19**:33288–33296. DOI:10.1021/acsnano.5c09066
40. Deng J., Yang Z., Wang H., et al. (2023). A systematic study of key elements underlying molecular property prediction. *Nat. Commun.* **14**:6395. DOI:10.1038/s41467-023-41948-6
41. Butler K.T., Davies D.W., Cartwright H., et al. (2018). Machine learning for molecular and materials science. *Nature* **559**:547–555. DOI:10.1038/s41586-018-0337-2
42. Wang Y. et al. (2022). Molecular contrastive learning of representations via graph neural networks. *Nat. Mach. Intell.* **4**:279–287. DOI:10.1038/s42256-022-00447-x
43. Li B. et al. (2024). Accelerating ionizable lipid discovery for mRNA delivery using machine learning and combinatorial chemistry. *Nat. Mater.* **23**:1002–1008. DOI:10.1038/s41563-024-01867-3
44. Han X. et al. (2024). In situ combinatorial synthesis of degradable branched lipidoids for systemic delivery of mRNA therapeutics and gene editors. *Nat. Commun.* **15**:1762. DOI:10.1038/s41467-024-45537-z
45. Han X. et al. (2025). Plug-and-play assembly of biodegradable ionizable lipids for potent mRNA delivery and gene editing in vivo. bioRxiv preprint 2025.02.25.640222. DOI:10.1101/2025.02.25.640222.
46. Wang W. et al. (2024). Artificial intelligence-driven rational design of ionizable lipids for mRNA delivery. *Nat. Commun.* **15**:10804. DOI:10.1038/s41467-024-55072-6
47. Wu K. et al. (2024). Data-balanced transformer for accelerated ionizable lipid nanoparticles screening in mRNA delivery. *Brief. Bioinform.* **25**:bbae186. DOI:10.1093/bib/bbae186
48. Sela M. et al. (2025). AI-validated brain-targeted mRNA lipid nanoparticles with neuronal tropism. *ACS Nano* **19**:36106–36128. DOI:10.1021/acsnano.4c15013
49. Rafiei M., Shojaei A. and Chau Y. (2025). Machine learning-assisted design of immunomodulatory lipid nanoparticles for delivery of mRNA to repolarize hyperactivated microglia. *Drug Deliv.* **32**:2465909. DOI:10.1080/10717544.2025.2465909
50. Segler M.H.S. et al. (2018). Planning chemical syntheses with deep neural networks and symbolic AI. *Nature* **555**:604–610. DOI:10.1038/nature25978
51. Yang X. et al. (2017). ChemTS: an efficient Python library for de novo molecular generation. *Sci. Technol. Adv. Mater.* **18**:972–976. DOI:10.1080/14686996.2017.1401424
52. Jensen J.H. (2019). A graph-based genetic algorithm and generative model/Monte Carlo tree search for the exploration of chemical space. *Chem. Sci.* **10**:3567–3572. DOI:10.1039/C8SC05372C
53. Iwata H. et al. (2023). VGAE-MCTS: A New Molecular Generative Model Combining the Variational Graph Auto-Encoder and Monte Carlo Tree Search. *J. Chem. Inf. Model.* **63**:7392–7400. DOI:10.1021/acs.jcim.3c01220
54. Yu T. et al. (2024). LipidBERT: A Lipid Language Model Pre-trained on METiS de novo Lipid Library. arXiv preprint arXiv:2408.06150. DOI:10.48550/arXiv.2408.06150.
55. Zeng J. et al. (2026). High-throughput, AI-assisted design and optimization of lipid nanoparticles for drug delivery. *J. Control. Release* **391**:114573. DOI:10.1016/j.jconrel.2025.114573
56. Hanna A.R., Issadore D.A. and Mitchell M.J. (2026). High-throughput platforms for machine learning-guided lipid nanoparticle design. *Nat. Rev. Mater.* **11**:50–64. DOI:10.1038/s41578-025-00831-0
57. Bae S.-H. et al. (2025). Rational Design of Lipid Nanoparticles for Enhanced mRNA Vaccine Delivery via Machine Learning. *Small* **21**:e2405618. DOI:10.1002/smll.202405618
58. Dilliard S.A. et al. (2021). On the mechanism of tissue-specific mRNA delivery by selective organ targeting nanoparticles. *Proc. Natl. Acad. Sci. U.S.A.* **118**:e2109256118. DOI:10.1073/pnas.2109256118
59. Liu S. et al. (2021). Membrane-destabilizing ionizable phospholipids for organ-selective mRNA delivery and CRISPR–Cas gene editing. *Nat. Mater.* **20**:701–710. DOI:10.1038/s41563-020-00886-0
60. Su K. et al. (2024). Reformulating lipid nanoparticles for organ-targeted mRNA accumulation and translation. *Nat. Commun.* **15**:5659. DOI:10.1038/s41467-024-50093-7
61. Lin Y. et al. (2026). Tissue-specific mRNA delivery and prime editing with peptide–ionizable lipid nanoparticles. *Nat. Mater.* **25**:133–145. DOI:10.1038/s41563-025-02320-9
62. Sago C.D. et al. (2018). High-throughput in vivo screen of functional mRNA delivery identifies nanoparticles for endothelial cell gene editing. *Proc. Natl. Acad. Sci. U.S.A.* **115**:E9944–E9952. DOI:10.1073/pnas.1811276115
63. Dobrowolski C. et al. (2022). Nanoparticle single-cell multiomic readouts reveal that cell heterogeneity influences lipid nanoparticle-mediated messenger RNA delivery. *Nat. Nanotechnol.* **17**:871–879. DOI:10.1038/s41565-022-01146-9
64. Peer D. et al. (2008). Systemic leukocyte-directed siRNA delivery revealing cyclin D1 as an anti-inflammatory target. *Science* **319**:627–630. DOI:10.1126/science.1149859
65. Ni H. et al. (2022). Piperazine-derived lipid nanoparticles deliver mRNA to immune cells in vivo. *Nat. Commun.* **13**:4766. DOI:10.1038/s41467-022-32281-5
66. Chen M.Z. et al. (2025). A versatile antibody capture system drives specific in vivo delivery of mRNA-loaded lipid nanoparticles. *Nat. Nanotechnol.* **20**:1273–1284. DOI:10.1038/s41565-025-01954-9
67. Hunter T.L. et al. (2025). In vivo CAR T cell generation to treat cancer and autoimmune disease. *Science* **388**:1311–1317. DOI:10.1126/science.ads8473
68. Kranz L.M. et al. (2016). Systemic RNA delivery to dendritic cells exploits antiviral defence for cancer immunotherapy. *Nature* **534**:396–401. DOI:10.1038/nature18300
69. Lei J. et al. (2024). Development of Mannosylated Lipid Nanoparticles for mRNA Cancer Vaccine with High Antigen Presentation Efficiency and Immunomodulatory Capability. *Angew. Chem. Int. Ed.* **63**:e202318515. DOI:10.1002/anie.202318515
70. Wang Q. et al. (2025). In vivo CD19 CAR T-cell therapy for refractory systemic lupus erythematosus. *N. Engl. J. Med.* **393**:1542–1544. DOI:10.1056/NEJMc2509522
71. Zhu Y. et al. (2024). Screening for lipid nanoparticles that modulate the immune activity of helper T cells towards enhanced antitumour activity. *Nat. Biomed. Eng.* **8**:544–560. DOI:10.1038/s41551-023-01131-0
72. Cao D. et al. (2025). Lipid nanoparticles for mRNA delivery in brain via systemic administration. *Sci. Adv.* **11**:eadw0730. DOI:10.1126/sciadv.adw0730
73. Ghosh S. et al. (2026). Multiplex screening of transferrin receptor-targeted lipid nanoparticles for systemic brain delivery using DNA barcodes. *Cell Biomaterials* **2**:100303. DOI:10.1016/j.celbio.2025.100303
74. You L. et al. (2018). Targeted brain delivery of rabies virus glycoprotein 29-modified deferoxamine-loaded nanoparticles reverses functional deficits in Parkinsonian mice. *ACS Nano* **12**:4123–4139. DOI:10.1021/acsnano.7b08172
75. Ogawa K. et al. (2022). Focused ultrasound/microbubbles-assisted BBB opening enhances LNP-mediated mRNA delivery to brain. *J. Control. Release* **348**:34–41. DOI:10.1016/j.jconrel.2022.05.042
76. Xu K. et al. (2024). Nose-to-brain delivery of nanotherapeutics: transport mechanisms and applications. *Wiley Interdiscip. Rev. Nanomed Nanobiotechnol.* **16**:e1956. DOI:10.1002/wnan.1956
77. Paunovska K. et al. (2019). Nanoparticles containing oxidized cholesterol deliver mRNA to the liver microenvironment at clinically relevant doses. *Adv. Mater.* **31**:e1807748. DOI:10.1002/adma.201807748
78. Shin J.E. et al. (2025). Transition temperature-guided design of lipid nanoparticles for effective mRNA delivery. *ACS Appl. Mater. Interfaces* **17**:28012–28024. DOI:10.1021/acsami.5c06464
79. Kim J.G. et al. (2026). Comprehensive engineering of ionizable lipid nanoparticles and mRNA elements for next-generation vaccines. *ACS Nano* **20**:7022–7045. DOI:10.1021/acsnano.5c19036
80. Zhang H. et al. (2025). Deep generative models design mRNA sequences with enhanced translational capacity and stability. *Science* **390**:eadr8470. DOI:10.1126/science.adr8470
81. Tang X. et al. (2024). A novel deep generative model for mRNA vaccine development: Designing 5′ UTRs with N1-methyl-pseudouridine modification. *Acta Pharm. Sin. B* **14**:1814–1826. DOI:10.1016/j.apsb.2023.11.003
82. Castillo-Hair S. et al. (2024). Optimizing 5′UTRs for mRNA-delivered gene editing using deep learning. *Nat. Commun.* **15**:5284. DOI:10.1038/s41467-024-49508-2
83. Chu Y. et al. (2024). A 5′ UTR language model for decoding untranslated regions of mRNA and function predictions. *Nat. Mach. Intell.* **6**:449–460. DOI:10.1038/s42256-024-00823-9
84. Xiong Y. et al. (2025). mRNABERT: advancing mRNA sequence design with a universal language model and comprehensive dataset. *Nat. Commun.* **16**:10371. DOI:10.1038/s41467-025-65340-8
85. Vadovics M. et al. (2025). Tailoring the adjuvanticity of lipid nanoparticles by PEG lipid ratio and phospholipid modifications. *Nat. Nanotechnol.* **20**:1312–1322. DOI:10.1038/s41565-025-01958-5
86. Han X. et al. (2023). Adjuvant lipidoid-substituted lipid nanoparticles augment the immunogenicity of SARS-CoV-2 mRNA vaccines. *Nat. Nanotechnol.* **18**:1105–1114. DOI:10.1038/s41565-023-01404-4
87. Chaudhary N. et al. (2024). Amine headgroups in ionizable lipids drive immune responses to lipid nanoparticles by binding to the receptors TLR4 and CD1d. *Nat. Biomed. Eng.* **8**:1483–1498. DOI:10.1038/s41551-024-01256-w
88. Zelkoski A.E. et al. (2025). Ionizable lipid nanoparticles of mRNA vaccines elicit NF-κB and IRF responses through toll-like receptor 4. *npj Vaccines* **10**:73. DOI:10.1038/s41541-025-01124-x
89. Ju Y. et al. (2022). Anti-PEG antibodies boosted in humans by SARS-CoV-2 lipid nanoparticle mRNA vaccine. *ACS Nano* **16**:11769–11780. DOI:10.1021/acsnano.2c04543
90. Paunovska K. et al. (2018). A direct comparison of in vitro and in vivo nucleic acid delivery reveals a weak correlation. *Nano Lett.* **18**:2148–2157. DOI:10.1021/acs.

nanolett.8b00432
91. Faria M. et al. (2018). Minimum information reporting in bio–nano experimental literature. *Nat. Nanotechnol.* **13**:777–785. DOI:10.1038/s41565-018-0246-4
92. Song S., Baek J. and Seo S. (2026). A Comprehensive Dataset of Lipid Nanoparticle Compositions and Properties for Nucleic Acid Delivery. *Sci. Data* **13**:139. DOI:10.1038/s41597-025-06456-w
93. Collins E. et al. (2026). Lipid Nanoparticle Database towards structure-function modeling and data-driven design for nucleic acid delivery. *Nat. Commun.* **17**:2464. DOI:10.1038/s41467-026-68818-1
94. Stegmann J.U. et al. (2023). Trustworthy AI for safe medicines. *Nat. Rev. Drug Discov.* **22**:855–856. DOI:10.1038/s41573-023-00769-4
95. Dikpati A. et al. (2020). Residual Solvents in Nanomedicine and Lipid-Based Drug Delivery Systems: a Case Study to Better Understand Processes. *Pharm. Res.* **37**:149. DOI:10.1007/s11095-020-02877-x
96. Shepherd S.J. et al. (2023). Throughput-scalable manufacturing of SARS-CoV-2 mRNA lipid nanoparticle vaccines. *Proc. Natl. Acad. Sci. U.S.A.* **120**:e2303567120. DOI:10.1073/pnas.2303567120

FUNDING AND ACKNOWLEDGMENTS

Not applicable.

AUTHOR CONTRIBUTIONS

H. Z. and J. X. conceived the review, performed the literature analysis, and drafted the manuscript. X. G., J. Z. and Z. H. critically revised the manuscript. All authors contributed to the manuscript and approved the final version.

DECLARATION OF INTERESTS

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

DATA AND CODE AVAILABILITY

Not applicable.